

Instantons and Rational Homology Spheres

Aliakbar Daemi, Mike Miller Eismeier

Abstract

In previous work, the second author defined ‘equivariant instanton homology groups’ $I^\bullet(Y, \pi; R)$ for a rational homology 3-sphere Y , a set of auxiliary data π , and a PID R . These objects are modules over the cohomology ring $H^{-*}(BSO_3; R)$.

We prove that the equivariant instanton homology groups $I^\bullet(Y; R)$ are independent of the auxiliary data π , and thus define topological invariants of rational homology spheres. Further, we prove that these invariants are functorial under cobordisms of 3-manifolds with a path between the boundary components.

For any rational homology sphere Y , we may also define an analogue of Floer’s irreducible instanton homology group of integer homology spheres $I_*(Y, \pi; R)$ which now depends on the auxiliary data π , unlike the equivariant instanton homology groups. However, our methods allow us to prove a precise “wall-crossing formula” for $I_*(Y, \pi; R)$ as the auxiliary data π moves between adjacent chambers. We use this to define an instanton invariant $\lambda_I(Y) \in \mathbb{Q}$ of rational homology spheres, conjecturally equal to the Casson–Walker invariant.

Our approach to invariance uses a novel technique known as a blown-up flow category. Given an obstructed cobordism $W : Y \rightarrow Y'$, which supports reducible instantons which can neither be cut out transversely nor be removed by a small change to the perturbation, we remove and replace a neighborhood of obstructed solutions in the moduli space of instantons. The resulting moduli spaces have a new type of boundary component, so do not define a chain map between the instanton chain complexes of Y and Y' . However, it does define a chain map between the instanton chain complex of Y and a sort of suspension of the instanton chain complex of Y' .

Contents

1	Introduction	3
2	An overview of the article	8
3	Moduli spaces of instantons	17
3.1	Linear analysis on 3-manifolds	17
3.2	The ASD equations on 4-manifolds	21
3.3	The case of cylinders	25
3.4	Compactifications of the moduli spaces	26

3.5	Reducibles and regularity	29
3.6	Families of metrics	35
3.7	Orientations	36
4	Geometric chains and instanton homology	40
4.1	Stratified-smooth spaces and the geometric chain complex	40
4.2	Flow categories and bimodules	43
4.3	The instanton chain complex and unobstructed cobordism maps	50
5	Obstructed cobordism maps	53
5.1	Obstructed cobordisms and obstructed gluing theory	53
5.2	Modified moduli spaces	59
5.3	The blown-up flow category	65
5.4	The bimodule of an obstructed cobordism	72
5.5	Invariance of instanton homology	77
5.5.1	Modified moduli spaces for homotopies	80
5.5.2	Homotopies of bimodules from modified moduli spaces	86
6	Functoriality of equivariant instanton homology	92
6.1	Cobordism maps in the pseudo-unobstructed case	92
6.2	Composites	99
6.3	The case $b_2^+ > 0$ via the blowup trick	104
7	Irreducible Floer homology and the Casson-Walker invariant	108
7.1	Floer’s complex from algebra	108
7.2	An instanton Casson–Walker invariant	112
	Appendix	116
A	Stratified-smooth spaces	116
A.1	Foundations of stratified-smooth spaces	116
A.2	Orientations on stratified-smooth spaces	124

A.3	The real blowup of a stratified-smooth space	131
A.4	Geometric homology from stratified-smooth spaces	137
A.5	Truncating the geometric chain complex	143

1 Introduction

Instanton homology originates in Floer’s landmark paper [Flo88], where he associates to a 3-manifold Y with the same integral homology groups as the 3-sphere (an integer homology 3-sphere) a finite-dimensional $\mathbb{Z}/8$ -graded abelian group, $I_*(Y)$. This group has many desiderata: its Euler characteristic returns twice the Casson invariant [Tau90], while a four-manifold X with boundary Y gives elements $D_X \in I_*(Y)$ generalizing Donaldson’s polynomial invariants for closed 4-manifolds; if a closed 4-manifold X can be decomposed as $X = X_1 \cup_Y X_2$ along an integer homology 3-sphere Y , then the Donaldson invariants of X may be computed by pairing $D_{X_1} \in I_*(Y)$ against $D_{X_2} \in I_*(-Y) =: I^*(Y)$.

Three-manifolds with trivial first Betti number (rational homology 3-spheres) provide an important extension of integer homology 3-spheres. The naive generalization of Floer’s instanton homology to rational homology spheres does not produce a well-defined invariant (but see Theorem 2 below). Instead, one needs to introduce a variation called the *equivariant* instanton homology groups. Such a construction was first attempted in [AB96], though it was not established that the result is an invariant of the homeomorphism type of 3-manifolds. An analogous and more transparently well-defined construction for integer homology spheres was carried out in [Don02, Section 7.3.4].

Our first theorem is that a variant of the Austin and Braam construction does indeed produce well-defined invariants of rational homology spheres.

Theorem 1. *For any PID R and any rational homology 3-sphere Y , there are three well-defined $\mathbb{Z}/8$ -graded groups*

$$I^+(Y; R), \quad I^-(Y; R), \quad I^\infty(Y; R),$$

which are invariants of the homeomorphism type of Y . These three groups are graded modules over the ring $H^{-}(BSO(3); R)$ and fit into an exact triangle*

$$\dots \xrightarrow{[-4]} I^+(Y; R) \xrightarrow{[3]} I^-(Y; R) \rightarrow I^\infty(Y; R) \xrightarrow{[-4]} \dots$$

in which each arrow is an $H^{-}(BSO(3); R)$ -module homomorphism, and the grading shifts are indicated above each arrow.*

More generally, we define invariants for 3-manifolds equipped with $SO(3)$ -bundles. The more general statement appears as Theorem 2.9.

Previously, instantons have been used to define *framed Floer homology* $I^\sharp(Y)$ for any 3-manifold Y [KM11]. In the case that Y is a rational homology sphere, the invariants of the present paper are expected to strictly contain more information than framed Floer homology. To be a bit more specific and analogous to the case of integer homology spheres in [Sca15], a connected-sum theorem would show that

our package of invariants recovers Kronheimer and Mrowka’s $I^\sharp(Y)$. A more precise statement appears as Remark 2.10. For readers familiar with Heegaard Floer homology [OS04], we remark that $I^\bullet(Y)$ is the formal counterpart of $\text{HF}^\bullet(Y)$ for $\bullet \in \{+, -, \infty\}$ whereas $I^\sharp(Y)$ is formally similar to $\widehat{\text{HF}}(Y)$.

The equivariant instanton Floer homology groups of Theorem 1 are constructed in [Mil19]¹, by a two-step process. First, a chain complex $\tilde{C}(Y, \pi)$ depending on a rational homology sphere Y and some auxiliary data π is constructed, so that this chain complex carries an action of the Lie group $SO(3)$. Next, an algebraic construction is applied to produce equivariant homology groups $I^\bullet(Y, \pi)$ for $\bullet \in \{+, -, \infty\}$. The auxiliary data π consists of a Riemannian metric on Y and a *regular* perturbation of the Chern-Simons functional. The novelty of this article is that the groups $I^\bullet(Y, \pi)$ are 3-manifold invariants and do not depend on π .

Each choice of π gives rise to a function $\sigma_\pi : \mathfrak{A}(Y) \rightarrow 2\mathbb{Z}$, called the *signature data* of π , where $\mathfrak{A}(Y)$ is the free orbits of the action of $\mathbb{Z}/2$ on $H^2(Y; \mathbb{Z})$ given by multiplication by ± 1 . The values of this function modulo 4 are independent of π , and are given by the dimension of certain cohomology groups. It is shown in [Mil19] that $I^\bullet(Y, \pi)$ is an invariant of (Y, σ_π) , and in this paper we show that indeed $I^\bullet(Y, \pi)$ does not depend on σ_π . See Section 2 for more details on the strategy of the proof.

There is also an extension of Floer’s instanton homology to rational homology spheres which is defined only in terms of irreducible (perturbed) flat connections. Unlike the invariants in Theorem 1, this invariant depends on the choice of signature data; we obtain a $\mathbb{Z}/8$ -graded abelian group $I_*(Y, \sigma)$ for any signature data function $\sigma : \mathfrak{A}(Y) \rightarrow 2\mathbb{Z}$. The method of the proof of Theorem 1 gives the following relation among $I_*(Y, \sigma)$ for different choices of σ .

Theorem 2. *Suppose $\sigma_0, \sigma_1 : \mathfrak{A}(Y) \rightarrow 2\mathbb{Z}$ are adjacent signature data functions in the following sense: $\sigma_1(\rho) = \sigma_0(\rho)$ for any $\rho \in \mathfrak{A}(Y)$ except one element ρ_0 where we have $\sigma_1(\rho_0) = \sigma_0(\rho_0) + 4$. Then there exists an exact triangle*

$$\begin{array}{ccc}
 I_*(Y, \sigma_0) & \longrightarrow & I_*(Y, \sigma_1) \\
 & \swarrow & \searrow \\
 & & \mathbb{Z}[\tilde{i}(\rho_0)]
 \end{array}$$

[−1]

where $\tilde{i}(\rho_0)$ is the integer given in Definition 3.14.

In Section 7, we state a conjecture which proposes a relation between the Euler characteristic of this invariant (plus a spectral correction term) and the Casson–Walker invariant [Wal92]. We verify this conjecture for lens spaces. For integer homology spheres, our claim is the main result of [Tau90].

Functoriality of the invariants

One of the most important features of Floer’s instanton homology for integer homology spheres is its connection to 4-manifolds. Given a negative-definite cobordism $W : Y \rightarrow Y'$ between integer homology

¹The definition of I^∞ and of the exact triangle above should be corrected following [Hel22, Section 5.4]

spheres with $b_1(W) = 0$, there is a well-defined induced map $I(W) : I_*(Y) \rightarrow I_*(Y')$, which gives I_* the structure of a functor over an appropriate cobordism category of 3- and 4-manifolds [Don02, Chapter 7] and [Frø02, Section 3.2].

The second major result of the present paper is related to the functoriality of the invariants of Theorem 1 with respect to cobordisms. Suppose $\text{Cob}_{3+1}^{\mathbb{Q}}$ is the category whose objects are rational homology 3-spheres with basepoints and whose morphisms between rational homology 3-spheres Y_1 and Y_2 are given by 4-dimensional smooth, oriented, and connected cobordisms $W : Y_1 \rightarrow Y_2$ together with a path γ running between basepoints of Y_i and an orientation of $H^1(W) \oplus H^{2,+}(W)$ (a homology orientation of W). The composition of morphisms is defined using composition of cobordisms. Let us also define $\text{Cob}_{3+1}^{\mathbb{Q},-}$ to be the subcategory with the same objects but whose morphisms are the negative-definite cobordisms. In this case, the vector space $H^{2,+}$ is trivial, and a homology orientation reduces to a choice of orientation on $H^1(W)$.

Theorem 3. *Let R be a PID in which 2 is invertible. Then there is a functor I^\bullet from the category $\text{Cob}_{3+1}^{\mathbb{Q}}$ to the category of $\mathbb{Z}/8$ -graded modules over the ring $H^{-*}(BSO(3); R)$ that specializes to the invariants of Theorem 1 at the level of objects. For an arbitrary PID R , there is a functor I^\bullet from the category $\text{Cob}_{3+1}^{\mathbb{Q},-}$ to the category of $\mathbb{Z}/8$ -graded modules over the ring $H^{-*}(BSO(3); R)$ satisfying a similar property.*

As with Theorem 1, there is an extension of this theorem to a more general class of 3-manifolds and $SO(3)$ -bundles, which is stated in Section 2. A special case of Theorem 3 is established in [Mil19] for *unobstructed cobordisms*; in the next section, we review the scope of the already-known functoriality results and describe our strategy to extend this to the level of generality stated above. In fact, this strategy is closely related to the method used in the proof of Theorem 1. We also expect that the assumption on invertibility of 2 in Theorem 3 can be dropped. In general, one expects that instanton theory provides a functor I^\bullet from $\text{Cob}_{3+1}^{\mathbb{Q}}$ for any PID R .

Remark. The methods in this paper are mostly unrelated to the specifics of the instanton theory. We expect that a similar set of arguments and constructions could be applied in the Seiberg-Witten setting, which would give a new way to define the S^1 -equivariant Seiberg-Witten Floer homology groups, at least in the case of rational homology spheres. The counterparts of the equivariant homology groups from Theorem 1 most likely coincide with those defined in [KM07]. For a rational homology sphere Y , there is an analogue of irreducible instanton homology, $\text{HM}_*^{\text{irr}}(Y, \mathfrak{s}, \pi)$, defined as the homology of a complex generated by irreducible solutions to the 3-dimensional Seiberg-Witten equations. We expect that an analogue of Theorem 2 holds for these irreducible monopole Floer groups; the authors are not aware of such a result in the literature.

Future directions

These results leave open a handful of foundational questions which are beyond the scope of this article. First off, the functoriality result in Theorem 3 asserts that $I^\bullet(Y; R)$ is functorial under negative-definite cobordisms for any PID R , and for arbitrary cobordisms when $1/2 \in R$.

Problem 1.1. *For any PID R , show that $I^\bullet(Y; R)$ is functorial for arbitrary cobordisms.*

We expect that this problem can be attacked using the ideas of this article (finding an algebraic home for obstructed gluing theory). On the other hand, extending functoriality on the object level seems significantly more difficult.

Problem 1.2. *Extend the definition of $I^\bullet(Y; R)$ to allow for 3-manifolds with $b_1(Y) > 0$.*

If $b_1(Y) \leq 2$ or $H_1(Y; \mathbb{Z})$ has no torsion except 2-torsion, then as in [AB96, Proposition 3.9] one can achieve transversality for moduli spaces on the cylinder $\mathbb{R} \times Y$, and it seems plausible that our proofs of invariance and functoriality extend to this setting. For $b_1(Y) > 2$, there will be obstructed reducibles on the cylinder; it already becomes difficult to define the differential on $\tilde{C}(Y)$. It is plausible that the ideas of this paper, used here to construct cobordism maps for obstructed cobordisms, could be helpful to make progress toward this goal. This will require a better understanding of the Uhlenbeck compactification. Understanding the Uhlenbeck compactification is also the main difficulty in making sense of Donaldson’s polynomial invariants [Don90, DK90] for negative definite cobordisms. In the case of non-negative definite cobordism between integer homology spheres, polynomial cobordism maps are constructed in [Don02].

Problem 1.3. *Extend $I^\bullet(Y; R)$ to a polynomial functor.*

There still remains the problem of developing good tools for calculation. Complete computations of the invariants of Theorem 1 only exist for some spherical manifolds. Two fundamental tools in Floer theory would help to remedy this situation.

Problem 1.4. *Construct a connected-sum quasi-isomorphism $\tilde{C}(Y) \otimes \tilde{C}(Y') \simeq \tilde{C}(Y \# Y')$, and use this to derive connected-sum formulas for I^\bullet and other related invariants.*

Problem 1.5. *Construct surgery exact triangles relating the equivariant instanton homology of Dehn surgeries on a knot.*

Frøyshov used instanton homology for integer homology spheres to define a numerical invariant h of integer homology spheres that is invariant with respect to homology cobordism relation [Frø02]. Similar to its counterparts in monopole Floer homology [Frø96, KM07] and Heegaard Floer homology [OS03], Frøyshov’s h -invariant has many interesting topological applications. In compare to h , it is easier to work with the invariant of [OS03], called the correction term, mainly because correction term can be generalized to rational homology spheres and its behavior with respect to Dehn surgery is better understood. The constructions of this paper provide the essential ingredients to address the following question.

Problem 1.6. *Find analogues of Frøyshov’s h -invariant for rational homology spheres.*

The Chern–Simons functionals on the space of $SU(2)$ -connections can be used to define an additional algebraic structure on equivariant instanton homology. In the absence of a perturbation of the Chern–Simons, this gives rise to a filtration of equivariant instanton complexes². In this paper, we do not attempt

²In general we need to pick a non-trivial perturbation of the Chern–Simons functional to define instanton homology, and this additional structure might not be quite a filtration. A more precise formulation of a closely related structure is given in [DS19, Subsection 7.3].

to flesh out the details of this construction. However, it is straightforward to see that the method of this paper implies that the additional structure on $\tilde{C}(Y)$ given by the Chern–Simons functional is also an invariant of Y . In particular, we expect that this structure can be employed to generalize Γ_Y from [Dae20], which is a refinement of Frøyshov’s invariant, to any rational homology sphere. Similarly, one should be able to generalize homology cobordism invariants of [NST19] using a similar strategy.

In the spirit of Problem 1.1, we should mention that the structure of the invariants are especially intricate over the ground ring $R = \mathbb{F}_2$; to compare, when $1/2 \in R$, the invariants described above are modules over $H^{-*}(BSO(3); R) \cong R[U]$, where $|U| = 4$. On the other hand, we have $H^{-*}(BSO(3); \mathbb{F}_2) \cong \mathbb{F}_2[w_2, w_3]$, where $|w_i| = i$.

Problem 1.7. *Find appropriate modifications of the previous invariants to the setting of \mathbb{F}_2 coefficients, and develop their properties.*

We hope to address some of these problems in future work.

Acknowledgements. The authors would like to thank Christopher Scaduto and Piotr Suwara for interesting conversations about the present work. The authors are grateful to the Simons Center for Geometry and Physics and the organizers of the 2019 PCMI Research Program “Quantum Field Theory and Manifold Invariants” where parts of the work were carried out. This work was completed while the first author was in residence at the Simons Laufer Mathematical Sciences Institute during Fall 2023 and supported by NSF Grant DMS-1928930.

AD: Add the grants either here or on the first page.

2 An overview of the article

In this section, we review the previous work on (equivariant) instanton homology of rational homology spheres [Mil19, AB95]. Then we review the novel techniques of the present paper and how they can be combined with the previous works to establish the theorem in the introductions. In fact, we can prove these results in the slightly more general context of *weakly admissible bundles*.

Weakly admissible bundles

A $U(2)$ -bundle on a 3-manifold Y is determined by its first Chern class. If w is an *oriented 1-cycle* (an embedded oriented 1-manifold) in Y , we may form in a canonical way a $U(2)$ -bundle $E = E_w$ over Y such that $c_1(E) = \text{PD}(w)$ and so that E carries a canonical trivialization over $Y \setminus w$. A pair (Y, w) is admissible if there is an embedded surface Σ in Y such that $w \cdot \Sigma$ is odd. This pair is weakly admissible if either $b_1(Y) = 0$ or (Y, w) is admissible.

We fix a connection B_0 on $\det(E)$ and define $\mathcal{A}(Y, w)$ as the space of all $U(2)$ connections on E that induce the connection B_0 on $\det(E)$. This is an affine space modeled on $\Omega^1(\text{ad}(E))$. Here $\text{ad}(E)$ is the vector bundle associated to E corresponding to the adjoint action of $U(2)$ on the Lie algebra $\mathfrak{su}(2)$, and $\Omega^1(\text{ad}(E))$ denotes the space of 1-forms on Y with values in $\text{ad}(E)$. We fix a basepoint y on Y and a trivialization of $\det(E)$ over y . A *framed connection* $\tilde{A} = (A, \Phi)$ on E is a pair of a connection A and a $U(2)$ -equivariant trivialization Φ of the fiber of $\text{ad}(E)$ above y . We write $\tilde{\mathcal{A}}(Y, w)$ for the space of all framed connections on E . Let $\mathcal{G}(Y, w)$ be the space of all automorphisms of E whose fiberwise determinant is 1. This space acts on $\mathcal{A}(Y, w)$ and $\tilde{\mathcal{A}}(Y, w)$ by taking pullback of connections (and their framing), and the quotient spaces are respectively denoted by $\mathcal{B}(Y, w)$ and $\tilde{\mathcal{B}}(Y, w)$; the gauge group action is free on $\tilde{\mathcal{A}}(Y, w)$, so that $\tilde{\mathcal{B}}(Y, w)$ is an infinite-dimensional manifold. These quotients are called the configuration spaces of connections and framed connections associated to (Y, w) . Changing framing of framed connections induces an action of $SO(3) = SU(2)/\{\pm I\}$ on $\tilde{\mathcal{B}}(Y, w)$ with the quotient space being $\mathcal{B}(Y, w)$.

There are three different possibilities for the stabilizer Γ_A of a connection $A \in \mathcal{A}(Y, w)$ with respect to the action of $\mathcal{G}(Y, w)$:

- (i) $\Gamma_A = \{\pm I\}$ where I is the trivial automorphism of E . Any such connection is called *irreducible*. The framed connections with the underlying connection A determines an orbit of the $SO(3)$ action of $\tilde{\mathcal{B}}(Y, w)$ which is diffeomorphic to $SO(3)$. We denote such orbits by the initial letters of the Greek alphabet α, β and γ .
- (ii) $\Gamma_A \cong S^1$. Any such connection is called *abelian*. The framed connections with the underlying connection A determines an $SO(3)$ -orbit diffeomorphic to $S^2 = SU(2)/S^1$ in $\tilde{\mathcal{B}}(Y, w)$. We usually denote such orbits by ρ .
- (iii) $\Gamma_A \cong SU(2)$. Any such connection is called *central*. The framed connections with the underlying connection A determines an $SO(3)$ -orbit consisting of one point in $\tilde{\mathcal{B}}(Y, w)$. We usually denote such orbits by θ .

In the case that A is not irreducible, we say A is reducible.

We may define a Chern-Simons functional on the configuration spaces of connections associated to (Y, w) . Fix a connection $A_0 \in \mathcal{A}(Y, w)$ and let A be any other connection in $\mathcal{A}(Y, w)$. Let \mathbf{A} be a connection on the pullback of E to $\mathbf{R} \times Y$ such that the restrictions of \mathbf{A} to $(-\infty, -1] \times Y$ and $[1, \infty) \times Y$ are respectively equal to the pullbacks of A and A_0 . Then we defined the value of the Chern-Simons functional at A to be the Chern-Weil integral

$$\frac{1}{8\pi^2} \int_{\mathbf{R} \times Y} \text{tr}(F_{\mathbf{A}} \wedge F_{\mathbf{A}}),$$

where $F_{\mathbf{A}}$ denotes the curvature of \mathbf{A} . Varying the connection \mathbf{A} by a homotopy does not change the value of the above integral, and changing A by an element of $\mathcal{G}(Y, w)$ changes the above expression by an integer. In particular, we have a well-defined functional $CS : \mathcal{B}(Y, w) \rightarrow \mathbb{R}/\mathbb{Z}$ that is independent of A_0 up to an overall translation. Precomposing CS with the map $\tilde{\mathcal{B}}(Y, w) \rightarrow \mathcal{B}(Y, w)$ which forgets the framing defines a functional on $\tilde{\mathcal{B}}(Y, w)$, which is still called the Chern-Simons functional.

A connection A represents a critical point of the Chern-Simons functional if the induced adjoint connection $\text{ad}(A)$ on the $SO(3)$ -bundle $\text{ad}(E)$ is flat. With a slight abuse of language, we call any such connection a flat connection associated to (Y, w) . We denote the space of all such flat connections by $\mathfrak{C}(Y, w)$, and write this set as a disjoint union based on the three possibilities for the stabilizers of flat connections:

$$\mathfrak{C}(Y, w) = \mathfrak{C}^*(Y, w) \sqcup \mathfrak{A}(Y, w) \sqcup \mathfrak{Z}(Y, w).$$

We can identify $\mathfrak{C}(Y, w)$ with the character variety

$$\chi(Y, w) := \{\phi : \pi_1(Y \setminus w) \rightarrow SU(2) \mid \phi \text{ maps any meridian of } w \text{ to } -I\} / SU(2).$$

Here the quotient is given by the conjugation action. The following lemma asserts that reducible flat connections associated to (Y, w) are determined by the cohomology groups of Y .

Lemma 2.1. *If (Y, w) is an admissible pair, then there is no reducible flat connection associated to (Y, w) . If $b_1(Y) = 0$, then the set of reducible flat connections form a discrete subset of $\mathcal{B}(Y, w)$, and are characterized as follows.*

- (i) *For any element θ of $\mathfrak{Z}(Y, w)$, there is a θ -parallel splitting of E into a direct sum of complex line bundles $L_{\theta}^{\oplus 2}$, where the two summands are canonically isomorphic. The map $\theta \mapsto c_1(L_{\theta})$ defines a canonical bijection*

$$\mathfrak{Z}(Y, w) \cong \{\xi \in H^2(Y; \mathbb{Z}) \mid 2\xi = \text{PD}(w)\}.$$

Moreover, when $\mathfrak{Z}(Y, E)$ is nonempty, the group $H^1(Y; \mathbb{Z}/2)$ acts on it transitively via $x \cdot \xi = \xi + \beta x$ with β being Bockstein homomorphism.

- (ii) *For any element ρ of $\mathfrak{A}(Y, w)$, there is a ρ -parallel splitting of E as a direct sum of S^1 -bundles L_{ρ} and L'_{ρ} . The map $\rho \mapsto \{c_1(L_{\rho}), c_1(L'_{\rho})\}$ defines a canonical bijection*

$$\mathfrak{A}(Y, w) \cong \{\{z_1, z_2\} \subset H^2(Y; \mathbb{Z}) \mid z_1 + z_2 = \text{PD}(w), z_1 \neq z_2\}.$$

Perturbations and signature data

To define something like a Morse complex for a functional, we need to assume its Hessian at any critical orbit has kernel equal to the tangent space at that orbit, in which case we say it is non-degenerate. This is usually not true for the Chern–Simons functional, so we will need to add a perturbation term.

The standard perturbations of the Chern–Simons functional is given by *holonomy perturbations* [Flo88, Formula (1b.9)]. Suppose $CS + \pi : \mathcal{B}(Y, w) \rightarrow \mathbb{R}/\mathbb{Z}$ is a perturbation of the Chern–Simons functional. A critical point of $CS + \pi$ is non-degenerate if the Hessian of the functional $CS + \pi$ at this connection, which can be regarded as an elliptic operator, has trivial kernel. (See Subsection 3.1 for more details on the definition of this operator.)

When $b_1(Y) = 0$, every element of $\mathfrak{Z}(Y, w)$ is non-degenerate. Moreover, the restriction of the Chern–Simons functional to the locus of abelian connections is non-degenerate. Motivated by this, we only consider perturbations π of the Chern–Simons functional which vanish on the abelian locus and on a neighborhood of the central locus. For any such perturbation π , reducible critical points of $CS + \pi$ agree with that of CS , and hence we have the following splitting of the critical points $\mathfrak{C}_\pi(Y, w)$ of $CS + \pi$:

$$\mathfrak{C}_\pi(Y, w) = \mathfrak{C}_\pi^*(Y, w) \sqcup \mathfrak{A}(Y, w) \sqcup \mathfrak{Z}(Y, w).$$

We can always find a perturbation term π as above such that all elements of $\mathfrak{C}_\pi(Y, w)$ are non-degenerate [Mil19, Theorem 3.6]. We call any such perturbation *nice*.

We associate a *signature data function* $\sigma_\pi : \mathfrak{A}(Y, w) \rightarrow 2\mathbb{Z}$ to any nice perturbation π . Any ρ in $\mathfrak{A}(Y, w)$ is a critical point of $CS + t\pi$ for any $t \in [0, 1]$. Thus the Hessian $H_t(\rho)$ of $CS + t\pi$ at the critical point ρ defines a 1-parameter family of Fredholm operators that are perturbations of elliptic operators with compact operators and act on the same space. Then $\sigma_\pi(\rho)$ is defined as

$$\sigma_\pi(\rho) = \dim \ker H_0(\rho) - 2\text{sf} \{H_{1-t}(\rho)\}_{t \in [0,1]},$$

the latter term being the spectral flow of the family $H_{1-t}(\rho)_{t \in [0,1]}$. When π is a small perturbation, this quantity can be understood as the signature of the restriction of $H_1(\rho)$ to $\ker H_0(\rho)$, which is a nondegenerate symmetric bilinear form. The quantity $\sigma_\pi(\rho)$ is equal modulo $4\mathbb{Z}$ to $\dim H_0(\rho)$.

Review of equivariant instanton homology

The basic object that instanton theory associates to a weakly admissible pair (Y, w) is the *tilde complex*³[Mil19]. Suppose π is a nice perturbation for (Y, w) and R is a ring. We may associate a relatively $\mathbb{Z}/8$ -graded chain complex of R -modules

$$\tilde{C}(Y, w, \pi; R), \tag{2.2}$$

equipped with an action of the dg-algebra

$$C_*(SO(3); R) \tag{2.3}$$

³This object is called the ‘framed instanton complex’ of (Y, w) in [Mil19]. We use the term ‘tilde-complex’ in this paper to avoid confusion with the Kronheimer and Mrowka’s $I^\sharp(Y)$.

of chains on the Lie group $SO(3)$. When there is no risk of confusion we drop π and R from our notation. Any chain complex with this dg-module structure is called an $SO(3)$ -complex.

The underlying $SO(3)$ -module of $\tilde{C}(Y, w, \pi; R)$ is defined in terms of the critical points of $CS + \pi$. The complex $\tilde{C}(Y, w, \pi)$ is the direct sum of $C_*(\alpha)$, where α runs over the critical orbits of $CS + \pi$ in the framed configuration space $\tilde{\mathcal{B}}(Y, w)$. Thus $\tilde{C}(Y, w, \pi)$ contains a summand $C_*(SO(3))$ for any element of $\mathfrak{C}_\pi^*(Y, w)$, a summand of $C_*(S^2)$ for any element of $\mathfrak{A}(Y, w)$ and a summand of $C_*(\text{point})$ for any element of $\mathfrak{Z}(Y, w)$. To define these R -modules, we use a certain geometric models for chains where $C_*(\alpha)$ is generated by ϕ where $\phi : P \rightarrow \alpha$ is a smooth map from a space P into α with P being a structured space analogous to a smooth manifold with corners. The definition of *geometric chains* used in our construction will be discussed in Section 4.1 and the foundational material carefully developed in Appendix A. The action of $C_*(SO(3))$ on $\tilde{C}(Y, w, \pi)$ is induced by the $SO(3)$ -action on $\tilde{\mathcal{B}}(Y, w)$.

For each two orbits α, β , (perturbed) *anti-self-duality* (ASD) equation determines a compact moduli space $\check{M}_\zeta^+(\alpha, \beta)$, equipped with an $SO(3)$ -action and $SO(3)$ -equivariant *restriction* maps to α and β . For a generic nice perturbation, these moduli spaces are smooth(ish), and any such perturbation is called *regular*. The appropriate analogue of the usual point-counting differential in Morse theory is instead a *fiber product differential*; for any $\phi : P \rightarrow \alpha$ representing an element of $C_*(\alpha)$, we have

$$\tilde{\partial}\phi = \partial_{gm}\phi + \sum_{\beta} \pm P \times_{\alpha} \check{M}_\zeta^+(\alpha, \beta),$$

where the first term is the geometric boundary of ϕ (again a chain in α), while each term in the sum is a chain in some $C_*(\beta)$. The fact that this operator squares to zero is encoded in boundary relations for the moduli spaces $\check{M}_\zeta^+(\alpha, \beta)$; the boundary components of $\check{M}_\zeta^+(\alpha, \beta)$ may be described as fiber products of smaller-dimensional moduli spaces. The definition of $\check{M}_\zeta^+(\alpha, \beta)$ and the determination of its boundary strata are recalled in Subsection 3.4, and the precise definition of the differential $\tilde{\partial}$ is given in Subsection 4.3. This construction deserves to be called the *Morse-Bott complex* of $CS + \pi$ on $\tilde{\mathcal{B}}(Y, w)$; a de Rham version was initially developed in [AB95], together with its equivariant homology.

A limited version of functoriality of tilde complexes with respect to cobordisms is established in [Mil19]. First let (Y, w, π) and (Y', w', π') be a weakly admissible pairs with chosen regular perturbations. Suppose $W : Y \rightarrow Y'$ is a cobordism and c is a properly oriented embedded surface in W such that the intersections of c with Y and Y' are respectively w and w' . Then c canonically determines a $U(2)$ bundle \mathbf{E}_c over W such that $c_1(\mathbf{E}_c) = \text{PD}(c)$. In particular, the restrictions of \mathbf{E} to Y and Y' are the $U(2)$ -bundles E_w and $E_{w'}$, respectively. The cobordism is also equipped with a path γ in W running between the chosen basepoints of Y and Y' . As with the basepoints, we usually suppress γ from our notation for cobordisms.

The ASD equation on W (after adding cylindrical ends) gives rise to a moduli space $M_z^+(W, c; \alpha, \alpha')$ for any pair of an orbit of π -flat connection α on Y and an orbit of π' -flat connection α' on Y' . An element of this moduli space is the gauge equivalence of a pair (A, Φ) of a (possibly broken) connection A on W , which satisfies a perturbation of the ASD equation compatible with the perturbation terms on the ends, and a framing Φ along γ which is parallel with respect to A . One might hope that after picking an appropriate perturbation of the ASD equation, $M_z^+(W, c; \alpha, \alpha')$ is a compact smooth(ish) space. However, there is an obstruction which can be expressed in terms of σ_π and $\sigma_{\pi'}$, the topology of W and the cohomology

class $\text{PD}(c)$, and when this obstruction is trivial, we say (W, c, π, π') is *unobstructed*. Otherwise, it is obstructed. The precise definition of $M_z^+(W, c; \alpha, \alpha')$ and unobstructedness of (W, c, π, π') will be given in Subsections 3.2 and 3.5 respectively. As an important example, let π and π' be two regular perturbations for a weakly admissible pair (Y, w) . Then $([0, 1] \times Y, [0, 1] \times w)$ as a cobordism from (Y, w) with perturbation term π to (Y, w) with perturbation term π' is obstructed if and only if $\sigma_\pi \leq \sigma_{\pi'}$.

The second main result of [Mil19] is that when (W, c, π, π') is an unobstructed cobordism with a path between the baespoints, a choice of an appropriate perturbation of the ASD equation and a homology orientation for W (an orientation of the vector space $H^1(W) \oplus H^{2,+}(W) \oplus H^1(Y)$), we may define a homomorphism

$$\tilde{C}_*(W, c) : \tilde{C}_*(Y, w, \pi) \rightarrow \tilde{C}_*(Y', w', \pi').$$

This homomorphism is $SO(3)$ -equivariant and is independent of the Riemannian metric on W and the perturbation of the ASD equation up to equivariant chain homotopy.

Finally, the third main result of [Mil19] is that given an algebraic object resembling $\tilde{C}_*(Y, w, \pi)$ (more precisely, a relatively graded dg-module equipped with a *periodic filtration*), one may construct *equivariant homology groups* which fit into an exact triangle. Applying this to $\tilde{C}(Y, w, \pi)$ one obtains the equivariant instanton homology groups I^+ , I^- , I^∞ of (Y, w, π) .

The definition of I^∞ in [Mil19, Section A.8] is flawed; the complex CI^∞ is defined in terms of a completion of an algebraic construction. However, the completion process described there does not have all the properties stated. Instead, one should use the more robust *full completion*, as in [Hel22, Section 5.4]; using this modified notion of completion, all the stated results of [Mil19, Section A.8] hold.

Usually, one proves that Floer homology groups are independent of perturbation, and hence topological invariants, by employing functoriality with respect to cobordisms. Unfortunately, the unobstructedness assumption is very strong. Such an argument can be used to show that equivariant instanton homology groups are invariants of (Y, w, σ) , where σ is the signature data of the perturbation. To show that these Floer homology groups do not depend on the choice of signature data in some subtle way, one also needs cobordism maps associated to the product cobordism $([0, 1] \times Y, [0, 1] \times w)$ when the signature data does not necessarily increase from the incoming end to the outgoing end. The primary purpose of the present paper is to extend the relevant functoriality properties to remedy this issue.

Results of this article

Suppose $(W, c) : (Y, w) \rightarrow (Y', w')$ is a cobordism with $b_1(W) = b^+(W) = 0$. We fix a regular perturbation π_i for (Y_i, w_i) . We also fix a Riemannian metric with cylindrical ends on W . The issue of obstructedness for (W, c, π, π') is related to the *reducible instantons*, the elements Λ in $M_z^+(W, c; \alpha, \alpha')$ which have non-trivial stabilizer with respect to the $SO(3)$ action. As is explained in Proposition 3.22, any such reducible Λ is determined by some cohomological data. Associated to any such reducible Λ , there is a quantity $N(\Lambda, \pi, \pi') \in 2\mathbb{Z}$ called the *normal index*. The expected dimension of the moduli space $M_z^+(W, c; \alpha, \alpha')$ containing Λ is equal to $N(\Lambda, \pi, \pi') + 2$. The quantity $N(\Lambda, \pi, \pi')$ can be computed in terms of some cohomological data and the signature data of π and π' (see Proposition 3.24). In particular, if Λ is asymptotic to the abelian reducibles ρ and ρ' on the two ends, then for regular perturbations π_1 and

π_2 for (Y, w) and regular perturbations π'_1 and π'_2 for (Y', w') we have

$$N(\Lambda, \pi_1, \pi'_1) - N(\Lambda, \pi_2, \pi'_2) = (\sigma_{\pi_1}(\rho) - \sigma_{\pi_2}(\rho)) - (\sigma_{\pi'_1}(\rho') - \sigma_{\pi'_2}(\rho')). \quad (2.4)$$

In the case that ρ or ρ' is central, the corresponding term on the right-hand side of the above identity vanishes. The tuple (W, c, π, π') is unobstructed exactly when $N(\Lambda, \pi, \pi')$ is non-negative for all possible choices of Λ .

Definition 2.5. We say (W, c, π, π') is *nearly unobstructed* if there is one reducible instanton Λ with $N(\Lambda, \pi, \pi') = -2$ and all the remaining reducibles have non-negative normal index. \diamond

Example 2.6. Suppose (π_1, π_2) are two regular perturbations of (Y, w) with *adjacent signature data*, meaning that $\sigma_{\pi_1} = \sigma_{\pi_2}$ except at one abelian flat connection ρ_0 , where $\sigma_{\pi_1}(\rho_0) = \sigma_{\pi_2}(\rho_0) + 4$. Then $(I \times Y, I \times w, \pi_1, \pi_2)$ is nearly unobstructed. On the other hand, $(I \times Y, I \times w, \pi_2, \pi_1)$ is unobstructed. \diamond

The following is our main technical theorem which allows us to obtain the results advertised in the introduction. The proof of this theorem will be given in Section 5.

Theorem 2.7. *Associated to any weakly admissible pair (Y, w) with a regular perturbation π and a choice of abelian flat connection ρ on Y , there is an $SO(3)$ -complex $S_\rho \tilde{C}(Y, w, \pi)$ and an $SO(3)$ -equivariant quasi-isomorphism*

$$S_* : \tilde{C}(Y, w, \pi) \rightarrow S_\rho \tilde{C}(Y, w, \pi).$$

Suppose $(W, c) : (Y, w) \rightarrow (Y', w')$ is a cobordism as above with path between the basepoints and π, π' are regular perturbations for $(Y, w), (Y', w')$ such that (W, c, π, π') is nearly unobstructed with exactly one obstructed reducible Λ running between abelian flat connections ρ and ρ' . Then there is an equivariant chain map

$$\tilde{C}(Y, w, \pi) \rightarrow S_{\rho'} \tilde{C}(Y', w', \pi'),$$

which is well-defined up to equivariant homotopy.

Since $SO(3)$ -equivariant quasi-isomorphisms induce isomorphisms on equivariant homology groups, the map S_* may be inverted at the level of equivariant homology. It follows that nearly unobstructed cobordisms give rise to maps on equivariant instanton homology. In Section 7 we use the above theorem to prove Theorem 1 from the introduction.

The theorem above, together with an analysis of composites, is nearly enough to prove the invariance of equivariant instanton homology. Suppose (π_1, π_2) are two regular perturbations of (Y, w) with *adjacent signature data*. Then Theorem 2.7 gives

$$\tilde{C}_*(Y, w, \pi_1) \rightarrow S_\rho \tilde{C}_*(Y, w, \pi_2). \quad (2.8)$$

Moreover, $(Y \times I, w \times I, \pi_2, \pi_1)$ is unobstructed and the construction of [Mil19] gives a cobordism map $\tilde{C}_*(Y, w, \pi_1) \rightarrow \tilde{C}_*(Y, w, \pi_2)$ and in Subsection 5.3 we shall show that this map can be lifted to a homomorphism

$$S_\rho \tilde{C}_*(Y, w, \pi_1) \rightarrow \tilde{C}_*(Y, w, \pi_2);$$

in Subsection 5.5 we will show that it is an equivariant homotopy inverse to the homomorphism in (2.8). Finally, for any two regular perturbations π, π' of (Y, w) , one may always find a sequence

$$\pi = \pi_0 < \cdots < \pi_n = \pi'_m > \cdots > \pi'_0 = \pi'$$

of regular perturbations such that (π_i, π_{i+1}) and (π'_{j+1}, π'_j) have adjacent signature data. This completes the proof of the following generalization of Theorem 1.

Theorem 2.9. *For any weakly admissible pair (Y, w) , there are relatively $\mathbb{Z}/8$ -graded groups,*

$$\tilde{I}(Y, w), \quad I^+(Y, w), \quad I^-(Y, w), \quad I^\infty(Y, w),$$

which are invariants of the oriented diffeomorphism type of Y and the element of $H^2(Y; \mathbb{Z}/2)$ determined by w . The first is a graded module over the ring $H_*(SO(3); R)$. The next three are graded modules over the ring $H^{-*}(BSO(3); R)$ and fit into an exact triangle

$$\cdots \xrightarrow{[-4]} I^+(Y, w) \xrightarrow{[3]} I^-(Y, w) \rightarrow I^\infty(Y, w) \xrightarrow{[-4]} \cdots$$

in which each arrow is an $H^{-*}(BSO(3); R)$ -module homomorphism, and the grading shifts are indicated above each arrow.

Remark 2.10. When (Y, w) is admissible, we have $I^\infty(Y, w) = 0$ and $I^+(Y, w) \cong I^-(Y, w)[-3]$ are canonically isomorphic to Floer's instanton homology groups $I_*(Y, w)$ for admissible bundles, while $\tilde{I}(Y, w) = 0$.

When Y is a rational homology sphere, we anticipate a relationship to Kronheimer and Mrowka's $I^\sharp(Y, w)$ — roughly, I_* of $(Y \# T^3, w \sqcup S^1)$ — which cannot be extracted directly from the homology groups above, but rather from the $SO(3)$ -complex \tilde{C} . Precisely, writing u for the action of $[SO(3)] \in C_3(SO(3); R)$ on \tilde{C} , a connected sum theorem should give an isomorphism between $\mathbb{Z}/4$ -graded groups $H(\tilde{C}(Y, w), \tilde{d} - 4u) \cong I^\sharp(Y, w)$. \diamond

Our next goal is to define cobordism maps for more general cobordisms (arbitrary cobordisms when $1/2 \in R$, and arbitrary negative-definite cobordisms otherwise). Were it not for difficulties with the Uhlenbeck compactification, it seems plausible that the techniques used in the proof of Theorem 2.7 could be used to construct induced maps for arbitrary cobordisms. Instead, we opt for a different approach. First we must be slightly more specific about the sort of issues one runs into in the definition of cobordism maps. When enumerating the reducible instantons on W , they come in three essentially different types:

- Abelian instantons Λ on W which are not central on the boundary.
- Abelian instantons Λ which restrict to central connections on the boundary. In this case the normal index of Λ is independent of the perturbations π and π' . If Λ is flat, then the normal index of Λ is -2 . Otherwise the normal index is non-negative and it is not a cause for concern. Flat abelian instantons which are central on the boundary are called *pseudocentral instantons*.
- Central connections, which are necessarily flat. These are cut out transversely when $b^+(W) = 0$ and not otherwise.

The first key insight is that because $I^\bullet(Y, w)$ and $I^\bullet(Y', w')$ are independent of the choice of perturbations π and π' , *we are free to choose the perturbation on the ends however we like*. In practice, this implies that the first type of abelian instanton is irrelevant: we may now choose perturbation data π and π' so that $N(\Lambda, \pi, \pi') \geq 0$ for any reducible Λ of the first type (see the identity in (2.4)). Therefore, the perturbation term over the cobordism can be chosen so that all such reducibles are cut out transversely. To resolve the issue related to pseudocentral reducibles, we remove a regular neighborhood of the elements of the moduli spaces containing the pseudocentral reducibles. An analysis of the boundary of these spaces allows to define an induced map between equivariant instanton homology groups. Resolving the issues related to the first two types of reducible instantons allows us to prove the second part of Theorem 3.

Central connections complicate the definition of cobordism maps in the case that $b^+(W) > 0$. It seems likely that a variation on the construction used in Theorem 2.7 will allow us to define the induced maps of such cobordism maps and hence complete the proof of Theorem 3. However, we do not attempt to carry this out here. Instead, we resolve the issue of obstructed reducibles with a blowup trick used to define Donaldson invariants past the realm they are usually defined; for instance, see [MM92, KM95, Kro05]. More precisely, we replace moduli spaces associated to (W, c) with the moduli spaces $(W \# \overline{\mathbb{C}\mathbb{P}^2}, c \cup D)$ cutting down by $\mu(D)$. Here D denotes the exceptional divisor in $\overline{\mathbb{C}\mathbb{P}^2}$. The new moduli spaces do not have any central flat connections. Moreover, if the cobordism map for (W, c) is already defined, then the above blowup process gives the same cobordism map up to a multiple of 2. In particular, if 2 is a unit in the ring R , then this multiple of 2 is harmless and hence it justifies that the blowup process gives a sensible definition of cobordism maps. This is why we require that 2 is invertible in the first part of Theorem 3.

In fact, we can slightly extend Theorem 3 to allow cobordisms between weakly admissible pairs. First we extend the definition of categories $\text{Cob}_{3+1}^{\mathbb{Q}}$ and $\text{Cob}_{3+1}^{\mathbb{Q}, -}$.

Definition 2.11. The category Cob_{3+1}^w has as its objects triples (Y, w, y) , where (Y, w) is a weakly admissible pair and $p \in Y$ is a basepoint. The morphisms $(Y, w, p) \rightarrow (Y', w', p')$ are quadruples (W, c, γ, o) , where (W, c) is an oriented connected cobordism $(Y, w) \rightarrow (Y', w')$, with γ a path between the basepoints, and o a homology orientation. Composition is given by gluing along the common boundary of two cobordisms.

The category $\text{Cob}_{3+1}^{w, -}$ has the same objects, and those morphisms for which either $b^+(W) = 0$ or (W, c) admits no central connections. \diamond

Then the functors I^\bullet extend to these categories, and the exact triangle is functorial.

Theorem 2.12. *Let R be a PID.*

- (i) *There is a functor I^\bullet from the category $\text{Cob}_{3+1}^{w, -}$ to the category of relatively $\mathbb{Z}/8$ -graded modules over the ring $H^{-*}(BSO(3); R)$ that specializes to the invariants of Theorem 2.9 at the level of objects; the induced maps commute with the structure maps in the exact triangle of Theorem 2.9.*
- (ii) *If 2 is invertible in R , there is a functor I^\bullet from the category Cob_{3+1}^w to the category of relatively $\mathbb{Z}/8$ -graded modules over $H^{-*}(BSO(3); R)$ which extends the functor from (i).*

Organization and strategy

In Section 3 below, we recall the relevant facts about ASD connections, perturbations, and moduli spaces thereof. To define the tilde-complex in (2.2), we need to have chain-level fiber product maps with respect to whatever kind of smooth structure the moduli spaces have. Such an apparatus is called a theory of geometric chains. In Appendix A, we define and present the foundations of *stratified-smooth spaces*. Our moduli spaces carry the structure of stratified-smooth spaces, and transverse intersections of stratified-smooth spaces are again stratified-smooth. We conclude this Appendix by verifying in A.4 that we can use stratified-smooth spaces to construct a homology theory of smooth manifolds canonically isomorphic to singular homology; for technical reasons, we truncate it to a bounded complex in A.5.

In Section 4, we use this apparatus to define the notion of *flow categories*, *bimodules between flow categories*, and *homotopy between bimodules*. These amount to collections of orbits α , as well as collections of geometric chains (corresponding to the moduli spaces $\check{M}_\zeta^+(\alpha, \beta)$ in the previous discussion) so that each moduli space satisfies an appropriate boundary relation. Using the above geometric chain apparatus, we associate a *flow complex* to a flow category, a chain map of flow complexes to a bimodule and a homotopy of chain maps to a homotopy between bimodules. We then explain how the tilde-complex of (2.2) and the homomorphisms associated to unobstructed cobordisms arise from this construction.

Section 5 is the technical heart of the paper. Given a nearly unobstructed cobordism (W, c, π, π') , the moduli spaces $M_z^+(W, c; \alpha, \alpha')$ will no longer be smooth manifolds with corners. However, if we delete a neighborhood of the ‘non-smooth points’, we are left with a stratified-smooth space which has new boundary strata, that we must investigate. While these new boundary strata are not explicit, we may correct them by adding a term interpolating to some more standard expressions (in terms of zero sets of sections of an *obstruction line bundle*). Doing so, we define *modified moduli spaces* which satisfy an explicit boundary relation — though not the right boundary relation to define a chain map between tilde-complexes.

Investigating these boundary relation leads us to the notion of *blown-up* flow category, which we explore in Section 5.3; its flow complex is the $SO(3)$ -complex $S_\rho \tilde{C}$ discussed above. After doing so, we verify in Section 5.4 that the modified moduli spaces do define bimodules with respect to these blown-up flow categories, resulting in Theorem 2.7. In Section 5.5 we prove that these compose in a reasonable way and thus prove Theorem 2.9.

With this theorem in hand, only a little work remains. We conclude by verifying that negative-definite cobordisms (or cobordisms without central connections) with pseudocentral reducibles still give chain maps well-defined up to homotopy in Section 6.1. In Section 6.2 we show that these compose functorially. We then discuss the blowup trick in Section 6.3, which gives us functoriality for arbitrary cobordisms when our coefficient ring contains $1/2$.

In the final section, we explain how to use our constructions to extract results about the extension of Floer’s *irreducible instanton homology* to rational homology spheres, and how to use this to construct a numerical invariant which conjecturally agrees with the Casson-Walker invariant.

3 Moduli spaces of instantons

3.1 Linear analysis on 3-manifolds

Let (Y, w) be a weakly admissible pair with the associated $U(2)$ -bundle E . For a non-negative integer k , let $\mathcal{A}_k(Y, w)$ and $\mathcal{G}_k(Y, w)$ be respectively the L_k^2 -completions of the space of connections $\mathcal{A}(Y, w)$ and the gauge group $\mathcal{G}(Y, w)$ introduced in Section 2. In particular, $\mathcal{A}_k(Y, w)$ can be identified with the affine space $A_0 + \Omega_k^1(\text{ad}(E))$ where A_0 is a smooth reference connection and $\Omega_k^1(\text{ad}(E))$ denotes the space of L_k^2 sections of the bundle $\Lambda^1 \otimes \text{ad}(E)$ over Y . Fix an integer $k \geq 2$ for the rest of the paper. The space $\mathcal{G}_{k+1}(Y, w)$ is a Banach Lie group acting smoothly on $\mathcal{A}_k(Y, w)$. We may thus define $\mathcal{B}_k(Y, w)$ to be the quotient of this action. We similarly define the space of framed connections $\tilde{\mathcal{A}}_k(Y, w)$ and the configuration space of framed connections $\tilde{\mathcal{B}}_k(Y, w)$.

The framed configuration space is a Banach manifold with smooth $SO(3)$ -action, and hence has a natural $SO(3)$ -equivariant tangent bundle $T_{[A, \Phi]} \tilde{\mathcal{B}}_k(Y, w)$. This may be identified with the quotient of $\Omega_k^1(Y, \text{ad}(E)) \oplus \mathfrak{so}(3)$ by the image of

$$(d_A, \text{ev}_y) : \Omega_{k+1}^0(Y, \text{ad}(E)) \rightarrow \Omega_k^1(Y, \text{ad}(E)) \oplus \mathfrak{so}(3).$$

Notice that this tangent bundle has natural Sobolev completions $T_j \tilde{\mathcal{B}}_k(Y, w)$ for all $1 \leq j \leq k$. (The point-evaluation component requires that our sections of $\text{ad}(E)$ are continuous, and hence the restriction that they are of regularity at least L_2^2 .)

We are now able to give a careful discussion of the Hessian operator. The formal gradient of the perturbed Chern-Simons functional defines an $SO(3)$ -equivariant section

$$\nabla(CS + \pi) : \tilde{\mathcal{B}}_k \rightarrow T_{k-1} \tilde{\mathcal{B}}_k.$$

When the perturbation is zero, its zeroes are the orbits of framed flat connections. At a given critical point $[A, \Phi]$ with stabilizer Γ , there is further a formal Γ -equivariant Hessian operator $\text{Hess}(CS + \pi) : T_k \tilde{\mathcal{B}}_k \rightarrow T_{k-1} \tilde{\mathcal{B}}_k$. This gives an acceptable notion of *nondegenerate* critical point: the kernel of this operator should be the tangent space to the orbit \mathcal{O} through $[A, \Phi]$ (in other words, as small as possible; by $SO(3)$ -invariance of $CS + \pi$ the tangent space $T_{[A, \Phi]} \mathcal{O}$ will always be a subspace of the kernel).

Alternatively we can define nondegeneracy of critical points in terms of *extended Hessian* operator. The extended Hessian operator for a connection A

$$\widehat{\text{Hess}}_A(CS) : L_k^2(\Omega^0 \oplus \Omega^1)(\text{ad}(E)) \rightarrow L_{k-1}^2(\Omega^0 \oplus \Omega^1)(\text{ad}(E))$$

is defined by the matrix $\begin{pmatrix} 0 & -d_A^* \\ -d_A & *d_A \end{pmatrix}$. This is a self-adjoint elliptic operator, and has discrete spectrum diverging to $\pm\infty$. If one instead works with $\widehat{\text{Hess}}_A(CS + \pi)$, then there is an additional term, given by a compact operator which factors as $L_k^2 \xrightarrow{\text{Hess}(\pi)} L_k^2 \hookrightarrow L_{k-1}^2$; then $\widehat{\text{Hess}}_A(CS + \pi)$ is again a self-adjoint Fredholm operator with discrete spectrum diverging to $\pm\infty$. The extended Hessian is now defined for all connections A and perturbations π , not just those for which $\nabla_A(CS + \pi) = 0$, and is defined with

the same domain and codomain for all (A, π) . It is thus a suitable home for a notion of spectral flow. It follows from a straightforward analysis that

$$\ker \widehat{\text{Hess}}_A(CS + \pi) = \mathfrak{g}_A \iff \ker \text{Hess}_{[A, \Phi]}(CS + \pi) = \text{TO}_{[A, \Phi]},$$

where \mathfrak{g}_A is the space of A -parallel sections of $\text{ad}(E)$, and thus the notion of nondegeneracy outlined above agrees with that of Section 2. This analysis is carried out in detail in [Mil19, Section 3.2].

As discussed in Section 2, we choose our perturbations to vanish on the reducible locus and in a neighborhood of the central connections, and therefore the set $\mathfrak{A}(Y, w)$ of gauge equivalence classes of abelian flat connections coincides with the set of gauge equivalence classes of abelian critical orbits of $CS + \pi$. However, the spectral behavior at these abelian flat connections depends on π . The following definition is meant to capture this dependence.

Definition 3.1. Let (Y, w) be a rational homology sphere with an oriented 1-cycle, equipped with a nice perturbation π . Its *spectral shift* is a function $S_\pi : \mathfrak{A}(Y, w) \rightarrow 2\mathbb{Z}$, defined to be the spectral flow

$$S_\pi(\rho) = \text{sf} \left(\left\{ \widehat{\text{Hess}}_\rho(CS + (1-t)\pi) \right\}_{t \in [0,1]} \right).$$

When defining the spectral flow between possibly degenerate operators A_0 and A_1 via the path $\{A_t\}_{t \in [0,1]}$, we have

$$\text{sf}(\{A_t\}) = \lim_{\varepsilon \rightarrow 0^+} \text{sf}(\{A_t + \varepsilon I\}).$$

That is, we compute the intersection number of the eigenvalues with the line $\lambda = -\varepsilon$ for sufficiently small $\varepsilon > 0$. \diamond

Each of these Hessian operators splits as a direct sum of an operator on the reducible locus (which is the same for all t because $\pi = 0$ on the reducible locus) and an operator normal to the reducible locus. The latter is S^1 -equivariant with respect to a weight two S^1 -action, and hence complex linear with respect to an appropriate complex structure. This spectral flow can be computed as twice the complex spectral flow of a family of complex-linear operators. In particular, it is an even integer.

There is a topological interpretation of the kernel of the unperturbed extended Hessian operator $\widehat{\text{Hess}}_\rho(CS)$. If ρ is a flat abelian connection, the corresponding adjoint representation $\text{ad}_\rho : \pi_1(Y) \rightarrow SU(2)/\pm I = SO(3)$ has image inside the subgroup $SO(2)$ induced by diagonal matrices $U(1) \subset SU(2)$, and thus defines a 1-dimensional complex representation \mathbb{C}_ρ of $\pi_1(Y)$. Then we have an isomorphism

$$\ker \widehat{\text{Hess}}_\rho(CS) \cong \mathbb{R} \oplus H^1(Y; \mathbb{C}_\rho),$$

where the term $H^1(Y; \mathbb{C}_\rho)$ is the first cohomology with coefficients in the local system \mathbb{C}_ρ .

We use this to define slight variation on the spectral shift.

Definition 3.2. A *signature data function* is a function $\sigma : \mathfrak{A}(Y, w) \rightarrow 2\mathbb{Z}$ so that

$$\sigma(\rho) \equiv \dim_{\mathbb{R}} H^1(Y; \mathbb{C}_\rho) \pmod{4}.$$

Given a spectral shift $S_\pi : \mathfrak{A}(Y, w) \rightarrow 2\mathbb{Z}$, the corresponding signature data function is denoted by σ_π and is defined as

$$\sigma_\pi(\rho) = \dim_{\mathbb{R}} H^1(Y; \mathbb{C}_\rho) - 2S_\pi(\rho). \quad \diamond$$

Remark 3.3. The quantity σ_π defined here is the same as the quantity σ_π in [Mil19, Definition 3.3]. When π is a small perturbation, the perturbed Hessian restricts to a nondegenerate S^1 -invariant bilinear form on $H^1(Y; \mathbb{C}_\rho)$, and $\sigma_\pi(\rho)$ is the signature of this bilinear form. In what follows, the perturbations π are not necessarily small, but we use the name ‘signature data’ by analogy. \diamond

Later, we will need to know that every signature data function arises as the signature data of some nice perturbation π . The remainder of this section will be devoted to a proof of this fact, recorded as a proposition below.

Proposition 3.4. *Let (Y, w) be a weakly admissible pair. For any given function $f : \mathfrak{A}(Y, w) \rightarrow 2\mathbb{Z}$, there is a nice perturbation π with $S_\pi = f$. Correspondingly, for any given signature data function σ , there is a nice perturbation π with $\sigma_\pi = \sigma$.*

On the way to proving this claim, we will need the following (a priori weaker) fact.

Lemma 3.5. *Let (Y, w) be as above and let ρ be a reducible flat connection. Then for any positive integer k there exists a perturbation π such that $S_\pi(\rho) \leq -k$, and π vanishes on the reducible locus and in a neighborhood of every other reducible flat connection. A similar claim holds if we replace $S_\pi(\rho) \leq -k$ with $S_\pi(\rho) \geq k$.*

Proof. We will construct this perturbation as a non-negative function which vanishes on the reducible locus (and in a neighborhood of each reducible flat connection but ρ). Adding this perturbation can only make eigenvalues larger, and we construct it so that the k smallest negative eigenvalues become positive. As a result, we will have that

$$\text{sf}(\text{Hess}'_\rho, \text{Hess}'_\rho + \text{Hess}'(\pi)) \geq k,$$

so that $S_\pi(\rho) \leq -k$.

Let V_k be the $SO(2)$ -invariant subspace of $T'_\rho \tilde{\mathcal{B}}$ spanned by the k smallest nonpositive eigenvalues of Hess'_ρ , with $-\lambda$ being the least of these eigenvalues; one may exponentiate this to a small $SO(2)$ -invariant disc $D_k \subset \tilde{\mathcal{B}}$, and consider the $SO(3)$ -invariant subset \tilde{D}_k containing D_k .

To construct such a perturbation, first recall the definition of holonomy perturbations. We construct these using a family of thickened loops $S^1 \times D^2 \rightarrow Y$ which all coincide on a neighborhood of $\{1\} \times D^2$; taking holonomy along n such thickened loops gives an $SO(3) = \text{Inn}(SU(2))$ -equivariant map $D^2 \times \tilde{\mathcal{B}}(Y, w) \rightarrow SU(2)^N$. We then apply a smooth $SO(3)$ -invariant function $f : SU(2)^N \rightarrow \mathbb{R}$ which vanishes on the reducible locus and in a neighborhood of $\{\pm I\}^N$, and then we integrate over the disc.

Given any compact $SO(3)$ -invariant submanifold of $\tilde{\mathcal{B}}$, such as \tilde{D}_k , we may choose the family of loops so that the holonomy map embeds \tilde{D}_k into $SU(2)^N$. Similarly we may assume that the image of any other flat reducible under this map is disjoint from \tilde{D}_k . Now choose an $SO(3)$ -invariant nonnegative function f on $SU(2)^N$ which vanishes on the reducible locus, vanishes near every other flat reducible, and on the image of D_k restricts to a large quadratic function Q — so large that

$$Q(\text{Hol}(\exp v)) \geq 2\lambda \|v\|^2$$

for all sufficiently small $v \in V_k$. Modifying the auxiliary data in this construction as appropriate (in particular, taking the radius of the thickened loops to be small) guarantees that for the corresponding perturbation π , we have

$$\pi(\text{Hol}(\exp v)) \geq 1.5\lambda\|v\|^2$$

for all sufficiently small $v \in V_k$. It follows that

$$\text{Hess}'_{\rho}(CS + \pi)(v) \cdot v = \text{Hess}'_{\rho}(v) \cdot v + \text{Hess}'_{\rho}(\pi)(v) \geq -\lambda\|v\|^2 + 1.5\lambda\|v\|^2 \geq \frac{\lambda}{2}\|v\|^2,$$

and thus we have made the Hessian positive-definite on the subspace V , taking these k non-negative eigenvalues to positive eigenvalues as desired.

The same argument can be run with f a negative function which restricts on a positive eigenspace to a negative-definite quadratic form to construct a perturbation so that the spectral flow is at least k . \square

Proof of Proposition 3.4. Fix an abelian flat connection ρ , and define

$$K_{\rho}(\pi) := \dim_{\mathbb{R}}(\ker(\text{Hess}'_{\rho}(CS + \pi))) \in 2\mathbb{N}.$$

Notice that the function $\pi \mapsto S(\pi) := S_{\pi}(\rho)$ is lower semicontinuous, and the function K measures its failure to be continuous: if $K_{\rho}(\pi) = 2m$, then for all sufficiently nearby π' , we have $S(\pi) \leq S(\pi') \leq S(\pi) + 2m$. In fact, any jump between $S(\pi)$ and $S(\pi) + 2m$ is realizable. Though stated under the assumption that $\pi = 0$, the argument of [Mil19, Proposition 3.9] shows that for any $0 \leq i \leq m$, there exists a nearby perturbation π' which vanishes near every reducible but ρ so that $S(\pi') = S(\pi) + 2i$, while $K_{\rho}(\pi') = 0$. That is, given a perturbation π , by a small change we may make $2i$ of the 0-eigenvalues of $\text{Hess}'_{\rho, \pi}$ negative, and $2(m - i)$ positive, leaving no 0-eigenvalues. (The additional perturbation is chosen to be very close to a small but nondegenerate $SO(2)$ -invariant quadratic form on \mathbb{R}^{2m} with $2i$ positive eigenvalues and $2(m - i)$ negative eigenvalues.)

Now suppose we want to find a regular perturbation π with $S_{\pi}(\rho) = f(\rho) = 2n$ where $n \geq 0$; the case $n < 0$ is similar. By Lemma 3.5, there exists some perturbation π_0 and π_1 that vanish near every reducible but ρ and has

$$S(\pi_0) \leq 2n \leq S(\pi_1).$$

Now consider the family of perturbations $\pi_t := t\pi_1 + (1 - t)\pi_0$. Although it might not be the case that there is some t with $S(\pi_t) = 2n$, we may find $t \in [0, 1]$ such that

$$S(\pi_t) \leq 2n \leq S(\pi_t) + K_{\rho}(\pi_t).$$

As we mentioned above, every jump in this range is realizable; that is, there exists a perturbation π_{ρ} which vanishes near all reducibles except ρ and is near to π_t , so that $S(\pi_{\rho}) = 2n$ and $K_{\rho}(\pi_{\rho}) = 0$.

For each ρ , choose as above a perturbation π_{ρ} which vanishes near all other reducible flat connections and has $S_{\pi}(\rho) = f(\rho)$ and $K_{\rho}(\pi) = 0$. Now set $\pi = \sum_{\rho} \pi_{\rho}$. Because each π_{ρ} vanishes near all reducibles except ρ , it follows that $S_{\pi}(\rho) = f(\rho)$ and $K_{\pi}(\rho) = 0$ for all ρ , as desired. This perturbation is such that all reducible flat connections are non-degenerate, but it is possible that there are irreducible critical points which are degenerate. However, nice perturbations are dense in the space of all perturbations, so choose a nice perturbation π' near π . So long as it is chosen sufficiently close to π , this will not change the value of S_{π} . This is the desired nice perturbation with $S_{\pi'}(\rho) = f(\rho)$ for all ρ . \square

Remark 3.6. As discussed above, the kernel of the unperturbed extended Hessian is isomorphic to $H^1(Y; \mathbb{C}_\rho)$. If $0 \leq S(\rho) \leq \dim_{\mathbb{R}} H^1(Y; \mathbb{C}_\rho)$ for all ρ , then the function S may be realized as the spectral shift S_π of an arbitrarily small perturbation. However, if $S(\rho)$ fails to lie between these bounds for any ρ , then S is not realizable as S_π for any small perturbation. Some of the perturbations constructed above are necessarily large — but zero on the reducible locus itself. This does not make any difference for the constructions in [Mil19]. There the only reason to require that 3-manifold perturbations are small is to ensure that the reducible critical set does not change. Here we enforce this with the condition on the functions defining the holonomy perturbations that $f_\pi = 0$ on the reducible locus as opposed to a smallness condition. \diamond

3.2 The ASD equations on 4-manifolds

We are now ready to move on to the configuration spaces and moduli spaces of instantons on cobordisms. Let $(W, c) : (Y, w) \rightarrow (Y', w')$ be a cobordism of weakly admissible pairs with $\mathbf{E}_c \rightarrow W$ being the $U(2)$ -bundle associated to the cycle c . Fix a collar neighborhood of ∂W in which the embedding of c is cylindrical. We will want to pass back and forth between W and the manifold

$$W^* = (-\infty, 0] \times Y \cup_Y W \cup_{Y'} [0, \infty) \times Y', \quad (3.7)$$

obtained by adding infinite cylindrical ends to the boundary components. Notice that W^* also comes equipped with an embedded surface c^* , which is cylindrical on the ends. In terms of the bundle \mathbf{E}_c , this means that its restriction to the ends is canonically isomorphic to the pullbacks of the bundles $E_w, E_{w'}$ which correspond to w, w' . For the definition of certain weighted Sobolev spaces below, we also fix a smooth function $\tau : W^* \rightarrow \mathbf{R}$ that is equal to the magnitude of the first component on the first and the third subspaces on the right hand subspace of (3.7).

For the purposes of framings, we usually suppose that Y and Y' are already equipped with basepoints y and y' , which are suppressed from notation. For the purpose of parallel transport, we will also assume that W is equipped with a choice of path $\gamma : [-1, 1] \rightarrow W$ so that $\gamma(-1) = y$ and $\gamma(1) = y'$, so that γ is cylindrical in the given collar neighborhood of the boundary. There is a natural smooth extension to a map $\mathbb{R} \rightarrow W^*$, which we also denote by γ in a slight abuse of notation. We consider γ to be a ‘base line’ on the cobordism, and $b = \gamma(0)$ to be a specific basepoint.

Now suppose that $\alpha \in \tilde{\mathcal{B}}_k(Y, w)$ and $\alpha' \in \tilde{\mathcal{B}}_k(Y', w')$ are fixed orbits of framed connections. Write $\mathcal{A}_k(\alpha)$ and $\mathcal{A}_k(\alpha')$ for the set of connections which lie in that gauge equivalence class. A *path of connections along (W, c) from α to α'* (or *homotopy class of path*), written $z : \alpha \rightarrow \alpha'$ where (W, c) is clear from context, is an element of the set

$$z \in \frac{\mathcal{A}_k(\alpha) \times \mathcal{A}_k(\alpha')}{\mathcal{G}_{k+3/2, \delta}(W, c)};$$

here the quotient is by the group of gauge transformations of (W^*, \mathbf{E}_{c^*}) of Sobolev class $L^2_{k+3/2}$ which exponentially decay to constant (but not necessarily parallel) gauge transformations on the ends, and the action is $\sigma \cdot (A, B) = (\sigma(-\infty) \cdot A, \sigma(+\infty) \cdot B)$. This set is in noncanonical bijection with \mathbb{Z} .

Fix an L^2_k reference connection \mathbf{A}_0 on W which restricts to a connection in α and α' along the corresponding ends; this gives rise to a path of connections $z : \alpha \rightarrow \alpha'$ along (W, c) by considering

the pair $(\mathbf{A}_0(-\infty), \mathbf{A}_0(+\infty)) = (A, A')$. By definition, gauge equivalence class of \mathbf{A}_0 determines a homotopy classes of path z . We then define

$$\mathcal{A}_{k,\delta}(W, c; A, A') = \mathbf{A}_0 + L_{k,\delta}^2\left(\Omega^1(W^*, \text{ad}(\mathbf{E}_c))\right).$$

Here $L_{k,\delta}^2$ is the space of functions of weighted Sobolev space where the weight is given by $e^{-\tau\delta}$ with τ being the function defined above. Later, we will take $\delta > 0$ to be sufficiently small with respect to choices of auxiliary data on ∂W . This only depends on the gauge equivalence classes of A, A' and not the choice of \mathbf{A}_0 .

There is a gauge group $\mathcal{G}_{k+1,\delta}(W, c; A, A')$ given by gauge transformations σ that $d_{\mathbf{A}_0}\sigma$ is of Sobolev class $L_{k,\delta}^2$. In particular, elements of this gauge group are asymptotic to stabilizers of A and A' on the ends. The quotient

$$\mathcal{B}_{k,\delta,z}(W, c; \alpha, \alpha') = \mathcal{A}_k(W, c; A, A')/\mathcal{G}_{k+1}(W, c; A, A')$$

depends only on the path of connections $z : \alpha \rightarrow \alpha'$, and not the choice of \mathbf{A}_0 or of (A, A') . We write $\mathcal{B}(W; \alpha, \alpha') = \bigsqcup_{z:\alpha\rightarrow\alpha'} \mathcal{B}_{k,z}(W, c; \alpha, \alpha')$, suppressing all of c , the Sobolev index k , and the disjoint union over paths from notation unless they are worth drawing attention to.

We may now follow the usual recipe to free up the action. We set

$$\tilde{\mathcal{A}}(W, c; A, A') = \mathcal{A}(W, c; A, A') \times \text{ad}(\mathbf{E})_{\gamma(0)},$$

where the latter term is the space of framings of the adjoint bundle at the basepoint $\gamma(0)$. This gives rise to the quotient $\tilde{\mathcal{B}}(W, c; \alpha, \alpha')$ which carries an action by $SO(3) = PU(2)$. Orbits again come in three types, corresponding to central, abelian, and irreducible connections. There is also an $SO(3)$ -equivariant vector bundle $\mathcal{V}_{k,\delta}^+ \rightarrow \tilde{\mathcal{B}}_k(W, c; \alpha, \alpha')$ defined using $L_{k,\delta}^2$ sections of $\Lambda^+ \otimes \text{ad}(\mathbf{E})$. There are natural completions of this space to vector bundles $\mathcal{V}_{j,\delta}^+$ for all $j \leq k$.

Remark 3.8. The choice of framing $\Phi \in \text{ad}(\mathbf{E})_{\gamma(0)}$ gives rise to a corresponding parallel framing Φ_t above $\text{ad}(\mathbf{E})_{\gamma(t)}$. If one prefers, the entire theory could be set up replacing the choice of framing at a basepoint $\Phi \in \text{ad}(\mathbf{E})_{\gamma(0)}$ with a choice of parallel framing over a base-path, and this perspective may be better equipped to give understanding of the bubbling phenomenon. \diamond

Now we introduce the ASD equations. Suppose that $(Y, w), (Y', w')$ are each equipped with metrics and regular perturbations, that (W, c) is equipped with a metric which is cylindrical on the fixed collar neighborhood of the boundary, and that the orbits α and α' are π, π' -flat connections. At this point, we should ensure that δ is chosen smaller than any $|\lambda|$, where λ is a non-zero eigenvalue of the perturbed extended Hessian operator associated to α and α' .⁴ Choose a smooth function $\beta : \mathbf{R} \rightarrow [0, 1]$ satisfying $\beta(t) = 1$ for $|t| \geq 1$ and $\beta(t) = 0$ for $|t| \leq 1/2$. Then on the two cylindrical ends one may assign to each connection a self-dual 2-form

$$(\widehat{\nabla}_{\pi})(\mathbf{A}) = \beta(t)(dt \wedge \nabla_{\mathbf{A}(t)}\pi)^+, \quad (\widehat{\nabla}_{\pi'})(\mathbf{A}) = \beta(t)(dt \wedge \nabla_{\mathbf{A}(t)}\pi')^+.$$

⁴If one is only interested in setting up a well-behaved Fredholm problem, one merely needs to ensure that δ is disjoint from this spectrum.

Then the map

$$\mathbf{A} \mapsto (F_{\mathbf{A}}^+)^0 - \widehat{\nabla}_{\pi}(\mathbf{A}) - \widehat{\nabla}_{\pi'}(\mathbf{A}) \quad (3.9)$$

descends to a smooth $SO(3)$ -equivariant section

$$F_{\pi, \pi'} : \widetilde{\mathcal{B}}_{k, \delta}(W, c; \alpha, \alpha') \rightarrow \mathcal{V}_{k-1, \delta}^+$$

We will have need to make additional perturbations on W^* . These are given in [Mil19, Definition 4.2], following the presentation in [Kro05]. The explicit form of these perturbations is irrelevant for our purposes, and we largely leave them as a black-box. For technical reasons in the proof of the transversality theorem⁵, and later the gluing process, we will later need to demand that these interior perturbations are chosen sufficiently small. These define smooth $SO(3)$ -equivariant sections of \mathcal{V}_+ with compact derivative. With a slight abuse of notation, we write F_{π} for the sum of the section above and these new perturbations.

We say that zeroes of F_{π} are π -instantons (or simply instantons, for short). Once again, these come in three types: central, abelian, irreducible. We write $M(W, c; \alpha, \alpha')$ for the moduli space of π -instantons. This moduli space is equipped with an $SO(3)$ -equivariant map

$$e_- \times e_+ : M(W, c; \alpha, \alpha') \rightarrow \alpha \times \alpha';$$

the map e_- is defined by sending $[\mathbf{A}, \Phi]$ to $[\mathbf{A}(-\infty), \text{Hol}_{\gamma(0) \rightarrow \gamma(-\infty)}^{\mathbf{A}} \Phi]$, using \mathbf{A} -parallel transport along γ to transport the framing at $\gamma(0) \in W$ to $(-\infty, y_1)$; the map e_+ is defined similarly. By equivariance and the fact that α and α' are orbits, either the left or right endpoint map separately is a submersion. The subspace of $M(W, c; \alpha, \alpha')$ given by instantons that represent a path z is denoted by $M_z(W, c; \alpha, \alpha')$.

We conclude by discussing the notion of the index, and thus expected dimension, of these moduli spaces. As with our discussion of the extended Hessian in the previous section, this discussion is easier before taking the quotient by gauge transformations.

Definition 3.10. If $\mathbf{A} \in \mathcal{A}_{k, \delta}(W, c; A, A')$ is a connection in the homotopy class z of paths between the π -flat connections A and A' at $\pm\infty$, we write $i(z)$ for the index of the elliptic complex

$$\Omega_{k+1, \delta, \text{ext}}^0(W, \text{ad}(\mathbf{E}_c)) \xrightarrow{d_{\mathbf{A}}} \Omega_{k, \delta}^1(W, \text{ad}(\mathbf{E}_c)) \xrightarrow{D_{\mathbf{A}, \pi}} \Omega_{k-1, \delta}^+(W, \text{ad}(\mathbf{E}_c)),$$

where the first term is the space of sections of $\text{ad}(\mathbf{E}_c)$ for which $d_{\mathbf{A}}\sigma \in L_{k, \delta}^2$, and $D_{\mathbf{A}, \pi}$ is the linearization of (3.9) (after adding the secondary perturbation on the interior of W^*). For a component $z : \alpha \rightarrow \alpha'$, write Γ_{α} for the stabilizer of α ; we also define

$$\widetilde{i}(z) = i(z) + \dim \Gamma_{\alpha},$$

the latter being called the *framed index*. ◇

Remark 3.11. Because any two \mathbf{A}_i lying in the same component of connections give rise to homotopic elliptic complexes, and any two gauge-equivalent \mathbf{A}_i give rise to isomorphic elliptic complexes, the quantities $i(z)$ and $\widetilde{i}(z)$ defined above only depend on the path z . ◇

⁵Precisely, see [Mil19, Proposition 4.37]; the argument achieves transversality of a parameterized moduli space along the locus where the part of the perturbation contained in W is zero, and uses this to show that generic **small** choices of interior perturbation are regular. In the gluing process, one must contend with this requirement that the perturbations are small, but this is comparatively less subtle.

We say \mathbf{A} is a regular π -instanton if $D_{\mathbf{A},\pi}$ is surjective. When \mathbf{A} is abelian the ASD operator $D_{\mathbf{A},\pi}$ splits as $D_{\mathbf{A},\pi}^{\text{red}} \oplus D_{\mathbf{A},\pi}^{\nu}$, a reducible part and a part normal to the reducible locus; the latter is called the *normal ASD operator*. In this case, nodedegeneracy of \mathbf{A} splits into two hypotheses: that both $D_{\mathbf{A},\pi}^{\text{red}} = D_{\mathbf{A}}^{\text{red}}$ and $D_{\mathbf{A},\pi}^{\nu}$ are surjective. The first condition is independent of the perturbation because our perturbations vanish along the reducible locus, and always holds by Hodge theory. The second condition is called being *regular normal to the reducible locus*, and is more difficult to achieve. The quantities $i(z)$ and $\tilde{i}(z)$ compute expected dimensions of the moduli spaces of instantons: if $M_z(W, c; \alpha, \alpha')$ consists of regular irreducible instantons, then

$$\begin{aligned} i(z) &= \dim M_z(W, c; \alpha, \alpha') / SO(3), \\ \tilde{i}(z) &= \dim \text{Fiber} (M_z(W, c; \alpha, \alpha') \rightarrow \alpha). \end{aligned}$$

A more detailed discussion can be found in [Mil19, Section 4.2].

The framed index is additive in the following sense. Let $(W_1, c_1) : (Y_1, w_1) \rightarrow (Y_2, w_2)$ and $(W_2, c_2) : (Y_2, w_2) \rightarrow (Y_3, w_3)$ be cobordisms, and suppose $\mathbf{A}_1 : \alpha_1 \rightarrow \alpha_2$ is a connection in the homotopy class z on (W_1, c_1) , while $\mathbf{A}_2 : \alpha_2 \rightarrow \alpha_3$ is a connection in the component w on (W_2, c_2) . Furthermore, suppose that the \mathbf{A}_i are pullback connections on the cylindrical ends and in the specified collar neighborhoods (achievable by a homotopy). Then write $W_{12} = W_1 \cup_{Y_2} W_2$; we may glue the connections \mathbf{A}_i to obtain a connection $\mathbf{A}_{12} : \alpha_1 \rightarrow \alpha_3$ in the homotopy class of the path that we denote by $z * w$. Then we have

$$\tilde{i}(z) + \tilde{i}(w) = \tilde{i}(z * w).$$

Reducing the calculation to $S^1 \times Y$, it follows quickly from the Atiyah-Singer index theorem that $\tilde{i}(\zeta) \in 8\mathbb{Z}$ for any homotopy class $\zeta : \alpha \rightarrow \alpha$; by additivity of index, and the fact $\pi_0 \tilde{\mathcal{B}}(W, c; \alpha, \alpha')$ is affine over $\pi_0 \tilde{\mathcal{B}}(\mathbb{R} \times Y, \mathbb{R} \times w; \alpha, \alpha')$, it follows that the mod 8 value of $i(z)$ is independent of the choice of the path z , and thus similarly with $\tilde{i}(z)$.

For an arbitrary path $z : \alpha \rightarrow \alpha'$, we define

$$\tilde{i}(W, c; \alpha, \alpha') = \tilde{i}(z) \pmod{8}, \quad i(W, c; \alpha, \alpha') = i(z) \pmod{8}.$$

Now suppose (Y, w) is a weakly admissible pair. Given any two null-cobordisms $(W_i, c_i) : (Y, w) \rightarrow (S^3, \emptyset)$, we can glue them together to obtain a cobordism $(\hat{W}, \hat{c}) : S^3 \rightarrow S^3$; we have

$$\tilde{i}(\hat{W}, \hat{c}; \theta, \theta) = \tilde{i}(W_2, c_2; \alpha, \theta) - \tilde{i}(W_1, c_1; \alpha, \theta).$$

The index theorem for the ASD operator over closed manifolds shows that $\tilde{i}(\hat{W}, \hat{c}) = 3(b^1(\hat{W}) - b^+(\hat{W})) + 2[\hat{c}]^2 \pmod{8}$.

Because b^1 and b^+ are additive under gluing along rational homology spheres, this leads to a definition of an absolute $\mathbb{Z}/2$ grading on $\mathfrak{C}(Y, w)$. When $w = \emptyset$ so that we may choose $c = \emptyset$, the term $2[\hat{c}]^2$ vanishes, so we in fact obtain an absolute $\mathbb{Z}/8$ grading.

Definition 3.12. When (Y, w) is a weakly admissible pair, there is an absolute $\mathbb{Z}/2$ grading on $\mathfrak{C}(Y)$: if $(W, c) : (Y, w) \rightarrow (S^3, \emptyset)$ is a null-cobordism of Y and $z : \alpha \rightarrow \theta$ a component of connections on W , we set $\tilde{i}(\alpha) = \tilde{i}(z) - 3(b^1(W) - b^+(W)) \pmod{2\mathbb{Z}}$. If $w = \emptyset$, taking $c = \emptyset$ the same formula defines an absolute $\mathbb{Z}/8$ grading. \diamond

In practice, we will only need instantons to be regular up to a certain index.

Definition 3.13. Let (W, c) be as above. We say that π is a regular perturbation if every π -instanton in a component z with $i(z) \leq 7$ is regular. \diamond

3.3 The case of cylinders

For the cylinder $(W, c) = I \times (Y, w)$, the above discussion becomes significantly simpler. The noncompact manifold W^* carries a canonical diffeomorphism to $\mathbb{R} \times Y$ with $c^* = \mathbb{R} \times w$, so that the associated bundle is the pullback bundle of the bundle $E \rightarrow Y$ associated to w . The path $\gamma : \mathbb{R} \rightarrow W^*$ which we will use for parallel transport is simply given by $\gamma(t) = (t, y)$, where $y \in Y$ is the chosen basepoint.

In this case, we use simpler perturbations than the previous section; we set

$$(\widehat{\nabla}_\pi)(\mathbf{A}) = (dt \wedge \nabla_{\mathbf{A}(t)}\pi)^+$$

with no bump function term, and we use no further perturbations. In particular, the manifold $\mathbb{R} \times Y$ carries a natural \mathbb{R} -action via $\phi_t(s, z) = (t + s, z)$. This induces a natural \mathbb{R} -action on $\widetilde{\mathcal{B}}(\mathbb{R} \times Y, \mathbb{R} \times c; \alpha, \beta)$ by pullback and parallel transport, with

$$\tau_t[\mathbf{A}, \Phi] = [\phi_t^* \mathbf{A}, \text{Hol}_{(0,y) \rightarrow (-t,y)} \Phi].$$

The endpoint maps to α and β are invariant under this action, as is the section F_π itself, as both metric and perturbation are translation-invariant.

It follows that the moduli space $M(\mathbb{R} \times Y; \alpha, \beta)$ carries a natural \mathbb{R} -action for which the endpoint maps are invariant. We write the quotient as

$$\check{M}(Y; \alpha, \beta) = M(\mathbb{R} \times Y; \alpha, \beta)/\mathbb{R};$$

then $\check{M}(Y; \alpha, \beta)$ is an $SO(3)$ -space with equivariant endpoint maps to α and β .

The set of components

$$\pi_0 \mathcal{B}(\mathbb{R} \times Y, \mathbb{R} \times w; \alpha, \beta)$$

is in canonical bijection with $\pi_0 \mathcal{G}(Y, w) \cong \mathbb{Z}$. When $M_\zeta(\mathbb{R} \times Y; \alpha, \beta)$ consists of nonconstant instantons — that is, either $\alpha \neq \beta$ or at least ζ is not the component of the constant trajectory — the \mathbb{R} -action is free.

When $\zeta : \alpha \rightarrow \beta$ is a given trajectory, the quantity $i(\zeta)$ depends modulo 8 only on α and β . We define

$$i(Y, w; \alpha, \beta) := i(\zeta) \pmod{8}, \quad \text{and similarly} \quad \widetilde{i}(Y, w; \alpha, \beta) := \widetilde{i}(\zeta) \pmod{8}.$$

We say that a perturbation π on (Y, w) is regular if all instantons with $i(\zeta) \leq 8$ are regular. Regular perturbations exist and in fact are open and dense⁶ in the space of all perturbations on (Y, w) by [Mil19, Theorem 4.34]. It follows by Proposition 3.4 that there exist regular perturbations on (Y, w) with arbitrary signature data (first choose a nice perturbation with desired signature data, then choose a nearby regular perturbation).

We can use this quantity to obtain an absolute grading on the set of π -flat connections.

⁶If one wanted to achieve transversality for all instantons, not merely those of index bounded above by a fixed quantity, then this would rather be a countable intersection of open and dense sets

Definition 3.14. When Y is a rational homology sphere, the set $\mathfrak{C}(Y, \emptyset)$ carries a natural $\mathbb{Z}/8$ -grading $\tilde{i}(\alpha) = \tilde{i}(Y; \alpha, \theta)$, where θ is the trivial connection. \diamond

This lifts the grading \tilde{i} of Definition 3.12 and has the property that

$$\tilde{i}(Y; \alpha, \beta) = \tilde{i}(\alpha) - \tilde{i}(\beta).$$

3.4 Compactifications of the moduli spaces

We move on to discussing the (framed) Uhlenbeck compactification of our moduli spaces. These are only well-behaved in a small dimension range (sufficient for our purposes); even then, the result is not quite a smooth manifold. We will define the compactifications and describe the extra structure they carry.

For the remainder of this section, we will abuse notation and write \mathbf{A} to denote an equivalence class of framed connection (so, properly, an equivalence class $[\mathbf{A}, \Phi]$, where Φ is a framing above a chosen basepoint). We suppose that the weakly admissible pair (Y, w) has been equipped with a regular perturbation π and a basepoint y , so that the framing is chosen at y .

Definition 3.15. We say that a broken framed π -instanton on (Y, w) between α and β is a sequence $\mathbf{A}_i \in \check{M}(Y, w; \alpha_i, \alpha_{i+1})$ of nonconstant framed π -instantons modulo translation for $0 \leq i \leq n$, with $\alpha_0 = \alpha$ and $\alpha_{n+1} = \beta$, so that

$$e_+(\mathbf{A}_i) = e_-(\mathbf{A}_{i+1})$$

for $0 \leq i \leq n$; that is, the framings agree when parallel transported to each intermediate critical orbit. If \mathbf{A}_i lies in the homotopy class of path $\zeta_i : \alpha_i \rightarrow \alpha_{i+1}$, we say that the homotopy class of the broken framed instanton \mathbf{A} is $\zeta = \zeta_1 * \cdots * \zeta_n$.

We write $\check{M}_\zeta^+(Y, w; \alpha, \beta)$ for the set of broken framed π -instantons running from α to β in the homotopy class ζ . \diamond

Since (Y, w) is fixed, we suppress them from notation for the rest of this section.

The moduli space $\check{M}_\zeta^+(\alpha, \beta)$ carries the structure of a stratified set, with strata corresponding bijectively to sequences

$$\alpha \xrightarrow{\zeta_0} \alpha_1 \xrightarrow{\zeta_1} \cdots \xrightarrow{\zeta_{n-1}} \alpha_n \xrightarrow{\zeta_n} \beta;$$

the corresponding stratum is given by the iterated fiber product

$$\check{M}_{\zeta_0}(\alpha, \alpha_1) \times_{\alpha_1} \cdots \times_{\alpha_n} \check{M}_{\zeta_n}(\alpha_n, \beta).$$

The moduli space $\check{M}_\zeta^+(\alpha, \beta)$ has a natural topology so that each stratum has the fiber product topology, given by the notion of *chain-convergence*. We say that a sequence (\mathbf{A}_k) of framed instantons for $k = 1, 2, \dots$ is *chain-convergent* to a broken trajectory $(\mathbf{B}_i)_{1 \leq i \leq n}$ if there is a sequence $t_k = (t_k^1, \dots, t_k^n) \in \mathbb{R}^n$ with each $\lim_{k \rightarrow \infty} t_k^i - t_k^{i-1} = \infty$ and so that the translates $\tau_{t_k^i}(\mathbf{A}_k)$ converge in the appropriate Sobolev topology on compact sets to \mathbf{B}_i .

There is a straightforward extension of this notion to chain-convergence of broken framed instantons (possibly with different numbers of components). If \mathbf{A}_n is a sequence of broken instantons in the homotopy class ζ which chain-converge to \mathbf{A} , then \mathbf{A} also lies in the homotopy class ζ .

We have the following theorem on the structure of this topological space; this is stated in slightly different form as [Mil19, Proposition 4.43]. The result can be straightforwardly obtained by either the implicit function theorem methods of [KM07, Chapter 19] (the method of proof preferred in [Mil19]), or by the Banach fixed-point methods of [DK90, Chapter 7.2], which we find more convenient when studying obstructed gluing theory in Section 5 below.

Proposition 3.16. *If $\zeta : \alpha \rightarrow \beta$ is a homotopy class with $i(\zeta) \leq 8$, then $\check{M}_\zeta^+(\alpha, \beta)$ is compact. Further, whenever ζ_0, \dots, ζ_n is a sequence of homotopy classes between the critical points whose concatenation is ζ , there are continuous $SO(3)$ -equivariant gluing maps*

$$(L, \infty]^n \times \check{M}_{\zeta_1}(\alpha, \alpha_1) \times_{\alpha_1} \cdots \times_{\alpha_n} \check{M}_{\zeta_n}(\alpha_n, \beta) \rightarrow \check{M}_\zeta^+(\alpha, \beta)$$

which are open embeddings onto a neighborhood of the corresponding stratum; these maps are local diffeomorphisms on each stratum separately.

In particular, when $i(\zeta) \leq 8$, the moduli space of possibly broken framed instantons \check{M}_ζ^+ is a compact topological manifold equipped with a stratification and a smooth structure on each stratum, as well as gluing maps giving continuous charts near the corners points which are smooth on each stratum. This is nearly the structure of a smooth manifold, but not quite, and we will need to take special care when taking transverse intersections or regular values of maps defined on such objects. Before discussing that, we record the structure on \check{M}_ζ^+ in the following definition.

Definition 3.17. A stratified-smooth manifold of dimension n consists of the following data:

- (i) A partially ordered set Δ with an order-preserving *dimension* function $d : \Delta \rightarrow \mathbb{Z}_{\geq 0}$ with the maximum value n ,
- (ii) A topological space P stratified by subsets P_s for $s \in \Delta$, with $\overline{P_s} = \bigcup_{t < s} P_t$,
- (iii) A smooth structure on each stratum P_s that turns it into a smooth manifold of dimension $d(s)$.

We demand that for each $x \in P_s$, there exists an open, stratum-preserving embedding

$$\phi : [0, \varepsilon]^{n-d(s)} \times \mathbb{R}^{d(s)} \rightarrow P,$$

with $\phi(0) = x$, so that ϕ induces an isomorphism of posets $\{0, 1\}^{n-d(s)} \rightarrow \Delta_{\geq s}$, and so that ϕ is a local diffeomorphism on each stratum.

If M is a smooth manifold and P is a stratified-smooth manifold, we say $f : P \rightarrow M$ is stratified-smooth if it is continuous and smooth on each stratum. We say that f is stratified-transverse to $S \subset M$ if $f|_{P_s}$ is transverse to S for all $s \in \Delta$. \diamond

In particular, a stratified-smooth manifold is a topological manifold with corners. Proposition 3.16 above asserts that $\check{M}_\zeta^+(\alpha, \beta)$ is a stratified-smooth $SO(3)$ -manifold, where Δ is the set of sequences

$(\zeta_0, \dots, \zeta_n)$ of homotopy classes of paths between various critical orbits $\zeta_j : \alpha_j \rightarrow \alpha_{j+1}$, with $\alpha_0 = \alpha$ while $\alpha_{n+1} = \beta$ and $\zeta_0 * \dots * \zeta_k = \zeta$. We say $t < s$ if s is obtained from t by concatenating some of its constituent homotopy classes.

The stratum corresponding to $s = (\zeta_0, \dots, \zeta_k)$ is

$$\check{M}_{\zeta_0}(\alpha, \alpha_1) \times_{\alpha_1} \dots \times_{\alpha_n} \check{M}_{\zeta_k}(\alpha_k, \beta),$$

with $d(s) = \sum_{j=0}^n (\tilde{i}(\zeta_j) - 1) + \dim \alpha$.

In our work later, we will need to look at transverse intersections of a stratified-transverse map with a smooth submanifold $S \subset M$. Unfortunately, such transverse intersections are *no longer* stratified-smooth manifolds; they need not even be topological manifolds in general. Nevertheless, the resulting structure (called a *stratified-smooth space*) is sufficient for all of our purposes. In the definition of the instanton chain complex, we will use a chain complex which calculates singular homology, whose generators are stratified-smooth spaces up to a certain notion of equivalence. We rapidly develop the foundations of stratified-smooth spaces in Appendix A. These are similar to, but simpler than, the notion of δ -chain; this notion has appeared in a handful of references, including [KM07, Lin18, Mil19].

Remark 3.18. The notation chosen here differs from the notation chosen for the same spaces in [Mil19], where $\check{M}_{\zeta}^+(\alpha, \beta)$ would have been written as $\overline{\mathcal{M}}_{E,z,k,\delta}(\alpha, \beta)$. We prefer the current notation both for brevity and for consistency with previous literature. \diamond

We now move on to a similar discussion for general cobordisms. Suppose $(W, c) : (Y, w) \rightarrow (Y', w')$ is equipped with a regular perturbation π_W , extending choices of regular perturbations π and π' on the ends, as well as a path γ from $y \in Y$ to $y' \in Y'$; we choose the framing in W above $\gamma(0)$, and parallel transport it along the path γ as necessary.

Definition 3.19. We say that a broken framed instanton on (W, c) between α and α' is a triple $(\mathbf{A}_i, \mathbf{B}, \mathbf{A}'_j)$:

- A (possibly empty) sequence $\mathbf{A}_i \in \check{M}_{\zeta_i}(\alpha_i, \alpha_{i+1})$ of nonconstant framed π -instantons (modulo translation) for $0 \leq i \leq n$, with $\alpha_0 = \alpha$ so that

$$e_+(\mathbf{A}_i) = e_-(\mathbf{A}_{i+1})$$

for $0 \leq i < n$;

- A framed π_W -instanton $\mathbf{B} \in M_z(W, c; \alpha_{n+1}, \beta_0)$, with $e_-(\mathbf{B}) = e_+(\mathbf{A}_n) \in \alpha_{n+1}$;
- A (possibly empty) sequence $\mathbf{A}'_i \in \check{M}_{\zeta'_i}(\alpha'_i, \alpha'_{i+1})$ of nonconstant framed π' -instantons (modulo translation) for $0 \leq i \leq m$, with $\alpha'_{m+1} = \alpha'$, so that

$$e_+(\mathbf{A}'_i) = e_-(\mathbf{A}'_{i+1})$$

for $0 \leq i < m$, and $e_-(\mathbf{A}'_0) = e_+(\mathbf{B}) \in \alpha'_0$.

We say this broken framed instanton represents the path of connections along (W, c) given by $\zeta_1 * \dots * \zeta_n * z * \zeta'_1 * \dots * \zeta'_m$.

We write $M_z^+(W, c; \alpha, \alpha')$ for the set of broken framed instantons on (W, c) running from α to α' in the homotopy class z . \diamond

Once again, this is a stratified set, stratified by iterated fiber products of the relevant moduli spaces; thus all strata inherit a natural smooth structure. The topology is similar to the case of a cylinder: we say that a sequence \mathbf{A}_k of framed instantons on (W, c) is chain-convergent to a broken framed instanton $(\mathbf{B}_i, \mathbf{A}, \mathbf{C}_j)$ if there are sequences $t_k = (t_k^1, \dots, t_k^n) \in \mathbb{R}_+^n$ and $s_k = (s_k^1, \dots, s_k^m) \in \mathbb{R}_-^m$ whose sequential differences go to $+\infty$ and $-\infty$, respectively, so that the framed instantons \mathbf{A}_k converge to \mathbf{A} on compact sets, and so that upon restricting \mathbf{A} to the respective ends, $\tau_{t_k^i}(\mathbf{A})$ converges on compact subsets⁷ of $\mathbb{R} \times Y_1$ to \mathbf{B}_i while $\tau_{s_k^j}(\mathbf{A}_k)$ converges on compact subsets of $\mathbb{R} \times Y_2$ to \mathbf{C}_j (both modulo gauge, and everything in sight equipped with a framing; in restricting to the ends we parallel transport the framing along a base-path γ to the basepoint on the end).

Then for $i(z) \leq 7$ there once again exist $SO(3)$ -equivariant gluing maps, stated explicitly in [Mil19, Proposition 4.42]. We record this as a proposition below.

Proposition 3.20. *Suppose (W, c) is a cobordism equipped with a regular perturbation. If $i(z) \leq 7$, then $M_z^+(W, c; \alpha, \alpha')$ is compact with the topology induced by chain-convergence. Furthermore, with the stratification described above, $M_z^+(W; \alpha, \beta)$ is equipped with the structure of a stratified-smooth manifold.*

3.5 Reducibles and regularity

We would like to construct regular perturbations π_W on cobordisms $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$, where π, π' are fixed regular perturbations on the boundary components. Unfortunately, this is not always possible, and in this section we will discuss the obstruction.

Achieving regularity at irreducible connections is standard. A generic perturbation has all irreducible instantons cut out regularly; one may see this by the arguments of [Don02, Section 5.5.1] except for those irreducible flat connections which are central on the ends, where one needs more care (and to introduce interior perturbations). It is more difficult to ensure regularity at the reducibles. If one of the boundary components is admissible, there are no reducible instantons whatsoever, so we may as well assume that both boundary components are rational homology spheres.

Convention 3.21. Unless otherwise specified, for the duration of this section, $(W, c, \pi_W) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is a cobordism between rational homology spheres equipped with regular perturbations π, π' , while π_W is a perturbation on the cobordism extending the chosen perturbations on the ends. The signature data functions of these perturbations are written σ_π and $\sigma_{\pi'}$, respectively. \diamond

First we recall the enumeration of reducible instantons on a cobordism $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$. Consider the set of components of the space of reducible connections on (W, c) ; we write $\mathfrak{A}(W, c)$ for the set of components which include no central connections, and $\mathfrak{Z}(W, c)$ for the set of components which include some central connection. The following enumeration of these components given in [Mil19, Proposition 1.7].

⁷Notice that the assumptions on t and s imply that for sufficiently large k , the domain of definition of these translates contains any particular compact subset of $\mathbb{R} \times Y$ or $\mathbb{R} \times Y'$.

Proposition 3.22. *Let $(W, c, \pi_W) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ be a cobordism between rational homology spheres. Then there are canonical bijections*

$$\begin{aligned}\mathfrak{Z}(W, c) &\cong \{x \in H^2(W; \mathbb{Z}) \mid 2x = \text{PD}(c)\}, \\ \mathfrak{A}(W, c) &\cong \{\{x, y\} \subset H^2(W; \mathbb{Z}) \mid x + y = \text{PD}(c), \ x \neq y\}.\end{aligned}$$

The central connections are regular if and only if $b^+(W) = 0$. As discussed in Section 3.2, for an abelian instanton Λ to be regular, its *normal ASD operator* D'_{Λ, π_W} must be surjective. This is a complex linear operator, so its index is an even integer; by homotopy invariance of index, this quantity depends only on Λ and the choice of π, π' on the ends.

Definition 3.23. Suppose Λ is an abelian connection on the cobordism (W, c, π_W) . We write $N(\Lambda; \pi, \pi')$ for the index of the normal ASD operator D'_{Λ, π_W} . \diamond

If Λ is to be cut out transversely, we need at the very least that $N(\Lambda; \pi, \pi') \geq 0$. (We may need more than this, as we will discuss later.) To make this discussion more concrete, we can compute these indices with the Atiyah–Patodi–Singer index theorem.

To state the calculation, we set some notation. For a cohomology class $x \in H^2(W)$, write $x_i \in H^2(Y_i)$ for its restriction to each boundary component. Notice that because the boundary components are rational homology spheres, the map $j : H_c^2(W; \mathbb{Q}) \rightarrow H^2(W; \mathbb{Q})$ is an isomorphism. Whenever $x \in H^2(W; \mathbb{Z})$ we write $x^2 \in \mathbb{Q}$ to mean the rational number obtained as

$$x^2 = \langle (j^{-1}x)^2, [W] \rangle,$$

where $[W] \in H_4^{\text{lf}}(W; \mathbb{Q})$ is the fundamental class in locally finite homology.

The following calculation is the main result of [Mil19, Section 4.5].

Proposition 3.24. *If $\Lambda : \alpha \rightarrow \alpha'$ is an abelian instanton in the component labeled by $\{x, y\}$, then $N(\Lambda; \pi, \pi')$ is the even integer*

$$\begin{aligned}N(\Lambda; \pi, \pi') &= -2(x - y)^2 + 2(b_1(W) - b^+(W)) \\ &\quad + \frac{\sigma_{\pi'}(\alpha') + \rho(\text{ad}_{\alpha'})}{2} - \frac{\sigma_{\pi}(\alpha) + \rho(\text{ad}_{\alpha})}{2} \\ &\quad + 1 - \frac{r(\alpha) + r(\alpha')}{2},\end{aligned}\tag{3.25}$$

where $\rho(\text{ad}_{\alpha})$ is the Atiyah–Patodi–Singer rho invariant of the extended Hessian operator $\widehat{\text{Hess}}_{\alpha}(CS)$, $r(\alpha) = 3$ if α is a central connection and $r(\alpha) = 1$ if α is abelian. Similar comments apply to $\rho(\text{ad}_{\alpha'})$ and $r(\alpha')$.

The quantity $-2(x - y)^2$ may be understood as eight times the *topological energy*

$$\left(c_2(\mathbf{E}_c) - \frac{1}{4}c_1(\mathbf{E}_c)^2 \right) = \frac{1}{8\pi^2} \int_W \text{tr} F_{\Lambda}^2,$$

where \mathbf{E}_c is the bundle associated to the geometric representative c .

The expression in the second line of (3.25) can be expressed in terms of the rho invariant of the perturbed extended Hessian $\widehat{\text{Hess}}_\alpha(CS + \pi)$, and the corresponding ρ -invariant for α' ; while this is simpler than the above expression, the statement of the proposition makes clear that π enters in the form of its signature data. Still, we record this perturbed ρ invariant as a definition.

Definition 3.26. Let (Y, w, π) be a rational homology sphere with an oriented 1-cycle and a nice perturbation. Then the perturbed rho invariant $\rho_\pi : \mathfrak{A}(Y, w) \rightarrow \mathbb{Q}$ is defined as

$$\rho_\pi(\alpha) = \rho(\text{ad}_\alpha) + \sigma_\pi(\alpha). \quad \diamond$$

Remark 3.27. In Section 7, we will use ρ_π to give an invariant of rational homology spheres which conjecturally computes twice the Casson-Walker invariant, in the spirit of Taubes' interpretation of the Casson invariant [Tau90] as a signed count of irreducible flat connections. For rational homology spheres, one must use the perturbed rho invariant as a correction term, analogous to Walker's definition [Wal92] of his invariant. \diamond

Next, we seek to find perturbations π_W so that all instantons are regular. To guarantee that this is possible, we will need to make assumptions on the normal index. Later in this article we will seek to weaken these assumptions in two different ways, which are both recorded in the definition below.

Definition 3.28. Suppose Λ is a component of abelian connections on $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ corresponding to the pair of cohomology classes $\{x, y\}$.

- We say Λ is unobstructed if either $-2(x - y)^2 < 0$ or $b_1(W) - b^+(W) \leq N(\Lambda; \pi, \pi')$; otherwise we say Λ is obstructed.
- We say Λ is pseudocentral if Λ is central on both ends and has $(x - y)^2 = 0$.

Quantifying over all components, there are three types of cobordism of interest to us.

- If every component of abelian connections over (W, c) is unobstructed, and either $b^+(W) = 0$ or (W, c) supports no central connections, then we say (W, c) is an unobstructed cobordism.
- If every component of abelian connections over (W, c) is either unobstructed or pseudocentral, and either $b^+(W) = 0$ or (W, c) supports no central connections, we say that (W, c) is a pseudo-unobstructed cobordism.
- If $b_1(W) = b^+(W) = 0$ and all but one component of abelian connections over (W, c) is unobstructed, while the one obstructed reducible is abelian on both ends and has $N(\Lambda; \pi, \pi') = -2$, we say that (W, c) is a nearly unobstructed cobordism. \diamond

Remark 3.29. The condition on $b^+(W)$ and central connections is to ensure that any central connection (which necessarily remains an instanton after perturbation) are automatically cut out transversely, so that we can focus our analysis on the more intricate case of abelian reducibles. If the perturbation π_W has

small interior part, a π_W -instanton necessarily has $\int \text{tr}(F_\Lambda^2) > -\varepsilon$, so that $-2(x - y)^2 \geq 0$, which is why we ignore those components with $-2(x - y)^2 < 0$.

We will see in Section 4.3 that unobstructed cobordisms induce chain maps between the tilde complexes.

A pseudocentral connection should be understood to be as close to central as an abelian connection can possibly be; a central connection on the cobordism restricts to central connections on the ends, and is also automatically flat, so that the quantity $(x - y)^2$ is automatically zero. We will see in Section 6.1 that pseudo-unobstructed cobordisms also induce maps on the tilde complexes, for which each pseudocentral connection contributes as if it were two central connections.

The last type of cobordisms — nearly unobstructed cobordisms — are analyzed in detail in Section 5, where we show that they also induce maps on instanton homology, but not between the tilde complexes (one must pass to a ‘suspension’). Using this, we will show that the only type of obstructed reducible we actually have to deal with are the pseudocentral connections. \diamond

It follows immediately from Proposition 3.24 and the fact that $\rho(\theta) = \sigma_\pi(\theta) = 0$ for central connections θ that the pseudocentral connections have particularly simple normal index.

Corollary 3.30. *If Λ is a pseudocentral connection on the cobordism $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$, then $N(\Lambda; \pi, \pi') = 2(b_1(W) - b^+(W) - 1)$.*

For each type of cobordism named in Definition 3.28, we can find special types of perturbations: either we can find regular perturbations, or at least perturbations which are sufficiently regular for the analysis we need to carry out.

Theorem 3.31. *Suppose we are given a cobordism $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$.*

- (i) *If (W, c) is unobstructed, there exists a regular perturbation π_W on (W, c) extending π and π' ; the part of π_W supported in the interior of W may be taken to be arbitrarily small. Given any two regular perturbations π_W with the same restrictions to the ends, there is a regular family $\pi_W(t)$ of perturbations interpolating between them.*
- (ii) *If (W, c) is pseudo-unobstructed, there exists a regular perturbation π_W on (W, c) extending π, π' with arbitrarily small interior part so that all instantons are regular, except possibly for the pseudocentral connections; the same is true for 1-parameter families of perturbations.*
- (iii) *If (W, c) is nearly unobstructed, there exist perturbations π_W with arbitrarily small interior part, so that all irreducible instantons are regular, all reducible instantons Λ with $N(\Lambda; \pi, \pi') \geq 0$ are regular, and so that for the reducible instanton with $N(\Lambda; \pi, \pi') = -2$, we have $\ker D_{\Lambda, \pi_W}^\nu = 0$; the same is true for 1-parameter families of perturbations. We call such perturbations nearly regular.*

Sketch of proof. First, fix a perturbation π_W with small interior part for which abelian instantons are cut out transversely within the reducible locus; arguing along the lines of [AB96, Section 3.3.1], one may guarantee that the central connections are isolated among reducible π_W -instantons (and that there are

exactly $2^{b_1(W)}$ of them). In doing so, we can focus on the reducible components which contain no central connections (henceforth, abelian components). A π_W -instanton has $\int \text{tr}(F_\Lambda^2) \approx \|F_{\mathbf{A}}\|^2$ (the error term coming from the perturbation on the interior; recall that the perturbations on the ends are assumed to vanish on the reducible locus itself), so that the topological energy of an instanton is nearly non-negative; because $-2(x - y)^2$ lies in $\frac{2}{|H_1(\partial W)|}$, so long as the interior part is taken sufficiently small, the quantity $-2(x - y)^2$ is non-negative for any abelian π_W -instanton. So we may focus our attention on those components with $-2(x - y)^2 \geq 0$.

Generically, the space of reducible π_W -instantons in each abelian component has dimension $b_1(W) - b^+(W)$; if $\pi_W(t)$ is a one-parameter family of perturbations, the space of abelian connections which are $\pi_W(t)$ -instantons for some t generically has dimension $b_1(W) - b^+(W) + 1$. Now

$$\Lambda \mapsto D'_{\Lambda, \pi_W}$$

defines a map from this space of abelian instantons to the space of complex linear Fredholm operators of index $N = N(\Lambda; \pi, \pi')$, and we want this map to miss the subspace of operators with either nontrivial kernel or nontrivial cokernel, depending on whether N is negative or non-negative. By [Mil19, Lemma 4.35], the space of operators with non-trivial kernel and cokernel has codimension $|N| + 2$, and the argument of [Mil19, Theorem 4.37] shows that for π_W sufficiently small, the map from the space of pairs (\mathbf{A}, π_W) of perturbations and abelian π_W -instantons is transverse to this locus.

Thus for generic small π_W and generic small 1-parameter families $\pi_W(t)$, we can avoid the locus of operators with nontrivial kernel and cokernel so long as $b_1(W) - b^+(W) + 1 < |N| + 2$, or equivalently, $b_1(W) - b^+(W) \leq |N|$. This is guaranteed for unobstructed and nearly unobstructed cobordisms by definition; further, in the unobstructed case, either $b_1(W) - b^+(W) \geq -1$ so that $N \geq 0$ and all normal ASD operators have trivial cokernel, or $b_1(W) - b^+(W) \leq -2$ and reducibles do not generically appear even in one-parameter families. The same argument applies in the case of pseudo-unobstructed cobordisms, except that one ignores the pseudocentral components. \square

We conclude this section with a more detailed discussion of pseudocentral connections. First, we can obtain more precise control over the perturbations on the pseudocentral connections than described above.

Proposition 3.32. *If (W, c) is a pseudo-unobstructed cobordism, we can obtain a perturbation π_W (and paths of perturbations between them) which is regular at all instantons except possibly the pseudocentral instantons, with the following properties at the pseudocentral instantons. The behavior depends on the quantity $c = b_1(W) - b^+(W)$.*

- (i) *If $c < -1$, then $b^+(W) > 1$, so that for a generic perturbation (and a generic path of perturbations), there are no abelian instantons whatsoever.*
- (ii) *If $c = -1$, then $b^+(W) > 0$, so that for a generic perturbation there are no abelian instantons whatsoever; for a generic path of perturbations, there will generically be an isolated set of pseudocentral instantons, which are such that the normal Fredholm operator has trivial kernel.*
- (iii) *If $c = 0$, then for a generic perturbation there are finitely many pseudocentral instantons, which may be assumed to have normal Fredholm operator with trivial kernel; the same is true in families.*

- (iv) If $c = 1$, then for a generic perturbation there is a 1-dimensional family of pseudocentral instantons, which are all regular; however, for a generic 1-parameter family of perturbations, there is a 2-dimensional family of pseudocentral instantons,
- (v) If $c > 1$, for a generic perturbation and path of perturbations, all pseudocentral instantons are regular.

In any of these cases, we call such a perturbation pseudo-regular.

Proof. The condition needed to ensure that all abelian instantons are cut out with trivial kernel or cokernel is $c \leq |N| + 1$, whereas the condition needed to ensure that all abelian instantons are cut out transversely in 1-parameter families is $c \leq |N|$. Corollary 3.30 asserts that $N = 2c - 2$. The inequality $c \leq |2c - 2| + 1$ holds for all c , whereas the inequality $c \leq |2c - 2|$ holds for all integers $c \neq 1$.

So for $c \neq 1$ generically one of the kernel and cokernel will be trivial for all pseudocentral instantons appearing in 1-parameter families. As for which it is, the cokernel is trivial if $N \geq 0$, while the kernel is trivial if $N \leq 0$; in terms of c , the former happens when $c \geq 1$ and the latter when $c \leq 1$. This gives the pentachotomy above. \square

As mentioned earlier, we will show in Section 6.1 that all pseudo-unobstructed cobordisms define well-defined induced maps on tilde complexes. Cases (i) and (v) require no further care; case (iv) will be shown irrelevant by index considerations; cases (ii)-(iii) will require the most work, but fall to the techniques of Section 5.

We conclude this section by pointing out that if one is free to change the perturbations on the ends, every cobordism is pseudounobstructed. We will use this in Section 6.1 to define induced maps on equivariant instanton homology for any cobordism (W, c) for which $b^+(W) = 0$ or (W, c) supports no central connections. We phrase the statement below slightly more generally (including the case when (Y, w) is non-trivial admissible, where it is tautologically true) for ease of reference.

Lemma 3.33. *Let $(W, c) : (Y, w) \rightarrow (Y', w')$ be a cobordism between weakly admissible pairs. Then there exist regular perturbations π, π' on (Y, w) and (Y', w') respectively so that $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is a pseudo-unobstructed cobordism.*

Proof. We saw in Proposition 3.4 that there exist regular perturbations with arbitrary signature data. Because the normal index only depends on the signature data on the ends, this is really a lemma about signature data functions. Recall that the set of signature data functions on Y is affine over the abelian group $F_{Y,w} = \text{Map}(\mathfrak{A}(Y, w), 4\mathbb{Z})$. Pick arbitrarily signature data functions σ_0, σ'_0 on Y and Y' . Write $r : \mathfrak{A}(W, c) \rightarrow \mathfrak{A}(Y, w)$ for the restriction and similarly for r' .

The normal index function computed in Proposition 3.24 may be considered as a function $N : \mathfrak{A}(W, \mathbf{E}) \times F_{Y,w} \times F_{Y',w'} \rightarrow 2\mathbb{Z}$, with

$$N(\Lambda, f, f') = N(\Lambda; \sigma_0 + f, \sigma'_0 + f') = N(\Lambda; \sigma_0, \sigma'_0) + \frac{f'(r'(\Lambda)) - f(r(\Lambda))}{2}.$$

Here, if $r(\Lambda)$ or $r'(\Lambda)$ are central, we interpret $f(r(\Lambda)) = 0$ and similarly $f'(r'(\Lambda)) = 0$.

Notice that $N(\Lambda, 0, 0)$ takes on a set of values which is bounded below, because for fixed π, π' , the formula given in Proposition 3.24 has only one unbounded term: the first one. (The sets $r\mathfrak{A}(W, c)$ and $r'\mathfrak{A}(W, c)$ are subsets of finite sets, so the latter terms have only finitely many output values.) The first term is always non-negative for instantons Λ . As a result, there is some nonnegative integer n so that $N(\Lambda, 0, 0) - (b_1(W) - b^+(W)) \geq -2n$ for all components of reducible connections Λ which might possibly contain an instanton.

So set $f(\alpha) = -4n$ for all $\alpha \in \mathfrak{A}(Y, w)$ and $f'(\alpha') = 4n$ for all $\alpha' \in \mathfrak{A}(Y', w')$. If Λ is not a pseudocentral reducible, then one of $r(\Lambda)$ or $r'(\Lambda)$ is non-central, $f'(r'(\Lambda)) - f(r(\Lambda))$ is either $4n$ or $8n$. As a result, $N(\Lambda, f, f') \geq (b_1(W) - b^+(W))$ for all reducibles other than pseudocentral connections.

Therefore $\sigma' = \sigma'_0 + f$ and $\sigma = \sigma_0 + f$ satisfy the desired hypotheses. \square

3.6 Families of metrics

The previous discussion will be enough to show that each (Y, w, π) gives rise to an instanton chain complex, and each unobstructed cobordism with regular perturbation (W, c, π_W) gives rise to a chain map between the instanton complexes. Even ignoring the ‘unobstructed’ caveat, this is not adequate. First, we want to ensure that the induced map is independent of the perturbation π on the interior up to chain-homotopy. Second, we want to ensure that the induced maps compose correctly.

We will need additional moduli spaces and structure theorems for them to prove these claims. Instead of dealing with the two issues separately, we introduce families of possibly *broken* metrics to handle both cases at once. Because we will be focused on independence and we are not worried about our homotopies being well-defined up to homotopy, we only discuss the case of metrics and perturbations parameterized by the interval, which we choose to write as $I = [0, \infty]$. If one chooses any regular metric and perturbation $g(0), \pi(0)$ and $g(\infty), \pi(\infty)$ with the interior part of the perturbation sufficiently small, there is always a regular family of metrics and perturbations $g(t), \pi(t)$ extending these (regular in the sense that a certain parameterized derivative operator is surjective for each $\pi(t)$ -instanton \mathbf{A} ; it need not be the case that each $\pi(t)$ is a regular perturbation).

Suppose

$$(W, c) : (Y_0, w_0) \rightarrow (Y_1, w_1), \quad (W', c') : (Y_1, w_1) \rightarrow (Y_2, w_2)$$

are cobordisms. We may define a family of composites

$$W \cup_L W' = W \cup_{Y_1} [-L, L] \times Y_1 \cup W'$$

where $L \in [0, \infty)$ is a length parameter. The resulting manifold has a natural smooth structure, metric, and perturbation, which agrees on the left end with that of W and on the right end with that of W' . So long as L is sufficiently large, if π and π' are regular perturbations, so is the composite perturbation $\pi \cup_L \pi'$.

We model the case $L = \infty$ slightly differently, as $W \cup_{Y_1} [0, \infty) \times Y_1 \sqcup (-\infty, 0] \times Y_1 \cup_{Y_1} W'$, and we call this a broken metric on the composite $W \cup_{Y_1} W'$. In this way the family of composites above defines a smooth family of (possibly broken) cylindrical end metrics on the composite, indexed by the

compact interval $[0, \infty]$. We say that a $\pi(\infty)$ -instanton from α_0 to α_2 is an element of

$$\bigsqcup_{\alpha_0 \xrightarrow{z} \alpha_1 \xrightarrow{z'} \alpha_2} M_z(W, c; \alpha_0, \alpha_1) \times_{\alpha_1} M_{z'}(W', c'; \alpha_1, \alpha_2),$$

and its homotopy class is $z * z'$. More generally, a broken $\pi(\infty)$ -instanton consists of a sequence $(\mathbf{A}_i^0, \mathbf{A}_W, \mathbf{A}_j^1, \mathbf{A}_{W'}, \mathbf{A}_k^2)$, where \mathbf{A}_i^j are broken framed instantons on (Y_j, w_j) for $1 \leq i \leq n_j$ (possibly $n_j = 0$), and \mathbf{A}_W and $\mathbf{A}_{W'}$ are instantons on W and W' , respectively. As usual, we demand that for each framed instanton in this sequence, the parallel transport of each framing to $+\infty$ coincides with the parallel transport of the next framing to $-\infty$.

Now if W is equipped with a family $g(t)$ of metrics (with $g(\infty)$ possibly broken) and $\pi(t)$ of perturbations (with $\pi(\infty)$ possibly broken), one may define a moduli space $M_z(W, c, I; \alpha, \beta)$, whose elements are pairs (t, \mathbf{A}) , where $t \in [0, \infty]$ and \mathbf{A} is a (gauge equivalence class of) framed $\pi(t)$ -instanton in the path z of connections along (W, c) . Precisely, the moduli space for the family is

$$M_z^+(W \cup W', c \cup c', I; \alpha, \beta) = \bigcup_{L \in I} M_z^+(W \cup_L W', c \cup_L c', \pi \cup_L \pi'; \alpha, \beta),$$

topologized as usual using chain-convergence.

Remark 3.34. The subspace

$$M_z^+(W \sqcup W', c \sqcup c', \pi(\infty); \alpha, \beta) \subset M_z^+(W \cup W', c \cup c', I; \alpha, \beta)$$

is not a stratified-smooth space, and is not the same as the fiber product $M_{z_+}^+(W_+; \alpha, \gamma) \times_{\gamma} M_{z_-}^+(W_-; \gamma, \beta)$. This is analogous to the fact that the subset of Δ^2 given by the union of positive-codimension strata is not a stratified-smooth space, and this subset is not the same as $\partial\Delta^2$; the natural map $\partial\Delta^2 \rightarrow \Delta^2$ is not one-to-one, but rather two-to-one on the corners of Δ^2 . \diamond

Then the main result of [Mil19, Section 4.8] is that the same structure results carry through for this parameterized moduli space.

Proposition 3.35. *If I is a regular family of (possibly broken) metrics and perturbations on (W, c) , and $z : \alpha \rightarrow \alpha'$ is a homotopy class with $i(z) \leq 6$, then $M_z^+(W, c, I; \alpha, \beta)$ carries the natural structure of a compact stratified-smooth manifold of dimension $\tilde{i}(z) + \dim \alpha + 1$.*

The boundary strata of this stratified-smooth manifold are as expected: two of them correspond to restricting the family to its endpoints, while the others correspond to factorizations on the two ends of (W, c) . In Section 4.3 below, these will be used to construct homotopies between the induced chain maps of cobordisms.

3.7 Orientations

We conclude our discussion by introducing orientations on the moduli spaces introduced thusfar. Any discussion of orientations is somewhat tedious, so we content ourselves to quoting the results of [Mil19,

Section 5] in a way which is hopefully accessible and easy to use, while making our orientation conventions clear.

First, there is the question of what data is needed to define an orientation on each of the relevant moduli spaces. This uses the notion of *homology orientation*: if (W, c) is a cobordism $(Y, w) \rightarrow (Y', w')$, a homology orientation is an orientation of the real vector space $H^1(W) \oplus H^{2,+}(W) \oplus H^1(Y)$.

Associated to each critical orbit α on (Y, w) and α' on (Y', w') are two-element sets $\Lambda(\alpha)$ and $\Lambda(\alpha')$. In Morse theory, these would correspond to an orientation of the unstable manifold at α ; the infinite-dimensional nature of our problem means this does not make sense. These are defined using orientations of determinant line bundles over the space of connections on a 4-manifold bounding (Y, w, α) .

Given a choice of elements of both $\Lambda(\alpha), \Lambda(\beta)$ on (Y, w) , this induces what is called a *fiber-orientation* on the moduli space $M_\zeta(Y; \alpha, \beta)$: an orientation of the fibers of the map $M_\zeta(Y; \alpha, \beta) \rightarrow \alpha$, which negates if we change either the element of $\Lambda(\alpha)$ or the element of $\Lambda(\beta)$. This choice is such that whenever we are given a choice of elements of $\Lambda(\alpha), \Lambda(\beta)$, and $\Lambda(\gamma)$ the gluing map

$$M_\zeta(Y; \alpha, \beta) \times_\beta M_{\zeta'}(Y; \beta, \gamma) \supset U \rightarrow M_{\zeta*\zeta'}(Y; \alpha, \gamma)$$

is a fiber-orientation preserving diffeomorphism from an open subset of the domain. Notice that here we do not reduce by the \mathbb{R} -action. Our convention is that \check{M}_ζ is given a fiber-orientation so that

$$\mathbb{R} \times \check{M}_\zeta^{\text{Fiber}} \cong M_\zeta^{\text{Fiber}}$$

is a fiber-orientation-preserving diffeomorphism.

If $(W, c) : (Y, w) \rightarrow (Y', w')$ is a cobordism equipped with a homology orientation and a chosen element of $\Lambda(\alpha)$ and $\Lambda(\alpha')$, then the moduli spaces $M_z(W; \alpha, \alpha')$ are all equipped with fiber-orientations; these fiber-orientations are negated if you swap any one of the given choices of element. These fiber-orientations have the property that given further choices of element $\Lambda(\beta)$ and $\Lambda(\beta')$, the maps

$$\begin{aligned} M_\zeta(Y, w; \beta, \alpha) \times_\alpha M_z(W, c; \alpha, \alpha') \supset U &\rightarrow M_{\zeta*z}(W, c; \beta, \alpha'), \\ M_z(W, c; \alpha, \alpha') \times_{\alpha'} M_{\zeta'}(Y', w'; \alpha', \beta') \supset U &\rightarrow M_{z*\zeta'}(W, c; \alpha, \beta') \end{aligned}$$

are fiber-orientation preserving diffeomorphisms.

In the next section, we will find it somewhat more convenient to recast the notion of fiber-orientation; a fiber-orientation is the same as an orientation of $M_\zeta(Y; \alpha, \alpha')$ which depends on a choice of orientation of α , and which negates if we negate the orientation on α . In this language, our convention is that all moduli spaces are oriented as $\alpha \times M_\zeta^{\text{Fiber}}$. In particular, our convention about reduction by the \mathbb{R} -action has us orient the reduced moduli spaces \check{M}_ζ so that there is an orientation-preserving open embedding from an open subset

$$\alpha \times \mathbb{R} \times \check{M}_\zeta^{\text{Fiber}} \supset U \rightarrow \alpha \times M_\zeta^{\text{Fiber}}.$$

When determining the orientation of boundary components of $\check{M}_\zeta^+(Y; \alpha, \beta)$ below, this convention introduces an additional sign of $(-1)^{\dim \alpha}$, corresponding to commuting the factor of \mathbb{R} past the factor of α .

Equipped with orientations, we may explicitly give boundary relations for all of the moduli spaces relevant to us, including their natural signs. Here ‘stratified diffeomorphism’ means an order-preserving bijection between the posets of strata and a homeomorphism between the spaces which covers said bijection and which is a diffeomorphism on each stratum; ‘orientation-preserving’ means that it is an oriented diffeomorphism on the top strata. When M is a stratified-smooth manifold, the closure $\overline{M_s}$ of each face M_s is again a stratified-smooth manifold; the expression ∂M means the disjoint union over $\overline{M_s}$, where M_s is a codimension-1 stratum equipped with the boundary orientation.

The following proposition determines the induced orientations on the boundary components of the moduli spaces that we introduced earlier. It should be understood that all moduli spaces are taken so that $i(z)$ is in the range where these are stratified-smooth manifolds (less than or equal to 8, 7, 6 respectively in the three cases), and the sums are taken over those moduli spaces in homotopy classes which concatenate to z .

Proposition 3.36. *Suppose (Y, w) and (Y', w') are weakly admissible pairs with regular perturbations π and π' . For each $\alpha \in \mathfrak{C}_\pi(Y, w)$, suppose an element of $\Lambda(\alpha)$ and an orientation for the orbit α is given. For each $\alpha' \in \mathfrak{C}_{\pi'}(Y', w')$, pick an element of $\Lambda(\alpha')$. Suppose $(W, c) : (Y, w) \rightarrow (Y', w')$ is an unobstructed cobordism.*

- (i) *Given $\alpha, \beta, \gamma \in \mathfrak{C}_\pi(Y, w)$ and paths ζ_0 from α to β and ζ_1 from β to δ with $\zeta = \zeta_0 * \zeta_1$, the boundary orientation and the fiber product orientation on the boundary component*

$$\check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_\beta \check{M}_{\zeta_1}^+(Y, w; \beta, \delta)$$

of $\check{M}_\zeta^+(Y, w; \alpha, \delta)$ are related by $(-1)^{\tilde{i}(\zeta_0) + \dim(\alpha)}$.

- (ii) *Suppose a regular perturbation π_W for the ASD equation on (W, c) is fixed. Given $\alpha, \beta \in \mathfrak{C}_\pi(Y, w)$, $\alpha' \in \mathfrak{C}_{\pi'}(Y', w')$ and paths ζ_0 from α to β and z_1 from β to α' (along (W, c)) with $z = \zeta_0 * z_1$, the boundary orientation and the fiber product orientation on the boundary component*

$$\check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_\beta M_{z_1}^+(W, c; \beta, \alpha')$$

*of $M_z^+(W, c; \alpha, \alpha')$ are related by $(-1)^{\dim(\alpha)}$. Similarly, given $\alpha \in \mathfrak{C}_\pi(Y, w)$, $\alpha', \beta' \in \mathfrak{C}_{\pi'}(Y', w')$ and paths z_0 from α to β' (along (W, c)) and ζ_1' from β' to α' with $z = z_0 * \zeta_1'$, then the boundary orientation and the fiber product orientation on the boundary component*

$$M_{z_0}^+(W, c; \alpha, \beta') \times_{\beta'} \check{M}_{\zeta_1'}^+(Y', w'; \beta', \alpha') \tag{3.37}$$

of $M_z^+(W, c; \alpha, \alpha')$ are related by $(-1)^{\tilde{i}(z_0) + \dim(\alpha) + 1}$. Note that an orientation of β' is necessary to define the orientation of $M_{\zeta_1'}^+(Y', w'; \beta', \alpha')$ and to define the fiber product orientation. However, the resulting orientation on (3.37) is independent of the orientation of β' .

- (iii) *Suppose a regular 1-parameter family of perturbations $\{\pi_W(t)\}_{t \in I}$ with $I = [0, 1]$ are chosen and $M_z^+(W, c, I; \alpha, \alpha')$ is the corresponding family moduli space. Then $M_z^+(W, c, \pi(0); \alpha, \alpha')$ and $M_z^+(W, c, \pi(1); \alpha, \alpha')$ give two boundary components of the family moduli space whose orientations respectively differs from the boundary orientation by a factor of $(-1)^{\dim \alpha - 1}$ and*

$(-1)^{\dim \alpha}$. Given $\alpha, \beta \in \mathfrak{C}_\pi(Y, w)$, $\alpha' \in \mathfrak{C}_\pi(Y', w')$ and paths ζ_0 from α to β and z_1 from β to α' with $z = \zeta_0 * z_1$, the boundary orientation and the fiber product orientation on the boundary component

$$\check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_\beta M_{z_1}^+(W, c, I; \beta, \alpha')$$

of $M_z^+(W, c, I; \alpha, \alpha')$ are related by $(-1)^{\tilde{i}(\zeta_0) + \dim(\alpha)}$. Similarly, given $\alpha \in \mathfrak{C}_\pi(Y, w)$, $\alpha', \beta' \in \mathfrak{C}_\pi(Y', w')$ and paths z_0 from α to β' and ζ_1' from β' to α' with $z = z_0 * \zeta_1'$, then the boundary orientation and the fiber product orientation on the boundary component

$$M_{z_0}^+(W, c, I; \alpha, \beta') \times_{\beta'} \check{M}_{\zeta_1'}^+(Y', w'; \beta', \alpha') \tag{3.38}$$

of $M_z^+(W, c; \alpha, \alpha')$ are related by $(-1)^{\tilde{i}(z_0) + \dim(\alpha)}$.

Remark 3.39. If one picks a complex line bundle η and a base connection A on it, then replacing the bundle E_c with $E_c \otimes \eta$, the map $A^{\det} \mapsto A^{\det} + A$ identifies the corresponding moduli spaces and all of their extra structure, *except for orientations*. If η is trivial on the ends of W , the orientations of our moduli spaces change by a uniform sign of $(-1)^{c_1^2(\eta)}$; this may be seen by reducing to the Kähler case using excision. See [Kro05, Section 4.2] for more details. It follows that up to sign, the induced map only depends on the associated $SO(3)$ -bundle E_c and not the geometric representative c . \diamond

4 Geometric chains and instanton homology

4.1 Stratified-smooth spaces and the geometric chain complex

The language we use to describe instanton homology below is the language of *geometric chains*. A theory of geometric chains is a flexible machine which, for each smooth manifold M , outputs a chain complex $C_*^{gm}(M)$ whose homology groups are naturally isomorphic to singular homology. This chain complex is freely generated by *probing maps*: smooth maps $\phi : P \rightarrow M$ so that P is a space of a certain pre-determined type, modulo certain pre-determined relations. Given a smooth map $f : M \rightarrow N$, we have an induced map $C_*^{gm}(M) \rightarrow C_*^{gm}(N)$, given at the level of generators by composition; at the level of homology, this agrees with the induced map in singular homology. Finally, given a correspondence $M \leftarrow W \rightarrow N$ between closed smooth manifolds with $r : W \rightarrow M$ a submersion, we can define chain-level *fiber-product maps*

$$\times_M W : C_*^{gm}(M) \rightarrow C_{*+\dim W-\dim M}^{gm}(N).$$

The differential of our instanton chain complex will include both the boundary operator on geometric chains as well as fiber-products with certain moduli spaces. Thus having some form of chain-level fiber product maps is absolutely crucial.

The flexibility to choose the class of probing spaces P and the relations satisfied by probing maps mean that it is easy to define a great many seemingly different theories of geometric chains. Already in [Fuk96, Lip14, Lin18, Mil19], there are four different theories of geometric chains, which are all canonically isomorphic to singular homology. We will need another, for a technical reason outlined below.

In Appendix A, we construct the particular class of probing spaces and equivalence relations used in this paper, called *compact oriented stratified-smooth spaces modulo collapse-equivalence*. In this section, we outline the salient features of our relevant geometric homology theory. The stated facts below are proven in the appendices. First, we'll define stratified-smooth spaces, our generalization of the notion of stratified-smooth manifolds from Definition 3.17.

Definition 4.1. A naive stratified-smooth space of dimension n consists of the following data:

- (i) A partially ordered set Δ with an order-preserving grading function $d : \Delta \rightarrow \mathbb{Z}_{\geq 0}$ with the maximum value n ,
- (ii) A topological space P stratified by subsets P_s for $s \in \Delta$, with $\overline{P_s} = \bigcup_{t < s} P_t$,
- (iii) A smooth structure on each stratum P_s .

An *oriented* naive stratified-smooth space is the additional data of an orientation on all top-dimensional strata. ◇

When P is compact, as all examples of interest to us will be, the set Δ is necessarily finite. Given a continuous map $\psi : [0, \varepsilon)^{n-d(s)} \times \mathbb{R}^{d(s)+\ell} \rightarrow \mathbb{R}^\ell$ which, on each stratum, is smooth and has zero as a regular value, the zero set $Z(\psi) = \psi^{-1}(0)$ is naturally a naive stratified-smooth space.

Definition 4.2. A stratified smooth space is a naive stratified-smooth space so that each point $p \in P$ has a neighborhood which is *locally isomorphic* to a naive stratified-smooth space of the form $Z(\psi)$.

We say this stratified-smooth space is oriented if it is also equipped with an orientation on the top strata so that these local isomorphisms can be chosen to be orientation-preserving. \diamond

The notion of ‘locally isomorphic’ is intricate, and includes the data of the stratification. The interested reader can find a significantly more detailed discussion, as well as some examples and non-examples, in Appendix A.1. Orientations are discussed in Appendix A.2. Because these are mostly a *technical necessity*, we do not go into further detail here. The moduli spaces of instantons $\check{M}_\zeta^+(\alpha, \beta)$ are always stratified-smooth *manifolds*, and the first time stratified-smooth spaces make an appearance geometrically is in Section 5.2. For now, they are only used to define a geometric chain complex whose generators include stratified-smooth manifolds, and there is little harm in imagining all ‘stratified-smooth spaces’ below are smooth manifolds with corners.

We shall define a *geometric chain complex* using stratified-smooth maps $\phi : P \rightarrow M$ from compact, connected, oriented stratified-smooth spaces as our probes into M , much the same as simplicial homology is constructed by probing M with continuous maps from simplices. We call such a pair (P, ϕ) a probe (in M). We will use the term *stratified-smooth chain (in M)* for a formal linear combination of probes.

When P is a stratified-smooth space, we will verify in Lemma A.13 that the closure \overline{P}_s of each stratum is also stratified-smooth. Then the boundary of a stratified-smooth probe is given as

$$\partial P = \sum_{d(s)=\dim P-1} \overline{P}_s,$$

where each boundary stratum is given an appropriate boundary orientation. This is a finite sum by compactness of P .

While the detailed foundations of stratified-smooth spaces and the geometric chain complex are carried out in the appendix, the following key points are all that is needed for the main text:

- If two probes $\phi : P \rightarrow M$ and $\eta : Q \rightarrow M$ are stratum-preserving diffeomorphic by an orientation-preserving diffeomorphism $f : P \rightarrow Q$ with $\eta f = \phi$, then we set $\phi = \eta$ as elements of the geometric chain complex $C_*^{gm}(M)$.
- If a probe $\phi : P \rightarrow M$ supports an orientation-reversing diffeomorphism $f : P \rightarrow P$ with $\phi f = \phi$, then we set $\phi = 0$ as an element of the geometric chain complex $C_*^{gm}(M)$.
- If a probe $\phi : P \rightarrow M$ has $\dim P > 0$ and

$$\text{rank}(d_x \phi_s : T_x P_s \rightarrow T_{\phi(x)} M) < \dim(P_s)$$

for all strata with $\dim P_s > 0$ and all $x \in P_s$, we say ϕ is degenerate, and we set $\phi = 0$ as an element of $C_*^{gm}(M)$.

- The most interesting relation, novel to this article and crucial in our definition of the proof of invariance of instanton homology, is *collapse-equivalence*. Let (P, ϕ) and (Q, η) be probes.

Suppose that there are loci $Z_P \subset P$ and $Z_Q \subset Q$ which are unions of embedded stratified-smooth subspaces of positive codimension, and *degenerate* boundary strata of P and Q , and suppose that there exists an orientation-preserving isomorphism $f : P \setminus Z_P \rightarrow Q \setminus Z_Q$ with $\eta f = \phi$. Then we say (Q, η) is *collapse-equivalent* to (P, ϕ) and set them equal in $C_*^{gm}(M)$.

- We use a truncation procedure to ignore high-dimensional chains, so that $C_*^{gm}(M)$ is concentrated in degrees $[0, \dim M]$; in degrees $d < \dim M$, this is simply $C_*^{gm}(M)$, while in degree $d = \dim M$, this is the quotient by the subgroup $\partial C_{\dim M+1}^{gm}$.
- Given smooth manifolds M and N , a probe $(f, g) : W \rightarrow M \times N$ is called an *oriented correspondence from M to N* . If f is a submersion and M is oriented, then the assignment

$$(P, \phi) \mapsto (P \times_M W, g\pi_2)$$

gives a map $\times_M W : C_*^{gm}(M) \rightarrow C_{*+\dim W-\dim M}^{gm}(N)$ which satisfies the relation

$$\partial(\phi \times_M W) = (\partial\phi) \times_M W + (-1)^{\dim P + \dim M} \phi \times_M (\partial W).$$

If either one of the orientation on M or W is negated, then the induced map $\times_M W$ is negated. Given a pair of correspondences $V \rightarrow L \times M$ and $W \rightarrow M \times N$ (with $V \rightarrow L$ and $W \rightarrow M$ submersions), the composite $(\phi \times_L V) \times_M W$ is equal to $\phi \times_L (V \times_M W)$. That is, the fiber product operation is associative.

- If $\phi : W \rightarrow M \times N$ is zero as an element of $C_*^{gm}(M \times N)$, then $\phi \times_M W = 0$ identically on the chain level. In particular, if $\dim W > \dim M + \dim N$, then $\phi \times_M W = 0$, and if W is collapse-equivalent to W' , then $\phi \times_M W = \phi \times_M W'$. It also follows that if V and W are composable correspondences (given by geometric chains in $L \times M$ and $M \times N$, respectively) so that

$$\dim V + \dim W - \dim M > \dim L + \dim N,$$

then $\phi \times_L (V \times_M W)$ is the zero operator at the chain level.

Example 4.3. The most important example of collapse-equivalence for our applications is obtained by real blowup, discussed formally in Appendix A.3. If $\psi : P \rightarrow E$ is a stratified-smooth section of a vector bundle which is transverse to the zero section, one may take the real blowup $B(\psi)$ of this section, replacing the zero locus with a sphere bundle over said zero locus. There is a natural map $B(\psi) \rightarrow P$ which is a collapse-equivalence. This will play an important role in Section 5 of this article, and we introduce collapse-equivalence so that we can consider a chain equal to its blowup. \diamond

In Appendix A.4, we verify that stratified-smooth spaces satisfy all the necessary requirements to give a suitable theory of geometric chains, canonically isomorphic to singular homology. We use the corresponding geometric chain complex freely below. In trying to understand the arguments of this article, it is almost entirely harmless to imagine that ‘stratified-smooth space’ may be read ‘smooth manifold with corners’. There is a perfectly good notion of collapse-equivalence for smooth manifolds with corners, for instance. In Appendix A.5 we add the truncation mentioned above.

Henceforth when M is a smooth manifold, we suppress *gm* from notation and write $C_*(M)$ for $tC_*^{gm}(M)$, the truncated geometric chains of Appendix A.5.

Remark 4.4. It seems plausible that with significant analytic work, we could prove that the instanton moduli spaces (and the modifications we discuss later) can be given the structure of smooth manifolds with corners. Assuming this, the theory of stratified-smooth chains would be wholly unnecessary, and we would be able to carry out our constructions with a more standard notion of geometric chains (like smooth manifolds with corners). However, all that we've actually proved is that these moduli spaces carry the structure of stratified-smooth manifolds. If we do not want to attempt to construct smooth atlases for the instanton moduli spaces, the development of the stratified-smooth theory is then forced upon us by our need for transverse intersections. This is a tradeoff, but we believe that in this case accepting stratified-smooth chains is the path of least resistance. \diamond

4.2 Flow categories and bimodules

With a suitable notion of geometric chains in hand, we will now define *flow categories and bimodules of geometric chains*. These are almost precisely the data arising from the moduli spaces of flowlines of a Morse–Bott function, though we make the drastic assumption that all endpoint maps are submersions. This is not necessary, but simplifies discussions of orientation and of the domains of various maps; it is also necessarily true in all of our cases of interest (which arise from G -equivariant Morse theory, where the critical submanifolds are orbits and endpoint maps are equivariant).

A flow category is precisely what we need to construct a Morse-Bott type complex in geometric chains.

Definition 4.5. A finitely-graded flow category \mathcal{C} in geometric chains is the following data, subject to the following restrictions:

- A finite set of objects $\text{Ob}(\mathcal{C})$, each equipped with the structure of a closed, connected, orientable smooth manifold. (Orientability is not strictly necessary, and is completely irrelevant over \mathbb{F}_2 , but makes our discussion of orientations on the moduli spaces simpler.)
- For each object $X \in \text{Ob}(\mathcal{C})$ a 2-element set $\Lambda(X)$ called the *orientation set*.
- A grading function $\tilde{i} : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{Z}/2N$ so that $\tilde{i}(X, Y) + \tilde{i}(Y, Z) = \tilde{i}(X, Z)$; when multiple flow categories are under consideration we will write $\tilde{i}_{\mathcal{C}}$ for this function. This is called the ‘fiber-dimension’, or ‘fiber-index’.
- For each pair X, Y of objects, for each orientation of X , and for each element of $\Lambda(X)$ and $\Lambda(Y)$, we have a geometric chain, written as

$$\mathcal{C}(X, Y) = \mathcal{C}_{XY} \in C_{\dim X + \tilde{i}(X, Y) - 1}(X \times Y),$$

so that the map $\mathcal{C}(X, Y) \rightarrow X$ is a submersion. We require that if we reverse any one of the orientation on X , the element of $\Lambda(X)$, or the element of $\Lambda(Y)$, the resulting geometric chain is the previous geometric chain with its orientation reversed.

- Given a sequence X, Y, Z of objects, we demand that we have equality

$$(-1)^{\dim X} \partial \mathcal{C}_{XY} = \sum_{Z \in \text{Ob}(\mathcal{C})} (-1)^{\tilde{i}(X, Z)} \mathcal{C}_{XZ} \times_Z \mathcal{C}_{ZY}$$

as elements of $C_*(X \times Y)$. ◇

Let us try to clarify a few points about this definition.

In Morse-Bott theory, to orient the moduli spaces of flowlines from X to Y , one needs to orient a number of things: X , the unstable bundle of X , and the stable bundle of Y . In the second bullet point, the sets $\Lambda(X)$ should be understood as analogous to the set of orientations on the stable bundle of X . While these are not canonical, once one makes a choice for each X (and orients each X) we obtain induced orientations on all moduli spaces of flowlines, with predictable change as we change any of our choices.

By the orientation assumption, and the way fiber product maps $\times_X \mathcal{C}_{XY}$ behave under orientation-reversal of X , one obtains a well-defined map

$$\times_X \mathcal{C}_{XY} : C_*(X) \otimes_{\mathbb{Z}/2} \mathbb{Z}[\Lambda(X)] \rightarrow C_*(Y)[\tilde{i}(X, Y) - 1] \otimes_{\mathbb{Z}/2} \mathbb{Z}[\Lambda(Y)],$$

where $\mathbb{Z}/2$ acts on the complex $C_*(X)$ and $C_*(Y)$ by negation/orientation-reversal and on $\Lambda(X)$ and $\Lambda(Y)$ by swapping the two elements; this map is independent of the orientation on X and does not require any choice of an element of $\Lambda(X)$ or $\Lambda(Y)$.⁸

In the fourth bullet point, the index should be read as an integer mod $2N$, so that $\mathcal{C}(X, Y)$ is an element of

$$\bigoplus_{k \equiv_{2N} \dim X + \tilde{i}(X, Y) - 1} C_k(X \times Y),$$

noting that this group vanishes when $k > \dim X + \dim Y$ so there are only finitely many nonzero terms in this sum. Here and in what follows \equiv_{2N} is shorthand for ‘equal modulo $2N$ ’.

Remark 4.6. This definition is a weakening of the standard definition [Zho20, Definition 2.8] of flow category in three ways. First, we only pay attention to these moduli spaces \mathcal{C}_{XY} up to a certain dimension, $\dim X + \dim Y$. Second, when looking at the fiber product $\mathcal{C}_{XZ} \times_Z \mathcal{C}_{ZY}$, we ignore all fiber product chains larger than $\dim X + \dim Y$. Third, usually one demands that the term $\bigsqcup_Z \mathcal{C}_{XZ} \times_Z \mathcal{C}_{ZY}$ is oriented diffeomorphic to $\partial \mathcal{C}_{XZ}$. Here, we only ask that they are equal as geometric chains; this means that if one is diffeomorphic to a *collapse* of the other, or differs from another by a degenerate chain, they still give the same element of the geometric chain complex. On the other hand, the submersive condition is much more strict than Zhou’s definition. If desired, it is not hard to remove this assumption in what follows, but all flow categories arising from G -invariant Morse theory (for compact Lie G) satisfy this condition. ◇

In practice, one constructs finitely-graded flow categories by providing actual stratified-smooth spaces (not their representatives in the chain complex, which involves a quotient by additional relations), and only doing so through a given dimension range. This use of a dimension range is precisely why we truncate our complex of geometric chains, so that the following lemma holds.

Lemma 4.7. *Suppose one has a set of objects X , a grading function \tilde{i} , and a collection of compact oriented stratified-smooth spaces \mathcal{C}_{XY} , only defined in the dimension range $\dim \mathcal{C}_{XY} \leq \dim X + \dim Y + 1$, which satisfy Definition 4.5 at the level of (non-truncated) geometric chains. Then these define a finitely-graded flow category in (truncated) geometric chains.*

⁸Equivalently but less canonically, one may suppress the sets $\Lambda(X)$ by choosing an element of each from the getgo.

Proof. Each compact oriented stratified-smooth space \mathcal{C}_{XY} gives rise to a corresponding element of the truncated geometric chain complex. We need to verify that the boundary relation (the final bullet point in Definition 4.5) still holds for all X, Y . It is useful to imagine for a moment that the complex is \mathbb{Z} -graded.

For $\dim \mathcal{C}_{XY} \leq \dim X + \dim Y$, the truncation procedure leaves both sides of the final bullet point unchanged; the result holds at the level of truncated geometric chains because it holds at the level of geometric chains.

For $\dim \mathcal{C}_{XY} = \dim X + \dim Y + 1$, the left-hand side is tautologically zero (this is the boundary of a chain which is zero at the chain level). As for the right-hand side, recall that in the definition of the (truncated) geometric chain complex, the relevant degree is defined to be the quotient

$$tC_{\dim X + \dim Y}^{gm}(X \times Y) = C_{\dim X + Y}^{gm}(X \times Y) / \partial C_{\dim X + \dim Y + 1}^{gm}(X \times Y).$$

Because the right-hand side is known to be the boundary of a $\dim X + \dim Y + 1$ -dimensional chain, it is zero on the chain level as well. Thus the relation reads $0 = 0$.

For $\dim \mathcal{C}_{XY} > \dim X + \dim Y + 1$, both sides of the equation are tautologically zero at the chain level, as $tC_d(X \times Y) = 0$ for $d > \dim X + \dim Y$.

In the finitely-graded case (where $\dim \mathcal{C}_{XY}$ need not be constant) the discussion above shows that the terms of dimension larger than $\dim X + \dim Y + 1$ are irrelevant, whereas the terms in dimension $\dim X + \dim Y + 1$ are still needed for verifying the boundary relation, even though these terms are zero at the chain level. \square

In the instanton setting, $\dim X + \dim Y$ is at most 6 and $2N = 8$, so this means we need those moduli spaces of dimension between 0 and 7, and the relevant moduli spaces have constant dimension.

The point of Definition 4.5 is that, using the fiber product constructions above with an appropriate theory of geometric chains, a flow category gives rise to a chain complex, the ‘flow complex’.

Definition 4.8. Suppose one has a finitely-graded flow category \mathcal{C} in geometric chains. There is an associated *flow complex* $CM_*(\mathcal{C})$ which comes equipped with a relative $\mathbb{Z}/2N$ -grading

$$CM_*(\mathcal{C}) = \bigoplus_{X \in \text{Ob}(\mathcal{C})} C_*(X)[\tilde{i}(X)] \otimes_{\mathbb{Z}/2} \mathbb{Z}[\Lambda(X)],$$

where the shift denotes that if $(P, \phi) \in C_*(X)$ and $(Q, \eta) \in C_*(Y)$ then their relative grading is

$$\text{gr}[(P, \phi), (Q, \eta)] = \dim P - \dim Q + \tilde{i}(X, Y).$$

One imagines that $\tilde{i}(X, Y) = \tilde{i}(X) - \tilde{i}(Y)$ for some illusory function $\tilde{i}(X)$.

The boundary operator is $\partial + d_M$, where ∂ is the boundary operator on geometric chains and

$$d_M \phi = \sum_Y (-1)^{\dim \phi} \phi \times_X \mathcal{C}_{XY}. \quad \diamond$$

When \mathcal{C} arises as the flow category of a Morse-Bott-Smale function on a Riemannian closed smooth manifold, this complex should be called the Morse-Bott complex, and its homology recovers the homology

of said manifold. The idea for this complex originates from a differential forms version used to develop equivariant cohomology in [AB95]. Not long after, Fukaya [Fuk96] used a variant defined over the integers to discuss Morse-Bott homology and develop a connected-sum theorem for instanton homology; our version here resembles Fukaya's, who also discussed geometric chains. The notation $CM_*(\mathcal{C})$ is chosen to reflect the Morse-theoretic origin of this complex.

Proposition 4.9. *The above construction defines a chain complex.*

Proof. We will suppress the sets $\Lambda(X)$ from the proof below by, for instance, picking an element of each from the start.

Observe that the differential has relative grading -1 : the geometric boundary operator ∂ has degree -1 , and because $\times_X \mathcal{C}_{XY} : C_*(X) \rightarrow C_*(Y)$ has degree $\tilde{i}(X, Y) - 1$, it has degree -1 when considered as a map

$$C_*(X)[\tilde{i}(X)] \rightarrow C_*(Y)[\tilde{i}(Y)],$$

so that d_M has degree -1 .

Next we should verify that the differential squares to zero. We have

$$\begin{aligned} d^2\phi &= \partial^2\phi + \sum_Y (-1)^{\dim \phi} \partial(\phi \times_X \mathcal{C}_{XY}) + (-1)^{\dim \phi - 1} \partial\phi \times_X \mathcal{C}_{XY} \\ &\quad + \sum_{Y,Z} (-1)^{\dim \phi + \tilde{i}(X,Z) - 1} (-1)^{\dim \phi} \phi \times_X \mathcal{C}_{XZ} \times_Z \mathcal{C}_{ZY}. \end{aligned}$$

The first term is zero, and the sign in the last term can be simplified to $\sum (-1)^{\tilde{i}(X,Z) - 1}$. Using the formula for the boundary of a fiber product boundary relation in the definition of flow category, the second term may be expanded out as

$$\begin{aligned} &(-1)^{\dim \phi} \left(\partial\phi \times_X \mathcal{C}_{XY} + (-1)^{\dim \phi + \dim X} \phi \times_X \partial\mathcal{C}_{XY} \right) \\ &= (-1)^{\dim \phi} \partial\phi \times_X \mathcal{C}_{XY} + \sum_Z (-1)^{\tilde{i}(X,Z)} \phi \times_X \mathcal{C}_{XZ} \times_Z \mathcal{C}_{ZY}, \end{aligned}$$

using our formulas for the boundary of a fiber product and for the boundary of \mathcal{C}_{XY} . Thus the second term of the sum is precisely the negative of the third plus fourth terms, and therefore $d^2\phi = 0$. \square

Remark 4.10. When we work with \mathbb{Z} -graded flow categories, the filtration by index \tilde{i} gives rise to a spectral sequence computing $HM_*(\mathcal{C})$. In the finitely-graded case, the discussion is slightly more subtle.

Again we suppress $\Lambda(X)$. Choose (arbitrarily) a lift of the relative grading to an absolute grading $\tilde{i}(X)$. This gives us a \mathbb{Z} -graded lift

$$\widetilde{CM}_*(\mathcal{C}) = \bigoplus_{\substack{X \in \text{Ob}(\mathcal{C}) \\ n \equiv_{2N} \tilde{i}(X)}} C_*(X)[n],$$

with the appropriate differential.

We may equip $\widetilde{CM}_*(\mathcal{C})$ with a *periodic filtration*: a filtration so that

$$F_{k+2N}\widetilde{CM}_{*+2N} = F_k\widetilde{CM}_*$$

This filtration is obtained by taking the sum only over those components with $n \leq k$. Because $\mathcal{C}_{XY} \rightarrow X$ is a submersion, $\dim \mathcal{C}_{XY} \geq \dim X$ for all X, Y , so that every nonzero term in the above sum has $\tilde{i}_{\mathcal{C}}(X, Y) \geq 1$. It follows that the differential preserves the filtration, and the associated graded differential is the geometric boundary operator (with no fiber product terms).⁹

As discussed in [Mil19, Section A.8], this gives rise to a $\mathbb{Z}/2N \times \mathbb{Z}$ -graded (relatively graded in the first coordinate, absolutely graded in the second) spectral sequence $(E^r, d^r) \Rightarrow HM^*(\mathcal{C})$ with

$$E^1 \cong \bigoplus_{X \in \text{Ob}(\mathcal{C})} H_*(X).$$

The differential on the E^1 page is given by fiber product with those moduli space components \mathcal{C}_{XY} of dimension $\dim X$ (that is, where $\tilde{i}_{\mathcal{C}}(X, Y) = 1$, so that $\text{Fiber}(\mathcal{C}_{XY} \rightarrow X)$ is a finite set). This filtration is degreewise finite, so the spectral sequence is strongly convergent, meaning that its E^∞ page is the associated graded group of the periodic filtration on $HM_*(\mathcal{C})$. \diamond

Just as flow categories give rise to chain complexes, it is natural to imagine that a similar structure gives rise to chain maps between flow complexes. This leads us to the following structure, which form the morphisms in a category sFlow (the ‘s’ for ‘submersive’).

Definition 4.11. Suppose \mathcal{C} and \mathcal{C}' are both finitely graded flow categories in geometric chains. A bimodule \mathcal{M} between these, written $\mathcal{M} : \mathcal{C} \rightarrow \mathcal{C}'$, is the following data, subject to the following restrictions.

- We are given a function $\tilde{i}_{\mathcal{M}} : \text{Ob}(\mathcal{C}) \times \text{Ob}(\mathcal{C}') \rightarrow \mathbb{Z}/2N$, which is affine in the sense that given $X, Z \in \mathcal{C}$ and $Y, W \in \mathcal{C}'$ we have

$$\tilde{i}_{\mathcal{C}}(X, Z) + \tilde{i}_{\mathcal{M}}(Z, W) + \tilde{i}_{\mathcal{C}'}(W, Y) = \tilde{i}_{\mathcal{M}}(X, Y).$$

- We also have, for each $(X, Y) \in \text{Ob}(\mathcal{C}) \times \text{Ob}(\mathcal{C}')$ and each orientation on X and choice of element of $\Lambda(X)$ and $\Lambda(Y)$, a geometric chain

$$\mathcal{M}(X, Y) = \mathcal{M}_{XY} \in C_{\dim X + \tilde{i}_{\mathcal{M}}(X, Y)}(X \times Y).$$

We demand that $\mathcal{M}_{XY} \rightarrow X$ is a submersion for all X, Y . Once again we demand that the orientation of \mathcal{M}_{XY} should reverse if we reverse any one of the orientation of X , the element of $\Lambda(X)$, or the element of $\Lambda(Y)$.

- We have an equality at the level of geometric chains in $X \times Y$

$$(-1)^{\dim X} \partial \mathcal{M}_{XY} = \sum_{Z \in \text{Ob}(\mathcal{C})} \mathcal{C}_{XZ} \times_Z \mathcal{M}_{ZY} + (-1)^{\tilde{i}_{\mathcal{M}}(X, W) + 1} \sum_{W \in \text{Ob}(\mathcal{C}')} \mathcal{M}_{XW} \times_W \mathcal{C}'_{WY}. \quad \diamond$$

⁹If we drop the submersive assumption, this spectral sequence fails to exist unless $\dim \mathcal{C}_{XY} \geq \dim X - 1$ for all X, Y , and the E^1 page has a twisted differential unless $\dim \mathcal{C}_{XY} \geq \dim X$ for all X, Y .

The fiber-dimension function $\tilde{i}_{\mathcal{M}}$ should be understood here as the analogue of degree of a chain map, in the absence of an absolute grading.

Remark 4.12. Given a generic correspondence $M \xleftarrow{f} W \xrightarrow{g} N$ between closed oriented smooth manifolds with Morse-Bott functions, there are moduli spaces of points in W which, projected to M (resp N), flow upward (downward) to a given critical submanifold. The collection of such moduli spaces $\mathcal{M}_W(X, Y)$ form a bimodule in the above sense. In a moment, we will see that a bimodule induces a chain map between the flow complexes. In the Morse-Bott case described here, the induced map on homology is $gf^!$, where $f^!$ is the umkehr map defined using Poincare duality on M and W . See, for instance, [KM07, Section 2.8] for the Morse case. \diamond

As expected, a bimodule induces a chain map between the flow complexes.

Proposition 4.13. *Given a bimodule $\mathcal{M} : \mathcal{C} \rightarrow \mathcal{C}'$, there is an induced chain map*

$$CM_*(\mathcal{M}) : CM_*(\mathcal{C}) \rightarrow CM_*(\mathcal{C}').$$

Given $\phi \in C_*(X)$ for $X \in \text{Ob}(\mathcal{C})$, we write

$$CM_*(\mathcal{M})\phi = \sum_{Y \in \text{Ob}(\mathcal{C}')} \phi \times_X \mathcal{M}_{XY}.$$

The proof that this is a chain map is a straightforward definition-push akin to the proof of Proposition 4.9, and we leave it to the reader.

Remark 4.14. In terms of the periodic filtration before, the chain map induced by a bimodule is a periodically filtered map; its behavior on relative gradings is described by $i_{\mathcal{M}}$. The induced map on the E^1 page of the spectral sequence is given by the induced map of $\times_X \mathcal{M}(X, Y)_{\dim X}$ on homology; here the subscript indicates that we only refer to those $\mathcal{M}(X, Y)$ of dimension $\dim X$, so that $\mathcal{M}(X, Y)$ has fiber-dimension zero. \diamond

We have the morphisms in our category of flow categories; for this to actually form a category, we should define identity and composite morphisms, and check the associativity and identity axioms.

Definition 4.15. Given a flow category \mathcal{C} , we say its *identity bimodule* $\mathbf{1}_{\mathcal{C}}$ has $\tilde{i}_{\mathbf{1}}(X, Y) = \tilde{i}(X, Y)$, and

$$\mathbf{1}_{\mathcal{C}}(X, Y) = \begin{cases} 0 & X \neq Y \\ X & X = Y \end{cases}.$$

Given two bimodules $\mathcal{M} : \mathcal{C} \rightarrow \mathcal{C}'$ and $\mathcal{N} : \mathcal{C}' \rightarrow \mathcal{C}''$, their composition $\mathcal{N} \circ \mathcal{M} : \mathcal{C} \rightarrow \mathcal{C}''$ is the bimodule with

$$\tilde{i}_{\mathcal{N} \circ \mathcal{M}}(X, Z) = \tilde{i}_{\mathcal{M}}(X, Y) + \tilde{i}_{\mathcal{N}}(Y, Z)$$

for any $Y \in \text{Ob}(\mathcal{C}'')$ and

$$(\mathcal{N} \circ \mathcal{M})(X, Z) = \sum_{Y \in \text{Ob}(\mathcal{C}'')} \mathcal{M}(X, Y) \times_Y \mathcal{N}(Y, Z). \quad \diamond$$

Remark 4.16. One gets a slightly different prediction for the identity bimodule from the perspective of Morse theory. There, the identity bimodule should arise as the moduli spaces corresponding to the identity correspondence, or equivalently, the compactified trajectory space of flowlines on M (the space of parameterized flowlines, including constant flowlines and broken flowlines). This would assign $\mathbf{1}_{\text{Morse}}(X, X) = X$, while $\mathbf{1}_{\text{Morse}}(X, Y) = [0, 1] \times \mathcal{C}(X, Y)$ when $X \neq Y$. However, note that $\mathbf{1}_{\text{Morse}}(X, Y)$ is degenerate as a chain in $X \times Y$: it factors through the projection $[0, 1] \times \mathcal{C}(X, Y) \rightarrow \mathcal{C}(X, Y)$, and this chain supports an orientation-reversing isomorphism $(t, \gamma) \mapsto (1 - t, \gamma)$, and hence is identically zero on the chain level for $X \neq Y$. So the identity bimodule above agrees with the one predicted by Morse theory, even at the chain level.

The phrasing above avoids discussion of this degenerate chain. \diamond

It is straightforward to verify the bimodule relations for the bimodules introduced above; it is easy to see that composition is associative and that the identity bimodule acts as the identity, thus proving the following proposition.

Proposition 4.17. *Submersive flow categories, with bimodules-up-to-isomorphism as morphisms, form a category sFlow with identity and composition as described above. The identity and composite bimodules act as the identity and composite morphisms on the flow complex.*

The sorts of bimodules (and thus chain maps) that arise from Morse–Bott theory depend on auxiliary data, such as metric structures. However, the induced maps on homology are well-defined. The usual argument chooses a 1-parameter family of auxiliary data, and uses this to construct a chain homotopy between the two chain maps. The definition below attempts to abstract this situation to the setting of flow categories.

Definition 4.18. Suppose we are given two bimodules $\mathcal{M}_0, \mathcal{M}_1 : \mathcal{C} \rightarrow \mathcal{C}'$ so that $\tilde{i}_{\mathcal{M}_0} = \tilde{i}_{\mathcal{M}_1}$ (that is, these two bimodules have the same ‘relative degree’). A homotopy $\mathcal{H} : \mathcal{M}_0 \Rightarrow \mathcal{M}_1$ is the data of, for each $(X, Y) \in \text{Ob}(\mathcal{C}) \times \text{Ob}(\mathcal{C}')$, and each triple of an orientation on X and element of $\Lambda(X)$ and $\Lambda(Y)$, a geometric chain

$$\mathcal{H}(X, Y) = \mathcal{H}_{XY} \in C_{\dim X + \tilde{i}_{\mathcal{M}}(X, Y) + 1}(X \times Y)$$

so that $\mathcal{H}_{XY} \rightarrow X$ is a submersion for all X, Y . As usual we demand that the orientation of \mathcal{H}_{XY} should reverse if we reverse any one of the orientation of X , the element of $\Lambda(X)$, or the element of $\Lambda(Y)$.

Finally, we demand an equality at the level of geometric chains in $C_*(X \times Y)$

$$\begin{aligned} (-1)^{\dim X} \partial \mathcal{H}(X, Y) &= \mathcal{M}_1(X, Y) - \mathcal{M}_0(X, Y) \\ &+ \sum_{Z \in \text{Ob}(\mathcal{C})} (-1)^{\tilde{i}_{\mathcal{C}}(X, Z)} \mathcal{C}_{XZ} \times_Z \mathcal{H}_{ZY} \\ &+ \sum_{W \in \text{Ob}(\mathcal{C}')} (-1)^{\tilde{i}_{\mathcal{M}}(X, W)} \mathcal{H}_{XW} \times_W \mathcal{C}'_{WY}. \end{aligned} \quad \diamond$$

Indeed, a straightforward calculation shows that this is precisely what is needed to define a chain homotopy between induced maps of bimodules.

Proposition 4.19. *A homotopy $\mathcal{H} : \mathcal{M}_0 \Rightarrow \mathcal{M}_1$ of bimodules induces a chain homotopy*

$$CM_*(\mathcal{H}) : CM_*(\mathcal{M}_0) \Rightarrow CM_*(\mathcal{M}_1).$$

It follows that the induced map of a bimodule on homology is well-defined up to bimodule homotopy.

The formula for the chain homotopy follows the same line as for the induced map, together with an extra sign factor: for $\phi \in C_*(X)$ we write

$$CM_*(\mathcal{H})\phi = \sum_{Y \in \text{Ob}(\mathcal{C}')} (-1)^{\dim \phi} \phi \times_X \mathcal{H}_{XY}.$$

If one likes, it follows from formal properties of this notion of homotopy that there is a well-defined homotopy category hsFlow and that taking the flow complex gives a functor

$$CM_* : \text{hsFlow} \rightarrow \text{hKom}_{2N}(\mathbb{Z})$$

to the homotopy category of $\mathbb{Z}/2N$ -graded complexes.

Suppose all of the flow categories and bimodules and so on above carry the action of a compact Lie group G , acting smoothly on all objects X and all moduli spaces \mathcal{C}_{XY} so that the correspondence map $\mathcal{C}_{XY} \rightarrow X \times Y$ is G -equivariant. Then $C_*(G)$ acts on all of the $C_*(X)$'s, compatibly with the fiber product maps, so that $CM_*(\mathcal{C})$ naturally carries the structure of a dg $C_*(G)$ -module. Further, a bimodule \mathcal{W} with a G -action compatible with those on the flow categories induces a $C_*(G)$ -equivariant chain map, and a G -equivariant homotopy between bimodules induces a $C_*(G)$ -equivariant chain homotopy between these maps.

The comment above goes through without change: there is a functor $\text{hsFlow}^G \rightarrow \text{hKom}_{2N}^G(\mathbb{Z})$ from the homotopy category of G -equivariant submersive flow categories and to the homotopy category of $\mathbb{Z}/2N$ -graded dg $C_*(G)$ -modules.

4.3 The instanton chain complex and unobstructed cobordism maps

Let (Y, w, π, y) be a weakly admissible pair with regular perturbation π and basepoint $y \in Y$. Then Proposition 3.36(i) shows that there is an associated flow category $\mathcal{I}(Y, w, \pi)$, with

- object set the set of $SO(3)$ -orbits of π -flat connections;
- orientation sets the $\Lambda(\alpha)$ discussed in Section 3.7;
- relative grading $\tilde{i}(Y; \alpha, \beta) = \tilde{i}(\zeta) \pmod{8}$, where $\zeta : \alpha \rightarrow \beta$ is any homotopy class of path of framed connections;
- morphisms the moduli spaces $\check{M}_\zeta^+(Y; \alpha, \beta)$, where $\zeta : \alpha \rightarrow \beta$ is the unique homotopy class with $0 \leq \tilde{i}(\zeta) \leq 7$.

As in Lemma 4.7, when verifying that this defines a flow category, we only need to define those moduli spaces with $\dim \check{M}(Y; \alpha, \beta) \leq \dim \alpha + \dim \beta + 1$. Because $\dim \check{M}_\zeta(Y; \alpha, \beta) = \dim \alpha + \tilde{i}(\zeta) - 1$, so we need those moduli spaces with $\tilde{i}(\zeta) \leq \dim \beta + 2$.

Because the dimension of an $SO(3)$ -orbit is at most 3, we need those moduli spaces with $\tilde{i}(\zeta) \leq 5$. Similarly, for the moduli spaces $M_z(W)$ on a cobordism (W, c) , we need those moduli spaces $M_z(W, c; \alpha, \alpha')$ with $\tilde{i}(z) \leq 4$; and for the moduli spaces $M_z(W, c, I; \alpha, \alpha')$ for a 1-parameter family of perturbations on a cobordism, we need those moduli spaces with $\tilde{i}(z) \leq 3$.

These are promised to us by Propositions 3.16, 3.20, and 3.35 respectively, so long as $i(z) \leq 8, 7, 6$ respectively. We have the relation

$$\tilde{i}(z) = i(z) + (3 - \dim \alpha) \geq i(z),$$

so that in particular we have well-defined stratified-smooth moduli spaces whenever $\tilde{i}(z) \leq 8, 7, 6$, regardless of what α is. In all cases, this means our moduli spaces are defined through fiber dimension $7 + (3 - \dim \alpha) = 10 - \dim \alpha \geq 7$, which is more than enough for our purposes.

Notice that $\mathcal{I}(Y, w, \pi)$ is an $SO(3)$ -equivariant flow category, so its flow complex is a dg-module over $C_*SO(3)$.

Definition 4.20. We define the *tilde complex* to be the associated flow complex

$$\tilde{C}_*(Y, w, \pi) = CM_*(\mathcal{I}(Y, w, \pi)),$$

and its structure as a $C_*SO(3)$ -module. ◇

The rest of Proposition 3.36 shows that unobstructed cobordisms $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ with homology orientation give rise to bimodules between the instanton flow categories $\mathcal{I}(Y, w, \pi)$ and $\mathcal{I}(Y', w', \pi')$, while 1-parameter families of (possibly broken) metrics and perturbations give rise to homotopies between these bimodules (and, in the case of broken metrics, their composites).

Thus unobstructed cobordisms give equivariant chain maps between the tilde complexes, which are well-defined up to $SO(3)$ -equivariant homotopy and compose as expected. We record this as a theorem.

Theorem 4.21. *If $(W, c, \pi_W) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is an unobstructed cobordism equipped with a homology orientation and a path γ between the basepoints, then there is an induced $C_*SO(3)$ -equivariant chain map*

$$\tilde{C}(W, c, \pi_W, \gamma) : \tilde{C}_*(Y, w, \pi) \rightarrow \tilde{C}_*(Y', w', \pi').$$

*This map is independent of π_W up to equivariant chain homotopy; the induced map of the cobordism $(\mathbb{R} \times Y, p^*E, p^*\pi)$ is the identity, and the composite of two unobstructed cobordism maps is equivariantly chain homotopic to the unobstructed cobordism map of the composite.*

Remark 4.22. The statement about the cylinder deserves comment, and is related to the situation of Remark 4.16. The claim follows from the fact that the moduli spaces on the cylinder are either $\check{M}^+(Y, w; \alpha, \beta) \times [0, 1]$ for $\alpha \neq \beta$ — which supports an orientation-reversing isomorphism, hence is zero on the chain level — or the set of constant trajectories α when $\alpha = \beta$. This means that $\mathcal{I}(\mathbb{R} \times Y)$ is the identity bimodule, hence induces the identity map on the flow complexes. ◇

Thus the tilde complex is a sort of functor under cobordism from a certain category of manifolds (Y, w, π, y) with morphisms the unobstructed cobordisms (with a homology orientation and path, but without a particular choice of perturbation) to a derived category of relatively $\mathbb{Z}/8$ -graded $C_*SO(3)$ -complexes.

Because we need to assume the cobordisms are unobstructed, the theorem above does not even imply that the homology $\tilde{I}(Y, w, \pi)$ is independent of the perturbation π . If the signature data are equal $\sigma_\pi = \sigma_{\pi'}$, then the complexes are equivariantly homotopy equivalent, but if $\sigma_\pi > \sigma_{\pi'}$ the cylinder is obstructed and we have not constructed a chain map $\tilde{C}_*(Y, w, \pi) \rightarrow \tilde{C}_*(Y, w, \pi')$; in fact, we will see in Corollary 7.5 that these complexes cannot possibly be equivariantly homotopy equivalent.

We now move on to remedying this unfortunate situation; our strategy is detailed in Section 2.

5 Obstructed cobordism maps

The primary goal of this section is to prove Theorem 2.7.

Convention 5.1. Throughout Section 5, we fix weakly admissible pairs (Y, w) and (Y', w') with regular perturbations π and π' , and a nearly unobstructed cobordism $(W, c) : (Y, w) \rightarrow (Y', w')$ as in the Definition 3.28. This means that $b_1(W) = b^+(W) = 0$ and (W, c) supports exactly one obstructed abelian instanton, which is asymptotic to abelian connections on (Y, w) and (Y', w') . We call this instanton $\Lambda : \rho \rightarrow \rho'$. We fix an perturbation π_W of the ASD equation for which all instantons but Λ are regular and for which the deformation operator of Λ has trivial kernel, as in Theorem 3.31. We write $\lambda : \rho \rightarrow \rho'$ for the path of connections over (W, c) given by Λ . \diamond

Due to the presence of the obstructed abelian instanton Λ , we cannot proceed as in subsection 4.3 to define a morphism $\tilde{C}(W, c, \pi_W)$: the standard gluing theory does not apply, and the compactified moduli spaces $M_z^+(W, c; \alpha, \alpha')$ are not stratified-smooth spaces anymore. In Subsection 5.1, we use *obstructed* gluing theory to analyze the moduli spaces $M_z^+(W, c; \alpha, \alpha')$ in the neighborhood of obstructed (broken) solutions.

In Section 5.2, we show that *modified moduli spaces* $N_z^+(W, c; \alpha, \alpha')$, obtained by deleting a neighborhood of any broken instanton factoring through Λ and adding a ‘correction term’, are indeed stratified-smooth spaces which satisfy an explicit boundary relation. However, they do not define bimodules between the instanton flow categories \mathcal{I} : the boundary relation is not the expected one.

Motivated by this analysis, we define the suspended flow category $\mathcal{B}_{\rho'}\mathcal{I}(Y', w', \pi')$ associated to $\mathcal{I}(Y', w', \pi')$ and ρ' , and then we show that the modified moduli spaces $N_z^+(W, c; \alpha, \alpha')$ give rise to a bimodule from $\mathcal{I}(Y, w, \pi)$ to $\mathcal{B}_{\rho'}\mathcal{I}(Y', w', \pi')$. The definition of the blown up flow category $\mathcal{B}_{\rho'}\mathcal{I}(Y', w', \pi')$ and the bimodule associated to (W, c) are respectively given in Subsections 5.3 and 5.4.

With these in hand, in Section 5.5 we apply these to $W = \mathbb{R} \times Y$, with π_0 and π_1 regular perturbations on (Y, w) with adjacent signature data $\sigma_{\pi_0} < \sigma_{\pi_1}$ (meaning they differ by the smallest possible amount, 4, at exactly one reducible ρ). To show that the natural continuation map $\tilde{C}(Y, w, \pi_0) \rightarrow \tilde{C}(Y, w, \pi_1)$ is a quasi-isomorphism, we use the constructions above to lift this to maps to and from from the *suspended complex* $S_{\rho}\tilde{C}(Y, w, \pi_0) \leftrightarrow \tilde{C}(Y, w, \pi_1)$, where the first complex is the flow complex of the blown-up flow category $\mathcal{B}_{\rho}\mathcal{I}(Y, w, \pi_0)$. Using similar ideas, we conclude by showing that these two maps are chain homotopy inverse to one another. Because the continuation map described above factors as

$$\tilde{C}(Y, w, \pi_0) \rightarrow S_{\rho}(Y, w, \pi_0) \rightarrow \tilde{C}(Y, w, \pi_1),$$

and it is easy to verify that the first map is a quasi-isomorphism, this gives the desired result (and with only slightly more work, invariance of instanton homology).

5.1 Obstructed cobordisms and obstructed gluing theory

Before we define the ‘modified moduli spaces’, we should understand the local structure of $M_z^+(W, c; \alpha, \alpha')$ in a neighborhood of the non-smooth locus of obstructed reducibles.

The moduli space $M_z^+(W, c; \alpha, \alpha')$ has three different types of obstructed elements. An *obstructed solution of type I* is a broken solution which belongs to the subspace

$$\check{M}_\zeta^+(Y, w; \alpha, \rho) \times_\rho \mathcal{O}_\Lambda \cong \check{M}_\zeta^+(Y, w; \alpha, \rho) \times \{\Lambda\}. \quad (5.2)$$

of a moduli space of the form $M_z^+(W, c; \alpha, \rho')$ where gluing the path ζ to the path over (W, c) represented by Λ is z . Similarly, an *obstructed solution of type II* is a broken solution in the subspace

$$\mathcal{O}_\Lambda \times_{\rho'} \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha') \cong \{\Lambda\} \times \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha') \quad (5.3)$$

of a moduli space of the form $M_z^+(W, c; \rho, \alpha')$. Finally an *obstructed solution of type III* is a broken solution in

$$\check{M}_\zeta^+(Y, w; \alpha, \rho) \times_\rho \mathcal{O}_\Lambda \times_{\rho'} \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha') \quad (5.4)$$

which is regarded as a subspace of $M_z^+(W, c; \alpha, \alpha')$ with $z = \zeta * \lambda * \zeta'$.

In the complement of these obstructed solutions, the stratified spaces $M_z^+(W, c; \alpha, \alpha')$ have the structure of stratified-smooth manifolds (as in Definition 3.17).

To study the behavior of the moduli space $M_z^+(W, c; \alpha, \alpha')$ in a neighborhood of an obstructed solution, we need to introduce a family of complex line bundles which are called *obstruction bundles*. The cokernel of the ASD operator for the elements of \mathcal{O}_Λ define a complex line bundle \mathcal{H}_Λ^+ over \mathcal{O}_Λ with an $SO(3)$ action that lifts the action on \mathcal{O}_Λ . Let $\mathcal{H}_{\alpha, \rho}$ be the line bundle over $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times_\rho \mathcal{O}_\Lambda$ pulled back from \mathcal{H}_Λ^+ . We call $\mathcal{H}_{\alpha, \rho}$ the *obstruction bundle* associated to Λ . Similarly, we have the complex line bundle $\mathcal{H}_{\rho', \alpha'}$ over the space $\check{M}_{\zeta'}^+(Y', w'; \rho', \alpha')$.

The following proposition, which is an instance of obstructed gluing theory, gives a description for a neighborhood of an obstructed solution of type I.

Proposition 5.5. *Suppose (\mathbf{A}_0, Λ) is an obstructed solution of type I in $M_z^+(W, c; \alpha, \rho')$ where $\mathbf{A}_0 \in \check{M}_\zeta^+(Y, w; \alpha, \rho)$ with $z = \zeta * \lambda$ given as above and $i(z) \geq -1$. Then there is an $SO(3)$ -invariant neighborhood U of \mathbf{A}_0 in $\check{M}_\zeta^+(Y, w; \alpha, \rho)$ and an $SO(3)$ -equivariant section $\psi_{\zeta\Lambda}^{\text{an}}$ of $\pi_1^*(\mathcal{H}_{\alpha, \rho})$ over $U \times (0, \infty]$ which vanishes on $U \times \{\infty\}$ and is transverse to the zero section over $\mathcal{H}_{\alpha, \rho} \times (0, \infty)$. There is also an $SO(3)$ -equivariant open embedding $\Phi_{\zeta\Lambda} : \left(\psi_{\zeta\Lambda}^{\text{an}}\right)^{-1}(0) \rightarrow M_z^+(W, c; \alpha, \rho')$ such that $\Phi_{\zeta\Lambda}$ maps $U \times \{\infty\} \subset \left(\psi_{\zeta\Lambda}^{\text{an}}\right)^{-1}(0)$ to $U \times \{\Lambda\} \subset M_z^+(W, c; \alpha, \rho')$ via the identity map on U .*

The superscript ‘an’ is short for *analytic*, as these sections are defined by the analytic machinery of obstructed gluing theory. In contrast, in the next section we will *choose* privileged sections $\psi_{\zeta\Lambda}^{\text{priv}}$.

For $i(z) = -1$, the section $\psi_{\zeta\Lambda}^{\text{an}}$ is nowhere vanishing and $M_z^+(W, c; \alpha, \rho')$ is empty, so there is no content in this case. However, later we need compatibility of the sections $\psi_{\zeta\Lambda}^{\text{an}}$ for different ζ and z , and it is necessary for us to include the case that $i(z) = -1$.

Sketch of the proof. This proposition can be verified using standard obstructed gluing theory results for which there is a large body of literature starting with the seminal work of Taubes in [Tau82]. Here we

sketch the general steps in the proof and refer the reader to [DK90, Chapter 7] (more precisely, the discussion following Theorem 7.2.24 and the discussion in Chapter 7.2.8) as well as [Don02, Chapter 4.5] for more details. The reader should note that the discussion of [DK90] is limited to the case that $Y = S^3$ and the arguments are worked out in a slightly different conformal model than cylindrical ends. However, working in the setup of 4-manifolds with cylindrical ends in fact simplifies the arguments, and the arguments of [DK90, Chapter 7] can be adapted to the case of a weakly admissible pair (Y, w) where $\mathfrak{C}_\pi(Y, w)$ is non-degenerate. For the ease of exposition, we shall also assume that the perturbation terms in the definition of the critical sets $\mathfrak{C}_\pi(Y, w)$, $\mathfrak{C}_\pi(Y', w')$ and the moduli spaces of instantons are trivial; it is straightforward to modify the arguments below to handle the general case.

As the first step in the proof, one needs to construct approximate solutions of the ASD equation by pre-gluing elements of $\check{M}_\zeta(Y, w; \alpha, \rho)$ and \mathcal{O}_Λ . Fix an $SO(3)$ -invariant pre-compact and open neighborhood U of \mathbf{A}_0 in $\check{M}_\zeta(Y, w; \alpha, \rho)$. Fix once and for all a connection $\hat{\rho}$ representing the gauge equivalence class of the reducible ρ , and similarly a connection $\hat{\Lambda}$ representing the gauge equivalence class of obstructed reducible connection on W . Fix positive real numbers T_0 and ε_0 which are respectively required to be large and small. For any $SO(3)$ -orbit \mathbf{A} in U , we may find a representative connection $\hat{\mathbf{A}}$ such that $\hat{\mathbf{A}}$ is in temporal gauge, $\hat{\mathbf{A}}$ is asymptotic to $\hat{\rho}$ and

$$\int_{[-T_0, \infty) \times Y} |F_{\hat{\mathbf{A}}}|^2 = \varepsilon_0.$$

The choice of the connection $\hat{\mathbf{A}}$ is unique up to the action of an element of the stabilizer of $\hat{\rho}$.

For any $T \in (0, \infty)$, we may glue the connection $\hat{\mathbf{A}}$ to $\hat{\Lambda}$ to define a connection $\hat{\mathbf{A}}\#_T\hat{\Lambda}$ on W . To perform this pre-gluing construction, let $\chi : \mathbb{R} \rightarrow [0, 1]$ be a smooth function such that $\chi(t) = 0$ for $t \leq -1$ and $\chi(t) = 1$ for $t \geq 0$. Then the connection which we call $\hat{\mathbf{A}}\#_T\hat{\Lambda}$ is given over the cylindrical end $(-\infty, 0) \times Y \subset W$ by

$$\chi(s+T)\hat{\Lambda} + (1 - \chi(s+T))\tau_{2T}^*(\hat{\mathbf{A}})$$

where $\tau_T : \mathbb{R} \times Y \rightarrow \mathbb{R} \times Y$ is the map given by $\tau_T(s, y) = (s+T, y)$. In the remaining part of W , $\hat{\mathbf{A}}\#_T\hat{\Lambda}$ agrees with $\hat{\Lambda}$.

Let us briefly discuss how to carry out this constructions for the whole neighborhood $U \subset \check{M}_\zeta$ in the space of framed connections. First, there is an open subset $\hat{U} \subset \tilde{\mathcal{A}}_\zeta(Y, w)$, consisting of framed connections $(\hat{\mathbf{A}}, \Phi)$ so that the equivalence class $[\hat{\mathbf{A}}, \Phi]$ lies in U , and $\hat{\mathbf{A}}$ satisfies the integral equality above, has asymptotic $\hat{\rho}$, and is in temporal gauge. Because the framing is chosen arbitrarily, this is an $SO(3)$ -invariant subset of the space of framed connections; and because of the $\Gamma_{\hat{\rho}}$ ambiguity in the lift $\hat{\mathbf{A}}$, the projection $\pi : \hat{U} \rightarrow U$ is a $\Gamma_{\hat{\rho}}$ -bundle.

We write $\mathcal{O}_{\hat{\rho}}$ for the $SO(3)$ -orbit of pairs $(\hat{\rho}, \Phi)$ where $\hat{\rho}$ is the fixed lift above and Φ is a chose of framing on (Y, w) ; we similarly write $\mathcal{O}_{\hat{\Lambda}}$ for pairs $(\hat{\Lambda}, \Phi_W)$ where $\hat{\Lambda}$ is the chosen lift above and Φ_W is chosen arbitrarily. Finally, the gluing map described above defines for each T a map $\hat{U} \times_{\mathcal{O}_{\hat{\rho}}} \mathcal{O}_{\hat{\Lambda}} \rightarrow \tilde{\mathcal{A}}_z(W, c)$, defined by

$$((\hat{\mathbf{A}}, \Phi), (\hat{\Lambda}, \Phi_W)) \mapsto (\hat{\mathbf{A}}\#_T\hat{\Lambda}, \Phi_W).$$

It is transparent from the construction above that this map is equivariant under the action of $\Gamma_{\hat{\rho}} \times_{\Gamma_{\hat{\rho}}} \Gamma_{\hat{\Lambda}} \cong S^1$, and hence descends to an $SO(3)$ -equivariant map from the quotient $U \times_{\mathcal{O}_{\hat{\rho}}} \mathcal{O}_{\hat{\Lambda}} \cong U$ to $\tilde{\mathcal{B}}_z(W, c)$.

Having described how to do this construction for families, we henceforth focus on one instanton at a time; the rest of the discussion can be carried out in families as well. To simplify notation, we write $\mathbf{A}\#_T\Lambda$ for an element of the S^1 -orbit of framed connections produced by the construction above, suppressing the framing from notation.

The framed connection $\mathbf{A}\#_T\Lambda$ is an approximate instanton, and we want to perturb it to obtain an actual solution of the ASD equation. To achieve this, we need to construct right inverses for the ASD operators. Since the connection \mathbf{A} is regular, the linearization of the ASD equation

$$d_{\mathbf{A}}^+ : \Omega^1(\mathbb{R} \times Y, \text{ad}(\mathbf{E}_w)) \rightarrow \Omega^+(\mathbb{R} \times Y, \text{ad}(\mathbf{E}_w)) \quad (5.6)$$

is surjective. After taking appropriate Sobolev completions of the above function spaces, we fix a right inverse $R_{\mathbf{A}}$ for (5.6). On the other hand, the operator d_{Λ}^+ has non-trivial cokernel. We fix a linear map

$$\sigma_{\Lambda} : H^+(\Lambda) \rightarrow \Omega^+(W, \text{ad}(\mathbf{E}_c))$$

which is supported away from the cylindrical ends of W so that the operator

$$d_{\Lambda}^+ + \sigma_{\Lambda} : \Omega^1(W, \text{ad}(\mathbf{E}_c)) \oplus H^+(\Lambda) \rightarrow \Omega^+(W, \text{ad}(\mathbf{E}_c)),$$

which maps (a, h) to $d_{\Lambda}^+(a) + \sigma_{\Lambda}(h)$ is surjective. Again after taking appropriate Sobolev completions of the space of 1-forms and self-dual forms on W , we may find a right inverse R_{Λ} . Linear gluing theory (see, for example, [Don02, Chapter 3.3]) allows us to obtain a right inverse $Q_{\mathbf{A}}^T$ for the operator

$$d_{\mathbf{A}\#_T\Lambda}^+ + \sigma_{\Lambda} : \Omega^1(W, \text{ad}(\mathbf{E}_c)) \oplus H^+(\Lambda) \rightarrow \Omega^+(W, \text{ad}(\mathbf{E}_c)),$$

which depends continuously on T . Note that here we use the assumption that the target of σ_{Λ} is away from the gluing region and hence the above map is well-defined. We also fix the notation

$$Q_{\mathbf{A},1}^T : \Omega^+(W, \text{ad}(\mathbf{E}_c)) \rightarrow \Omega^1(W, \text{ad}(\mathbf{E}_c)), \quad Q_{\mathbf{A},H}^T : \Omega^+(W, \text{ad}(\mathbf{E}_c)) \rightarrow H^+(\Lambda),$$

for the components of $Q_{\mathbf{A}}^T$. (As before, we need to consider appropriate Sobolev completions of the function spaces in the domain and the co-domain of the above operators.)

To perturb the connections $\mathbf{A}\#_T\Lambda$ to ASD connections, first we consider the weaker equation

$$F^+(\mathbf{A}\#_T\Lambda + Q_{\mathbf{A},1}^T(\phi)) + \sigma_{\Lambda}(Q_{\mathbf{A},H}^T(\phi)) = 0.$$

A solution to this equation is equivalent to a fixed point of the operator

$$Z_{\mathbf{A}}^T(\phi) = -F^+(\mathbf{A}\#_T\Lambda) - (Q_{\mathbf{A},1}^T(\phi) \wedge Q_{\mathbf{A},1}^T(\phi))^+.$$

One can apply the contraction mapping principle to conclude that this equation has a unique solution for any $\mathbf{A} \in U$ and any $T \in (0, \infty)$. (Note that we are still using the assumption that T_0 is a large constant, so that $-F^+(\mathbf{A}\#_T\Lambda)$ may be assumed to be arbitrarily small.) If ϕ_0 is the solution of this equation, then $Q_{\mathbf{A},H}^T(\phi_0)$ determines an S^1 -equivariant map from $\tilde{U} \times (0, \infty)$ to $H^+(\Lambda)$. This induces a section $\psi_{\zeta\Lambda}^{\text{an}}$ of the bundle $\mathcal{H}_{\alpha,\rho}$ over $U \times (0, \infty)$, and the honest solutions of the ASD equation are given by the zeros of this section. This section extends continuously as zero over the space $U \times \{\infty\}$. The desired embedding $\Phi_{\zeta\Lambda}$ is given by mapping (\mathbf{A}, T) in the zero set of $\psi_{\zeta\Lambda}^{\text{an}}$ to $\mathbf{A}\#_T\Lambda + \phi_0$. \square

We need a global version of Proposition 5.5. In particular, we wish to consider an extension of the above proposition to the case that $\mathbf{A} \in \check{M}_\zeta(Y, w; \alpha, \rho)$ is replaced with a broken trajectory in $\check{M}_\zeta(Y, w; \alpha, \rho)$. This proposition is the first place we are obligated to use stratified-smooth spaces, as here we take the zero set of a continuous section of a vector bundle which is smooth on each stratum; a priori, there is no reason to believe the result is a stratified-smooth manifold. However, the definition of stratified-smooth spaces is chosen so that such zero sets are always stratified-smooth.

Proposition 5.7. *Suppose z and ζ are given as above and $-1 \leq i(z) \leq 6$. Then the following holds.*

- (i) *There is an $SO(3)$ -equivariant continuous section $\psi_{\zeta\Lambda}^{\text{an}}$ of $\pi_1^*(\mathcal{H}_{\alpha,\rho})$ over $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times (0, \infty]$ which vanishes on the subspace $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times \{\infty\}$, and whose restriction to each stratum of $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times (0, \infty)$ is smooth and transverse to the zero section. In particular, the intersection of $(\psi_{\zeta\Lambda}^{\text{an}})^{-1}(0)$ and $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times (0, \infty)$ is again a stratified-smooth space.*
- (ii) *There is an injective open continuous map $\Phi_{\zeta\Lambda} : (\psi_{\zeta\Lambda}^{\text{an}})^{-1}(0) \rightarrow M_z^+(W, c; \alpha, \rho')$ such that $\Phi_{\zeta\Lambda}$ maps the subspace $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times \{\infty\} \subset (\psi_{\zeta\Lambda}^{\text{an}})^{-1}(0)$ to $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times \{\Lambda\} \subset M_z^+(W, c; \alpha, \rho')$ via the identity map. Moreover, the restriction of $\Phi_{\zeta\Lambda}$ to the intersection of $(\psi_{\zeta\Lambda}^{\text{an}})^{-1}(0)$ and $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times (0, \infty)$ gives an open embedding of stratified-smooth spaces.*
- (iii) *For $\beta \in \mathfrak{C}_\pi(Y, w)$, let $\zeta_0 : \alpha \rightarrow \beta$ and $\zeta_1 : \beta \rightarrow \rho$ be paths of connections on (Y, w) with composite $\zeta_0 * \zeta_1 = \zeta$. For any $[x, y]$ in the subspace*

$$\check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_\beta \check{M}_{\zeta_1}^+(Y, w; \beta, \rho) \tag{5.8}$$

of $\check{M}_\zeta^+(Y, w; \alpha, \rho)$, we have

$$\psi_{\zeta\Lambda}^{\text{an}}([x, y, \Lambda], t) = \psi_{\zeta_1\Lambda}^{\text{an}}([y, \Lambda], t), \tag{5.9}$$

and if $\psi_{\beta\Lambda}^{\text{an}}([y, \Lambda], t) = 0$, then

$$\Phi_{\zeta\Lambda}([x, y, \Lambda], t) = (x, \Phi_{\zeta_1\Lambda}([y, \Lambda], t)).$$

To interpret (5.9), note that the restriction of $\mathcal{H}_{\alpha,\rho}$ to the subspace of $\check{M}_\zeta^+(Y, w; \alpha, \rho)$ given in (5.8) is induced by the pull-back of $\mathcal{H}_{\alpha,\rho}$. In particular, the equivariant section $\psi_{\beta\Lambda}^{\text{an}}$ of $\mathcal{H}_{\alpha,\rho}$ determines a section of $\mathcal{H}_{\alpha,\rho}$ restricted to (5.8).

Sketch of the proof. This proposition is a global version of Proposition 5.5 and can be verified by constructing the section $\psi_{\zeta\Lambda}^{\text{an}}$ and the map $\Phi_{\zeta\Lambda}$ in a neighborhood of points in dimension i strata of $\check{M}_\zeta^+(Y, w; \alpha, \rho)$ by induction on i . The induction step can be addressed by an argument similar to Proposition 5.5. The transversality of $\psi_{\zeta\Lambda}^{\text{an}}$ follows from our assumption on regularity of the elements of $M_z^+(W, c; \alpha, \rho')$ away from obstructed solutions. The compatibility assumption in (iii) can be verified as in [DK90, Proposition 7.2.64]. \square

There is an analogue of Proposition 5.7 for obstructed solutions of type II in the moduli spaces of the form $M_z^+(W, c; \rho, \alpha')$. In particular, a neighborhood of such obstructed solutions in $M_z^+(W, c; \rho, \alpha')$ are described using an $SO(3)$ -equivariant section $\psi_{\Lambda\zeta'}^{\text{an}}$ of $\mathcal{H}_{\rho', \alpha'} \times [-\infty, 0)$ over $\check{M}_{\zeta'}^+(Y', w'; \rho', \alpha') \times [-\infty, 0)$ and a map $\Phi_{\Lambda\zeta'} : \left(\psi_{\Lambda\zeta'}^{\text{an}}\right)^{-1}(0) \rightarrow M_z^+(W, c; \rho, \alpha')$. These satisfy the analogues of properties (i)-(iii) in Proposition 5.7.

The following proposition concerns the obstructed solutions of type III and describes the behavior of the moduli spaces $M_z^+(W, c; \alpha, \alpha')$ in a neighborhood of such obstructed solutions. The pull-back of the bundle \mathcal{H}_Λ determines a complex line bundle over

$$X_{\zeta, \zeta'} := \check{M}_\zeta^+(Y, w; \alpha, \rho) \times_\rho \mathcal{O}_\Lambda \times_{\rho'} \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha'),$$

which is denoted by $\mathcal{H}_{\alpha, \alpha'}$. It is clear from the definition of $\mathcal{H}_{\alpha, \alpha'}$ that it is the pullback of the bundles $\mathcal{H}_{\alpha, \rho}$ and $\mathcal{H}_{\rho', \alpha'}$ with respect to the obvious projection maps. Thus sections of $\mathcal{H}_{\alpha, \rho}$ and $\mathcal{H}_{\rho', \alpha'}$ can be also regarded as sections of $\mathcal{H}_{\alpha, \alpha'}$, and this convention is used in the statement of the following proposition about the behavior of the moduli spaces around obstructed solutions of type III.

We omit the proof, as it follows essentially the same lines as Proposition 5.7.

Proposition 5.10. *Suppose z, ζ and ζ' are given as above and $0 \leq i(z) \leq 7$. Then the following holds.*

- (i) *There is an $SO(3)$ -equivariant section $\psi_{\zeta\Lambda\zeta'}^{\text{an}}$ of $\pi_1^*(\mathcal{H}_{\alpha, \alpha'})$ over $X_{\zeta, \zeta'} \times (0, \infty] \times [-\infty, 0)$ whose restriction to $X_{\zeta, \zeta'} \times ((0, \infty] \times [-\infty, 0)) \setminus \{(\infty, -\infty)\}$ is transverse to the zero section. For any $[x, \Lambda, x'] \in X_{\zeta, \zeta'}$, we have*

$$\psi_{\zeta\Lambda\zeta'}^{\text{an}}([x, \Lambda, x'], t, -\infty) = \psi_{\zeta\Lambda}^{\text{an}}([x, \Lambda], t), \quad \psi_{\zeta\Lambda\zeta'}^{\text{an}}([x, \Lambda, x'], \infty, t) = \psi_{\Lambda\zeta'}^{\text{an}}([\Lambda, x'], t).$$

In particular, $\psi_{\zeta\Lambda\zeta'}^{\text{an}}$ vanishes on $X_{\zeta, \zeta'} \times \{(\infty, -\infty)\}$ and the intersection of the complement of this space and $\left(\psi_{\zeta\Lambda\zeta'}^{\text{an}}\right)^{-1}(0)$ is a stratified-smooth space.

- (ii) *There is an injective open continuous map $\Phi_{\zeta\Lambda\zeta'} : \left(\psi_{\zeta\Lambda\zeta'}^{\text{an}}\right)^{-1}(0) \rightarrow M_z^+(W, c; \alpha, \alpha')$ such that $\Phi_{\zeta\Lambda\zeta'}$ maps $X_{\zeta, \zeta'} \times \{(\infty, -\infty)\}$ to $X_{\zeta, \zeta'} \subset M^+(W; \alpha, \alpha')$ via the identity map. Moreover, the restriction of $\Phi_{\zeta\Lambda\zeta'}$ to the intersection of $X_{\zeta, \zeta'} \times ((0, \infty] \times [-\infty, 0)) \setminus \{(\infty, -\infty)\}$ and $\left(\psi_{\zeta\Lambda\zeta'}^{\text{an}}\right)^{-1}(0)$ gives an open embedding of stratified-smooth spaces.*

- (iii) *For $\beta \in \mathfrak{C}_\pi(Y, w)$, let $\zeta_0 : \alpha \rightarrow \beta$ and $\zeta_1 : \beta \rightarrow \rho$ be paths of connections on (Y, w) with composite $\zeta_0 * \zeta_1 = \zeta$. For any $[x, y, \Lambda, x']$ in the subspace*

$$\check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_\beta \check{M}_{\zeta_1}^+(Y, w; \beta, \rho) \times_\rho \mathcal{O}_\Lambda \times_{\rho'} \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha')$$

of $X_{\zeta, \zeta'}$, we have

$$\psi_{\zeta\Lambda\zeta'}^{\text{an}}([x, y, \Lambda, x'], t, t') = \psi_{\zeta_1\Lambda\zeta'}^{\text{an}}([y, \Lambda, x'], t, t'),$$

and if $\psi_{\beta\Lambda\zeta'}^{\text{an}}([y, \Lambda, x'], t, t') = 0$, then

$$\Phi_{\zeta\Lambda\zeta'}([x, y, \Lambda, x'], t, t') = (x, \Phi_{\zeta_1\Lambda\zeta'}([y, \Lambda, x'], t, t')).$$

A similar claim holds for any $\beta' \in \mathfrak{C}_\pi(Y', w')$, paths of connections $\zeta'_0 : \rho' \rightarrow \beta'$ and $\zeta'_1 : \beta' \rightarrow \alpha'$ on (Y', w') with composite $\zeta'_0 * \zeta'_1 = \zeta'$, and any $[x, \Lambda, x', y']$ in the following subspace of $X_{\zeta, \zeta'}$:

$$\check{M}_\zeta^+(Y, w; \alpha, \rho) \times_\rho \mathcal{O}_\Lambda \times_{\rho'} \check{M}_{\zeta'_0}^+(Y', w'; \rho', \beta') \times_{\beta'} \check{M}_{\zeta'_1}^+(Y', w'; \beta', \alpha')$$

(iv) Let $[x, \Lambda, x'] \in X_{\zeta, \zeta'}$. If $\psi_{\zeta\Lambda}^{\text{an}}([x, \Lambda], t) = 0$, we have

$$\Phi_{\zeta\Lambda\zeta'}([x, \Lambda, x'], t, -\infty) = [\Phi_{\Lambda\zeta'}([x, \Lambda], t), x'].$$

Similarly, if $\psi_{\Lambda\zeta'}^{\text{an}}([\Lambda, y], t) = 0$, we have

$$\Phi_{\zeta\Lambda\zeta'}([x, \Lambda, x'], \infty, t) = [x, \Phi_{\Lambda\zeta'}([\Lambda, x'], t)].$$

5.2 Modified moduli spaces

In the previous section, we investigated the local structure of the obstructed reducibles in $M_z^+(W, c; \alpha, \alpha')$. In this subsection, we modify the moduli spaces $M_z^+(W, c; \alpha, \alpha')$ so that the resulting moduli spaces are stratified-smooth, and so that we have explicit control over their boundary components. This modification involves two steps: removing a neighborhood of obstructed solutions and adding several new components to the moduli spaces. The first step of this process is performed using the description of neighborhoods of obstructed solutions given in the last subsection. The second step involves choosing some additional sections of the obstruction bundles, which we will fix next. Throughout this section z is a path of connections along (W, c) with $i(z) \leq 6$.

We write $\psi_{\zeta, \Lambda}^T$ and $\psi_{\Lambda, \zeta'}^T$ for the restrictions of the sections $\psi_{\zeta, \Lambda}^{\text{an}}$ and $\psi_{\Lambda, \zeta'}^{\text{an}}$ over $\check{M}_\zeta^+(Y, w; \alpha, \rho) \times \{T\}$ and $\check{M}_{\zeta'}^+(Y', w'; \rho', \alpha') \times \{-T\}$ for some large value $T \in (0, \infty)$ that will be specified later. We fix an $SO(3)$ -equivariant section $\psi_{\zeta\Lambda}^{\text{priv}}$ of $\mathcal{H}_{\alpha, \rho}$ over $\check{M}_\zeta^+(Y, w; \alpha, \rho)$.

Unlike $\psi_{\zeta\Lambda}^T$, the section $\psi_{\zeta\Lambda}^{\text{priv}}$ is not determined by analysis, but is rather pre-chosen (and independent of the choice of T). In particular, we regard $\psi_{\zeta\Lambda}^{\text{priv}}$ as an additional information attached to (Y, w) , rather than being determined by the cobordism (W, c) .

We also need to fix a section that interpolates between $\psi_{\zeta\Lambda}^T$ and $\psi_{\zeta\Lambda}^{\text{priv}}$. Let $\psi_{\zeta\Lambda}^{\text{htpy}}$ be a section of $[0, 1] \times \mathcal{H}_{\alpha, \rho}$ over $[0, 1] \times \check{M}_\zeta^+(Y, w; \alpha, \rho)$ whose restrictions to $\{0\} \times \check{M}_\zeta^+(Y, w; \alpha, \rho)$ and $\{1\} \times \check{M}_\zeta^+(Y, w; \alpha, \rho)$ are respectively equal to $\psi_{\zeta\Lambda}^T$ and $\psi_{\zeta\Lambda}^{\text{priv}}$.

Similarly, we fix $\psi_{\Lambda\zeta'}^{\text{priv}}$ and $\psi_{\Lambda\zeta'}^{\text{htpy}}$ that are respectively sections of $\mathcal{H}_{\rho', \alpha'}$ and $[0, 1] \times \mathcal{H}_{\rho', \alpha'}$. Moreover, the restriction of $\psi_{\Lambda, \alpha'}^{\text{htpy}}$ to $\{0\} \times \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha')$ and $\{1\} \times \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha')$ are respectively equal to $\psi_{\Lambda\zeta'}^T$ and $\psi_{\Lambda\zeta'}^{\text{priv}}$. The sections fixed in this paragraph are required to satisfy the following properties:

- (s.1) The number T is chosen such that $\psi_{\zeta\Lambda}^T$ and $\psi_{\Lambda\zeta'}^T$ are stratum-wise transverse to the zero section for any ζ and ζ' . We also demand that the various sections ψ^{priv} and ψ^{htpy} are stratum-wise transverse to the zero section.

(s.2) For any pair ζ, ζ' , we may regard $t \cdot \psi_{\zeta\Lambda}^{\text{priv}} + (1-t) \cdot \psi_{\Lambda\zeta'}^{\text{priv}}$ with $t \in [0, 1]$ as a section of $\pi_1^*(\mathcal{H}_{\alpha, \alpha'})$ defined over the stratified-smooth manifold $[0, 1] \times X_{\zeta, \zeta'}$. This section is required to be transverse to the zero section on each stratum.

(s.3) For $\beta \in \mathfrak{C}_\pi(Y, w)$, let $\zeta_0 : \alpha \rightarrow \beta$ and $\zeta_1 : \beta \rightarrow \rho$ be paths of connections on (Y, w) with composite $\zeta_0 * \zeta_1 = \zeta$. The restriction of $\psi_{\zeta\Lambda}^{\text{priv}}$ to

$$\check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_{\beta} \check{M}_{\zeta_1}^+(Y, w; \beta, \rho) \subset \check{M}_{\zeta}^+(Y, w; \alpha, \rho) \quad (5.11)$$

is given by $\psi_{\zeta_1\Lambda}^{\text{priv}}$. Similarly, the restriction of $\psi_{\zeta\Lambda}^{\text{htpy}}$ to

$$[0, 1] \times \check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_{\beta} \check{M}_{\zeta_1}^+(Y, w; \beta, \rho) \subset [0, 1] \times \check{M}_{\zeta}^+(Y, w; \alpha, \rho)$$

is given by $\psi_{\zeta_1\Lambda}^{\text{htpy}}$.

(s.4) For $\beta' \in \mathfrak{C}_\pi(Y', w')$, let $\zeta'_0 : \rho' \rightarrow \beta'$ and $\zeta'_1 : \beta' \rightarrow \alpha'$ be paths of connections on (Y', w') with composite $\zeta'_0 * \zeta'_1 = \zeta'$. The restriction of $\psi_{\Lambda\zeta'}^{\text{priv}}$ to

$$\check{M}_{\zeta'_0}^+(Y', w'; \rho', \beta') \times_{\beta'} \check{M}_{\zeta'_1}^+(Y', w'; \beta', \alpha') \subset \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha') \quad (5.12)$$

is given by $\psi_{\Lambda\zeta'_0}^{\text{priv}}$. Similarly, the restriction of $\psi_{\Lambda\zeta'}^{\text{htpy}}$ to

$$[0, 1] \times \check{M}_{\zeta'_0}^+(Y', w'; \rho', \beta') \times_{\beta'} \check{M}_{\zeta'_1}^+(Y', w'; \beta', \alpha') \subset [0, 1] \times \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha')$$

is given by $\psi_{\Lambda\zeta'_0}^{\text{htpy}}$.

We shall also need some variations of the sections $\psi_{\zeta\Lambda\zeta'}^{\text{an}}$. First let

$$A_T = \left\{ (t, t') \in (0, \infty) \times [-\infty, 0) \mid \frac{1}{t'} - \frac{1}{t} = -\frac{1}{T} \right\},$$

which is oriented as an inverse image of $1/t' - 1/t$, so that its boundary has the point $(\infty, -T)$ with negative sign and the point $(T, -\infty)$ with positive sign. (This is named A_T because it is an arc between these two points.) This is oriented diffeomorphic to the unit interval $[0, 1]$ via the map $(t, t') \mapsto T/t = 1 + T/t'$.

Define $\psi_{\zeta\Lambda\zeta'}^T$ be the restriction of $\psi_{\zeta\Lambda\zeta'}^{\text{an}}$ to $A_T \times X_{\zeta, \zeta'}$. Let $\psi_{\zeta\Lambda\zeta'}^{\text{priv}}$ be the section over $A_T \times X_{\zeta, \zeta'}$ defined by

$$\psi_{\zeta\Lambda\zeta'}^{\text{priv}}([x, y], t, t') = \frac{T}{t} \cdot \psi_{\zeta\Lambda}^{\text{priv}}(x) - \frac{T}{t'} \cdot \psi_{\Lambda\zeta'}^{\text{priv}}(y).$$

Via the diffeomorphism $A_T \cong [0, 1]$ above, this map is identified with the section $s\psi_{\zeta\Lambda}^{\text{priv}} + (1-s)\psi_{\Lambda\zeta'}^{\text{priv}}$. Our assumption (s.2) on the sections $\psi_{\zeta\Lambda}^{\text{priv}}$ and $\psi_{\Lambda\zeta'}^{\text{priv}}$ imply that $\psi_{\zeta\Lambda\zeta'}^{\text{priv}}$ is transverse to the zero section on each stratum.

Finally, let $\psi_{\zeta\Lambda\zeta'}^{\text{htpy}}$ be a section of $\pi_1^*(\mathcal{H}_{\alpha, \alpha'})$ over $[0, 1] \times A_T \times X_{\zeta, \zeta'}$ that satisfies the following properties:

(c.1) As another requirement on T , we demand that $\psi_{\zeta\Lambda\zeta'}^T$ is transverse to the zero section on each stratum. We also demand that the section $\psi_{\zeta\Lambda\zeta'}^{\text{htpy}}$ is also transverse to the zero section on each stratum.

(c.2) The restriction of $\psi_{\zeta,\Lambda,\zeta'}^{\text{htpy}}$ to $\{0\} \times A_T \times X_{\zeta,\zeta'}$ and $\{1\} \times A_T \times X_{\zeta,\zeta'}$ are respectively equal to $\psi_{\zeta\Lambda\zeta'}^T$ and $\psi_{\zeta\Lambda\zeta'}^{\text{priv}}$.

(c.3) For any $[x, y] \in X_{\zeta,\zeta'}$ and $t \in [0, 1]$, we have

$$\psi_{\zeta\Lambda\zeta'}^{\text{htpy}}(t, T, -\infty, [x, y]) = \psi_{\zeta\Lambda}^{\text{htpy}}(t, [x]), \quad \psi_{\zeta\Lambda\zeta'}^{\text{htpy}}(t, \infty, T, [x, y]) = \psi_{\Lambda\zeta'}^{\text{htpy}}(t, [y]).$$

(c.4) For $\beta \in \mathfrak{C}_\pi(Y, w)$, let $\zeta_0 : \alpha \rightarrow \beta$ and $\zeta_1 : \beta \rightarrow \rho$ be paths of connections on (Y, w) with composite $\zeta_0 * \zeta_1 = \zeta$. Let $[x, y, \Lambda, x']$ be an element of the subspace

$$\check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_{\beta} \check{M}_{\zeta_1}^+(Y, w; \beta, \rho) \times_{\rho} \mathcal{O}_{\Lambda} \times_{\rho'} \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha') \subset X_{\zeta,\zeta'}.$$

Then for any $(s, s') \in A_T$ and $t \in [0, 1]$, we have

$$\psi_{\zeta\Lambda\zeta'}^{\text{htpy}}(t, s, s', [x, y, \Lambda, x']) = \psi_{\zeta_1\Lambda\zeta'}^{\text{htpy}}(t, s, s', [y, \Lambda, x']).$$

Similarly, for any $\beta' \in \mathfrak{C}_\pi(Y', w')$, paths of connections $\zeta'_0 : \rho' \rightarrow \beta'$ and $\zeta'_1 : \beta' \rightarrow \alpha'$ on (Y', w') with composite $\zeta'_0 = \zeta'_1 = \zeta'$ and any element $[x, \Lambda, x', y']$ of

$$\check{M}_{\zeta}^+(Y, w; \alpha, \rho) \times_{\rho} \mathcal{O}_{\Lambda} \times_{\rho'} \check{M}_{\zeta'_0}^+(Y', w'; \rho', \beta') \times_{\beta'} \check{M}_{\zeta'_1}^+(Y', w'; \beta', \alpha') \subset X_{\zeta,\zeta'},$$

we have

$$\psi_{\zeta\Lambda\zeta'}^{\text{htpy}}(t, s, s', [x, \Lambda, x', y']) = \psi_{\zeta\Lambda\zeta'_0}^{\text{htpy}}(t, s, s', [x, \Lambda, x']),$$

for any $(s, s') \in A_T$ and $t \in [0, 1]$.

Now, we are ready to define the modified moduli spaces. Fix $\alpha \in \mathfrak{C}_\pi(Y, w)$, $\alpha' \in \mathfrak{C}_\pi(Y', w')$ and a path $z : \alpha \rightarrow \alpha'$ of connections along (W, c) with $i(z) \leq 6$. Recall that we have an obstructed abelian instanton $\Lambda : \rho \rightarrow \rho'$, which determines a subspace of obstructed solutions in the moduli space $M_z^+(W, c; \alpha, \alpha')$. If $\alpha' = \rho'$, then this subspace consists of obstructed solutions of type I given as $\check{M}_{\zeta}^+(Y, w; \alpha, \rho) \times_{\rho} \mathcal{O}_{\Lambda}$ with z being the composition of ζ and the path determined by Λ . A neighborhood $U_{\zeta\Lambda}$ of such obstructed solutions is given as

$$\Phi_{\zeta\Lambda} \left((\psi_{\zeta\Lambda}^{\text{an}})^{-1}(0) \cap \left(\check{M}_{\zeta}^+(Y, w; \alpha, \rho) \times (T, \infty] \right) \right) \subset M_z^+(W, c; \alpha, \rho').$$

We will verify below that the complement of $U_{\zeta\Lambda}$ in $M_z^+(W, c; \alpha, \rho')$ is an oriented compact stratified-smooth space; its top stratum is an open subset of $M_z^+(W, c; \alpha, \rho')$, and we give it the same orientation.

The zeros of $\psi_{\zeta\Lambda}^{\text{htpy}}$ over $[0, 1] \times \check{M}_{\zeta}^+(Y, w; \alpha, \rho)$ gives another stratified-smooth space $M_{\zeta\Lambda}^{\text{htpy}}$ associated to the path $z = \zeta * \lambda$. To define an orientation on $M_{\zeta\Lambda}^{\text{htpy}}$, first fix the product orientation on $[0, 1] \times \check{M}_{\zeta}^+(Y, w; \alpha, \rho)$. Since $\psi_{\zeta\Lambda}^{\text{htpy}}$ is a transverse section of a complex bundle, its zero set inherits an orientation from $[0, 1] \times \check{M}_{\zeta}^+(Y, w; \alpha, \rho)$. We multiply this orientation by ± 1 as appropriate so that $M_{\zeta\Lambda}^{\text{htpy}}$ has

$(\psi_{\zeta\Lambda}^T)^{-1}(0)$ as a boundary component with the opposite orientation as the corresponding boundary component of $M_z^+(W, c; \alpha, \rho') \setminus U_{\zeta\Lambda}$.

In the case $\alpha = \rho$, the moduli space $M_z^+(W, c; \alpha, \alpha')$ consists of obstructed solutions of type II given as $\mathcal{O}_\Lambda \times_{\rho'} \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha')$. We define the neighborhood $U_{\Lambda\zeta'}$ of these obstructed solutions as the image of

$$\Phi_{\Lambda\zeta'} \left((\psi_{\Lambda\zeta'}^{\text{an}})^{-1}(0) \cap \left(\check{M}_{\zeta'}^+(Y', w'; \rho', \alpha') \times [-\infty, T) \right) \right) \subset M_z^+(W, c; \rho, \alpha').$$

The compact stratified-smooth space given by the complement of $U_{\Lambda\zeta'}$ is oriented following the same convention as in the unobstructed case. We also define $M_{\Lambda\zeta'}^{\text{htpy}}$ to be the stratified-smooth space given as the zeros of $\psi_{\Lambda\zeta'}^{\text{htpy}}$ over $[0, 1] \times \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha')$. This space is again a compact stratified-smooth space, and its top stratum is oriented using the same convention as in the previous paragraph: one of its boundary components is canonically diffeomorphic to a boundary component in the complement of $U_{\Lambda\zeta'}$, and the orientations on these two boundary components should be opposite to one another.

For $\alpha \neq \rho$ and $\alpha' \neq \rho'$, the space $M_z^+(W, c; \alpha, \alpha')$ contains obstructed solutions of type III given as $X_{\zeta, \zeta'}$ where the composition $\zeta * \lambda * \zeta' = z$. The index assumption $i(z) \leq 6$ implies that there are unique choices of ζ and ζ' such that $X_{\zeta, \zeta'}$ is non-empty. A neighborhood $U_{\zeta\Lambda\zeta'}$ of these obstructed solutions can be obtained by forming the image of

$$\Phi_{\zeta\Lambda\zeta'} \left((\psi^{\text{an}})^{-1}(0) \cap \left(\left\{ (s, s') \in (0, \infty) \times [-\infty, 0) \mid \frac{1}{s'} - \frac{1}{s} \leq -\frac{1}{T} \right\} \times X_{\zeta, \zeta'} \right) \right) \subset M_z^+(W, c; \rho, \alpha').$$

The complement of this open neighborhood is a compact stratified-smooth space and can be oriented following the same conventions as in the unobstructed case. Let also $M_{\zeta\Lambda\zeta'}^{\text{htpy}}$ be the stratified-smooth space given as the zero set of $\psi_{\zeta\Lambda\zeta'}^{\text{htpy}}$ over $[0, 1] \times A_T \times X_{\zeta, \zeta'}$. If the domain of this section is given the product orientation, then we fix the orientation on $M_{\zeta\Lambda\zeta'}^{\text{htpy}}$ which differs from the standard orientation on the zero set of a complex line bundle by the appropriate factor ± 1 so that the appropriate boundary component has the opposite sign as the corresponding boundary component of the complement of $U_{\zeta\Lambda\zeta'}$.

The *modified moduli space* $N_z^+(W, c; \alpha, \alpha')$ is the $SO(3)$ -space given as

$$N_z^+(W, c; \alpha, \alpha') := (M_z^+(W, c; \alpha, \alpha') \setminus U_\Lambda) \sqcup M_\Lambda^{\text{htpy}}. \quad (5.13)$$

For the sake of uniformity, we have dropped the notation ζ and ζ' from our notations for the spaces $U_{\zeta\Lambda}$, $U_{\Lambda\zeta'}$ and $U_{\zeta\Lambda\zeta'}$, as well as from the spaces $M_{\zeta\Lambda}^{\text{htpy}}$, $M_{\Lambda\zeta'}^{\text{htpy}}$ and $M_{\zeta\Lambda\zeta'}^{\text{htpy}}$.

Proposition 5.14. *The modified moduli space $N_z^+(W, c; \alpha, \alpha')$ is a compact oriented stratified-smooth space. The boundary of $N_z^+(W, c; \alpha, \alpha')$ is the union of the following compact orientable smooth-stratified spaces.*

- (i) *If $\beta \in \mathfrak{C}_\pi(Y, w)$, and $\zeta_0 : \alpha \rightarrow \beta$ and $z_1 : \beta \rightarrow \alpha'$ are paths of connections along (Y, w) and (W, c) respectively with composite $\zeta_0 * z_1 = z$, then*

$$\check{M}_{\zeta_0}^+(Y, w; \alpha, \beta) \times_\beta N_{z_1}^+(W, c; \beta, \alpha')$$

is a boundary component of $N_z^+(W, c; \alpha, \alpha')$. The boundary orientation and the fiber product orientation on this boundary component are related to each other by $(-1)^{\dim(\alpha)}$.

(ii) If $\beta' \in \mathfrak{C}_{\pi'}(Y', w')$, and $z_0 : \alpha \rightarrow \beta'$ and $\zeta_1 : \beta' \rightarrow \alpha'$ are paths of connections along (W, c) and (Y', w') respectively with composite $z_0 * \zeta_1 = z$, then

$$N_{z_0}^+(W, c; \alpha, \beta') \times_{\beta'} \check{M}_{\zeta_1}^+(Y', w'; \beta', \alpha')$$

is a boundary component of $N_z^+(W, c; \alpha, \alpha')$. The boundary orientation and the fiber product orientation on this boundary component are related to each other by $(-1)^{\tilde{i}(z_0) + \dim(\alpha) + 1}$.

Depending on the type of α and α' , the remaining boundary components are given as in one of the following three cases:

(l) If $\alpha' = \rho'$ and $\zeta : \alpha \rightarrow \rho$ is a path of connections on (Y, w) with composite $\zeta * \lambda = z$, then the boundary components of $N_z^+(W, c; \alpha, \alpha')$ contain

(l-i) two copies of $(\psi_{\zeta\Lambda}^T)^{-1}(0)$ with opposite orientations;

(l-ii) a copy of $(\psi_{\zeta\Lambda}^{\text{priv}})^{-1}(0)$ whose orientation differs from the standard orientation of the zero section of a complex line bundle over an oriented space by a factor of $(-1)^{\dim(\alpha)}$.

(r) If $\alpha = \rho$ and $\zeta' : \rho' \rightarrow \alpha'$ is a path of connections on (Y', w') with composite $\lambda * \zeta' = z$, then the boundary components of $N_z^+(W, c; \alpha, \alpha')$ contain

(r-i) two copies of $(\psi_{\Lambda\zeta'}^T)^{-1}(0)$ with opposite orientations;

(r-ii) a copy of $(\psi_{\Lambda\zeta'}^{\text{priv}})^{-1}(0)$ whose orientation differs from the standard orientation of the zero section of a complex line bundle over an oriented space by a factor of $-1 = (-1)^{\dim(\rho) + 1}$.

(c) If $\alpha \neq \rho$ and $\alpha' \neq \rho'$, and $\zeta : \alpha \rightarrow \rho$, $\zeta' : \rho' \rightarrow \alpha'$ are paths of connections on (Y, w) and (Y', w') with composite $\zeta * \lambda * \zeta' = z$, then the boundary components of $N_z^+(W, c; \alpha, \alpha')$ contain

(c-i) two copies of $(\psi_{\zeta\Lambda\zeta'}^T)^{-1}(0)$ with opposite orientations;

(c-ii) a copy of $(\psi_{\zeta\Lambda\zeta'}^{\text{priv}})^{-1}(0)$, whose orientation differs from the standard orientation of the zero set of the section $s\psi_{\zeta\Lambda}^{\text{priv}} + (1-s)\psi_{\Lambda\zeta'}^{\text{priv}}$ over $\check{M}_{\zeta}^+(Y, w; \alpha, \rho) \times_{\rho} ([0, 1] \times \check{M}_{\zeta'}^+(Y', w'; \rho', \alpha'))$ by a factor of $(-1)^{\dim \alpha - 1}$.

Sketch of proof. What we will verify is that N_z^+ is a stratified-smooth space, and we will confirm the relevant boundary orientations; the actual enumeration of boundary strata is straightforward. The fact that $M_{\Lambda}^{\text{htpy}}$ is an oriented stratified-smooth space is the content of Proposition A.29.

As for the complement of U_{Λ} , one should define a function $L : M_z^+(W, c; \alpha, \alpha') \rightarrow (-\infty, 0)$ so that on a slight thickening of U_{Λ} it takes the following form in the three different cases. In the first case,

$L(x, \Lambda, t) = -1/t$. In the second case, $L(\Lambda, y, t') = 1/t'$. In the third case, $L(x, t, \Lambda, t', y) = 1/t' - 1/t$. In all cases, the set U_Λ is defined by the property $L > -1/T$. We extend L arbitrarily to the whole of M_z^+ , so long as $L^{-1}[-1/T, 0) = \overline{U_\Lambda}$ remains true. Then Proposition A.31 guarantees that the complement of U_Λ is an oriented stratified-smooth space with orientation induced as a codimension-0 subspace of M_z^+ .

What remains is to confirm the stated boundary orientations; because M^{htpy} is defined so that its positive boundary component is oriented by the same conventions as $(\psi^T)^{-1}(0)$, it suffices to compute the boundary orientation on the stratum $L^{-1}(-1/T)$ using Proposition A.31. Instead of computing the orientation on the zero locus using the convention that the last factor should map via dL in an orientation-preserving way to \mathbb{R} , the sign $(-1)^{\dim X-1}$ shows that we can compute this zero set using the convention that the *first* factor maps via dL in an orientation-preserving way to \mathbb{R} (with no additional sign). We ignore the discussion of taking the zero set of ψ^T , as taking the zero set of a complex bundle induces an orientation in a canonical way which commutes with all other orientation-related operations.

Now the conventions of Section 3.7 have $M_\zeta^+(Y, w; \alpha, \rho)$ oriented as

$$\alpha \times \mathbb{R} \times \check{M}_\zeta^{\text{Fiber}}(Y, w; \alpha, \rho).$$

Commuting the \mathbb{R} factor to the front adds a sign of $(-1)^{\dim \alpha}$. The map L is then given by $L(t, a, x) = -1/t$; the differential of this is $dL(\dot{t}, \dot{a}, \dot{x}) = t^{-2}dt$, so dL gives an orientation-preserving isomorphism from the first factor to the last. Therefore $L^{-1}(-1/T)$ is oriented as

$$(-1)^{\dim \alpha} \alpha \times \check{M}_\zeta^{\text{Fiber}} \cong (-1)^{\dim \alpha} \check{M}_\zeta(Y, w; \alpha, \rho).$$

This justifies the first sign. The second is similar: this space is oriented as $\rho' \times \mathbb{R} \times \check{M}_{\zeta'}^{\text{Fiber}}(Y', w'; \rho', \alpha')$. Moving the factor of \mathbb{R} to the front introduces a sign of $(-1)^{\dim \rho'} = 1$. The map is given by $L(t', a', A') = 1/t'$, so that $dL = -(t')^{-2}dt'$; this is orientation-reversing, which introduces an additional sign of -1 , giving the statement in the proposition.

Finally, the set $L^{-1}(-1/T)$ over $X_{\zeta, \zeta'}$ is oriented as the zero set of a map from the following space to \mathbb{R} :

$$\alpha \times \mathbb{R} \times \check{M}_\zeta^{\text{Fiber}}(Y, w; \alpha, \rho) \times \mathbb{R} \times \check{M}_{\zeta'}^{\text{Fiber}}(Y', w'; \rho', \alpha').$$

Commuting all factors of \mathbb{R} to the front induces a sign of $(-1)^{2 \dim \alpha + \tilde{i}(\zeta) - 1}$; passing to the boundary face $L^{-1}(-1/T)$ gives us

$$(-1)^{\tilde{i}(\zeta)} A_T \times \alpha \times \check{M}_\zeta^{\text{Fiber}}(Y, w; \alpha, \rho) \times \check{M}_{\zeta'}^{\text{Fiber}}(Y', w'; \rho', \alpha'),$$

where the extra sign appears because A_T is oriented as the inverse image of $L^{-1}(-1/T)$, but placing the dL factor *first* is the opposite orientation convention on A_T . Finally, commuting A_T back to the middle (and recalling that the oriented diffeomorphism from A_T to $[0, 1]$ sends our section to the one stated in the Proposition) gives us the sign of $(-1)^{\dim \alpha - 1}$, as stated. \square

Remark 5.15. Because the map $\psi_{\zeta \Lambda \zeta'}^{\text{priv}}$ may be identified with the section $t \cdot \psi_{\zeta \Lambda}^{\text{priv}} + (1 - t) \cdot \psi_{\Lambda \zeta'}^{\text{priv}}$ over $\check{M}_\zeta(Y, w; \alpha, \rho) \times [0, 1] \times \check{M}_{\zeta'}(Y', w'; \rho', \alpha')$, it can be described entirely in terms of these sections. We give an alternate description of this zero set in Lemma A.39 in the appendix, which we will use in Section 5.4 below when verifying certain boundary relations in the definition of bimodule. \diamond

5.3 The blown-up flow category

In the previous section, we defined modified moduli spaces $N_\zeta^+(W, c; \alpha, \alpha')$ attached to (W, c) . We would like to use the machinery of Section 4.2 to say that these define a bimodule (as in Definition 4.11) between flow categories, and hence a chain map between their flow complexes. This amounts to a boundary relation on these moduli spaces.

Inspecting the boundary relation of Proposition 5.14, we see that these moduli spaces do not define a bimodule $\mathcal{I}(Y, w, \pi) \rightarrow \mathcal{I}(Y', w', \pi')$ between the framed instanton flow categories. However, we still expect to be able to extract a chain map from these moduli spaces, using a *slightly different flow category* as the codomain. We start by describing this slightly different flow category, via a construction called the *blown-up flow category*.

Before doing so, we should give a brief introduction to a construction carried out in Appendix A.3; see Construction A.32 for more details. Suppose we are given a stratified-smooth space with a map to a smooth manifold $\phi : P \rightarrow M$, a complex line bundle $L \rightarrow M$, and a lift $\psi : P \rightarrow L$ of ϕ to a section of L transverse to the zero section. Then we may define its *real blowup*, a stratified-smooth space $B(\psi)$ with a natural map to the sphere bundle $S(L)$; there is a natural projection map $\pi : B(\psi) \rightarrow P$, and if considered as a geometric chain in M , the chain P is *collapse-equivalent* to its blowup $B(\psi)$ (a notion described in more detail in Appendix A.3; what matters here is that collapse-equivalence is preserved by fiber products, and that two collapse-equivalent stratified-smooth spaces define the same element in the geometric chain complex). This chain in $S(L)$ satisfies the boundary relation

$$\partial B(\psi) = B(\psi|_{\partial P}) - (-1)^{\dim P} Z(\psi) \times_M S(L),$$

where $Z(\psi)$ is the zero set of this section and the map from the second chain to $S(L)$ is projection to the second factor.

This construction allows us to lift chains in M to chains in the sphere bundle $S(L)$, which we use in the construction of the blown-up flow category below.

Construction 5.16. Let \mathcal{C} be an $SO(3)$ -equivariant, finitely graded flow category as in Definition 4.5, whose objects are all $SO(3)$ -orbits isomorphic to one of $SO(3)$, S^2 , or the singleton $*$. We write $\text{Ob}(\mathcal{C}) = \mathfrak{C}^* \sqcup \mathfrak{A} \sqcup \mathfrak{J}$ for the decomposition of the set of objects into these three types of orbit; \mathfrak{C}^* is the set of objects isomorphic to $SO(3)$, while \mathfrak{A} is the set of objects isomorphic to S^2 and \mathfrak{J} is the set of singleton objects.

Further, assume that each $\mathcal{C}(\alpha, \beta)$ is a free $SO(3)$ -manifold.

Choose an element $\rho \in \mathfrak{A}$. Fix an $SO(3)$ -equivariant complex line bundle $\mathcal{H}_\rho \rightarrow \rho$, where $SO(3)$ acts freely on any nonzero vector; one may take $\mathcal{H}_\rho = T\rho \cong SO(3) \times_{SO(2)} \mathbb{C}$.

Choose, for each $\beta \in \text{Ob}(\mathcal{C})$, a lift $\psi_{\rho\beta} : \mathcal{C}(\rho, \beta) \rightarrow \mathcal{H}_\rho$ and which is transverse to 0. We demand these are compatible in the sense that, on the fiber product

$$\mathcal{C}(\rho, \alpha) \times_\alpha \mathcal{C}(\alpha, \beta) \subset \partial \mathcal{C}(\rho, \beta),$$

the restriction of $\psi_{\rho\beta}$ coincides with the pullback of $\psi_{\rho\alpha}$ to the fiber product.¹⁰

¹⁰Verifying that such $\psi_{\rho\beta}$ exist is a straightforward induction; the key point is that every $\mathcal{C}(\rho, \beta)$ is $SO(3)$ -free, so there is no

These choices in hand, we define the blown-up flow category $\mathcal{B}_\rho(\mathcal{C})$ (or simply \mathcal{B} when the flow category \mathcal{C} and object ρ are clear from context) as follows; the construction depends on the family of sections $\psi_{\rho\beta}$, but we suppress it from notation.

First we will describe the object space. Each object $\alpha \in \text{Ob}(\mathcal{C}) \setminus \{\rho\}$ is left unchanged and gives an object of \mathcal{B} , with the same relative grading. The chosen orbit ρ gives rise to two objects of \mathcal{B} . The first is $S_\rho = S(\mathcal{H}_\rho)$, the unit sphere bundle of $\mathcal{H}_\rho \rightarrow \rho$. This space is isomorphic to $SO(3)$, and will be given relative grading one lower than ρ . The second is ρ itself, but it will have relative grading 2 less than ρ . To distinguish, we will write this as $\bar{\rho}$. That is,

$$\text{Ob}(\mathcal{B}) = (\text{Ob}(\mathcal{C}) \setminus \{\rho\}) \sqcup S_\rho \sqcup \bar{\rho},$$

where the relative grading is given by the original relative grading plus a shift factor:

$$\tilde{i}_{\mathcal{B}}(\alpha, \beta) = \tilde{i}_{\mathcal{C}}([\alpha], [\beta]) + s(\alpha) - s(\beta).$$

If $\alpha \in \text{Ob}(\mathcal{B})$, one should interpret

$$[\alpha] = \begin{cases} \rho & \alpha = S_\rho, \bar{\rho} \\ \alpha & \text{else} \end{cases}$$

The shift factor is $s(\alpha) = 0$ for $\alpha \neq \bar{\rho}, S_\rho$ but $s(\bar{\rho}) = -2$ and $s(S_\rho) = -1$.

We move on to defining the morphism spaces $\mathcal{B}(\alpha, \beta)$ in our blown-up flow category. Let us introduce some notation:

- Write $Z(\rho, \beta) = Z(\psi_{\rho\beta}) = \psi_{\rho\beta}^{-1}(0)$ for the stratified-smooth space given by the zero locus of $\psi_{\rho\beta}$; this space has dimension two less than $\mathcal{C}(\rho, \beta)$.
- Write $B(\rho, \beta)$ for the real blowup $B(\psi_{\rho\beta})$ discussed above; this has the same dimension as $\mathcal{C}(\rho, \beta)$. This is equipped with canonical maps $\psi/\|\psi\| : B(\rho, \beta) \rightarrow S_\rho$ and $e_+\pi : B(\rho, \beta) \rightarrow \beta$.

Then we define the morphism spaces $\mathcal{B}_\rho(\alpha, \beta)$ as in the following matrix:

$$\mathcal{B}(\alpha, \beta) = \left(\begin{array}{c|ccc} & \alpha \neq \rho & S_\rho & \bar{\rho} \\ \hline \beta \neq \rho & \mathcal{C}(\alpha, \beta) & -B(\rho, \beta) & Z(\rho, \beta) \\ S_\rho & \mathcal{C}(\alpha, \rho) \times_\rho S_\rho & 0 & 0 \\ \bar{\rho} & 0 & S_\rho & 0 \end{array} \right)$$

Here the columns correspond to the critical orbit you start at, and the rows correspond to the critical orbit you end at. \diamond

Remark 5.17. The fiber product map $\times_{S_\rho} S_\rho : C_*(S_\rho) \rightarrow C_*(\bar{\rho})$ arising from the term in entry (3, 2) of the above matrix is the induced map of $\pi : S_\rho \rightarrow \rho$; that is, if $\phi : P \rightarrow S_\rho$ is a geometric chain, then $\phi \times_{S_\rho} S_\rho = \pi\phi : P \rightarrow \bar{\rho}$ is the composite of ϕ with projection to the base. (Thus, one can find a copy of the algebraic mapping cone of $\pi_* : C_*(S_\rho) \rightarrow C_*(\bar{\rho})$ inside the Morse chain complex of this flow category.) \diamond

trouble arising from the demand for equivariant transversality.

In a moment we will verify that this does indeed give a flow category. First, we restate the properties of real blowups for the particular spaces $B(\rho, \beta) = B(\psi_{\rho\beta})$. First, $B(\rho, \beta)$ sits in a canonical commutative diagram

$$\begin{array}{ccc} B(\rho, \beta) & \xrightarrow{\psi/\|\psi\| \times e_+} & S_\rho \times \beta \\ \downarrow & & \downarrow \pi \times 1 \\ \mathcal{C}(\rho, \beta) & \xrightarrow{e_- \times e_+} & \rho \times \beta \end{array}$$

When considered as a geometric chain downstairs in $\rho \times \beta$ by following the projection $\pi \times 1$, Lemma A.37 asserts that this space is blowup-equivalent to $\mathcal{C}(\rho, \beta)$, so define the same elements in the geometric chain complex of $\rho \times \beta$, and induce the same chain-level fiber product maps. Lastly, by the same lemma we have

$$\partial B(\rho, \beta) = B(\psi_{\rho\beta} |_{\partial \mathcal{C}(\rho, \beta)}) - (-1)^{\dim \mathcal{C}(\rho, \beta)} Z(\rho, \beta) \times_\rho S(\mathcal{H}_\rho).$$

By commuting the odd-dimensional factor $S(\mathcal{H}_\rho)$ to the left, the last term simplifies to $-S_\rho \times_\rho Z(\rho, \beta)$. As for the first term: expanding out using the boundary formula for $\partial \mathcal{C}(\rho, \beta)$ from Definition 4.5, and using the fact that $\psi_{\rho\beta} = \psi_{\rho\alpha} \times 0$ over $\mathcal{C}(\rho, \alpha) \times_\alpha \mathcal{C}(\alpha, \beta)$, the whole boundary relation simplifies to

$$\partial B(\rho, \beta) = -S_\rho \times_\rho Z(\rho, \beta) + \sum_{\alpha \in \mathcal{C} \setminus \rho} (-1)^{\tilde{i}_{\mathcal{C}(\rho, \alpha)}} B(\rho, \alpha) \times_\alpha \mathcal{C}(\alpha, \beta).$$

(Recall here that $\dim \mathcal{C}(\rho, \alpha) = \tilde{i}_{\mathcal{C}(\rho, \alpha)} - 1$.) This in hand, we are prepared to establish that \mathcal{B} is a flow category.

The reader should have on-hand Definition 4.5 while reading the argument. Below is a diagram depicting the various moduli spaces in the definition of this flow category, as well as where they run from and to; when reading the argument below, keep in mind that the goal is to show that the boundary components of $\mathcal{B}(\alpha, \beta)$ are a sum over $(-1)^{\tilde{i}(\alpha, \gamma)} \mathcal{B}(\alpha, \gamma) \times_\gamma \mathcal{B}(\gamma, \beta)$, and this sum can be understood as running over all two-step paths from α to β in the diagram. The moduli space $\mathcal{B}(\alpha, \gamma)$ decorates the arrow from α to γ , while the moduli space $\mathcal{B}(\gamma, \beta)$ decorates the arrow from γ to β .

$$\begin{array}{ccc} & \mathcal{C}(\alpha, \beta) & \\ & \downarrow \cap & \\ & \alpha/\beta & \\ \mathcal{C}(\alpha, \rho) \times_\rho S_\rho & \nearrow & \nwarrow Z(\rho, \beta) \\ & -B(\rho, \beta) & \\ S_\rho & \xrightarrow{S_\rho} & \bar{\rho} \end{array}$$

The term α/β means that an arrow leaving this orbit correspond to moduli spaces domain α , while arrows running to this orbit correspond to moduli spaces with codomain β .

Theorem 5.18. *The objects, relative gradings, and morphism spaces constituting $\mathcal{B}_\rho(\mathcal{C})$, constructed as above, define a flow category as in Definition 4.5.*

Proof. Because all of these spaces are $SO(3)$ -spaces and all of the objects of \mathcal{B} are $SO(3)$ -orbits, it follows that all endpoint maps are still submersions. We leave it to the reader to confirm that

$$\dim \mathcal{B}(\alpha, \beta) = \dim \alpha + \tilde{i}_{\mathcal{B}}(\alpha, \beta) - 1$$

for all α, β . Now let us check the boundary relations in the definition of flow category, one-by-one. When proving the flow category relations below, we immediately replace every occurrence of $\mathcal{B}(\alpha, \beta)$ with its definition in terms of \mathcal{C} from the matrix in Construction 5.16.

- If $\alpha, \beta \neq \bar{\rho}, S_\rho$, then the flow category relations stipulate that

$$\begin{aligned} (-1)^{\dim \alpha} \partial \mathcal{C}(\alpha, \beta) &= \sum_{\gamma \neq \rho} (-1)^{\tilde{i}_{\mathcal{B}}(\alpha, \gamma)} \mathcal{C}(\alpha, \gamma) \times_\gamma \mathcal{C}(\gamma, \beta) \\ &\quad + (-1)^{\tilde{i}_{\mathcal{B}}(\alpha, S_\rho)} \mathcal{C}(\alpha, \rho) \times_\rho S_\rho \times_{S_\rho} (-B(\rho, \beta)). \end{aligned}$$

When $\gamma \neq \rho$, the expression $\tilde{i}_{\mathcal{B}}(\alpha, \gamma)$ is simply $\tilde{i}_{\mathcal{C}}(\alpha, \gamma)$. Because $\tilde{i}_{\mathcal{B}}(\alpha, S_\rho) = \tilde{i}_{\mathcal{C}}(\alpha, \rho) - 1$, the terms in the last sum simplify to

$$(-1)^{\tilde{i}_{\mathcal{C}}(\alpha, \rho)} \mathcal{C}(\alpha, \rho) \times_\rho B(\rho, \beta).$$

Because $B(\rho, \beta)$ is collapse-equivalent to $\mathcal{C}(\rho, \beta)$ as a chain in $\rho \times \beta$, this fiber product is equal to $\mathcal{C}(\alpha, \rho) \times_\rho \mathcal{C}(\rho, \beta)$ in the geometric chain complex. Thus the flow category relation for $\mathcal{B}(\alpha, \beta)$ asks that we have an equality of elements of $C_*(\alpha \times \beta)$

$$(-1)^{\dim \alpha} \partial \mathcal{C}(\alpha, \beta) = \sum_{\gamma \in \text{Ob}(\mathcal{C})} (-1)^{\tilde{i}_{\mathcal{C}}(\alpha, \gamma)} \mathcal{C}(\alpha, \gamma) \times_\gamma \mathcal{C}(\gamma, \beta).$$

This is precisely what it means when we assume that \mathcal{C} is a flow category to begin with.

- If $\alpha = S_\rho$ and $\beta \neq \bar{\rho}, S_\rho$, the desired boundary relation is

$$\begin{aligned} (-1)^{\dim S_\rho} (-\partial B(\rho, \beta)) &= (-1)^{\tilde{i}_{\mathcal{B}}(S_\rho, \bar{\rho})} S_\rho \times_\rho Z(\rho, \beta) \\ &\quad + \sum_{\gamma \neq \rho} (-1)^{\tilde{i}_{\mathcal{B}}(S_\rho, \gamma)} -B(\rho, \gamma) \times_\gamma \mathcal{C}(\gamma, \beta). \end{aligned}$$

After simplifying the signs and simplifying the last terms, we are trying to show that

$$\partial B(\rho, \beta) = -S_\rho \times_\rho Z(\rho, \beta) + \sum_{\gamma} (-1)^{\tilde{i}_{\mathcal{C}}(\rho, \gamma)} B(\rho, \gamma) \times_\gamma \mathcal{C}(\gamma, \beta).$$

This is precisely the boundary relation for the real blowup recalled above.

- The boundary relation for $\alpha = \bar{\rho}$ and $\beta \neq \bar{\rho}$, S_ρ follows immediately from the flow category relation for \mathcal{C} (take the zero locus of L on both sides of the flow category relation). The case $\alpha \neq S_\rho, \bar{\rho}$ and $\beta = S_\rho$ is similarly tautological.
- The boundary of $\mathcal{B}(\alpha, \bar{\rho})$ is tautologically empty. On the right-hand side of the flow category relation, the only factorization $\mathcal{B}(\alpha, \beta) \times_\beta \mathcal{B}(\beta, \bar{\rho})$ which is not tautologically empty is $\beta = S_\rho$. This means that the flow category relation in this case is

$$0 = \mathcal{C}(\alpha, \rho) \times_\rho S_\rho \times_{S_\rho} S_\rho = \pm \mathcal{C}(\alpha, \rho) \times_\rho S_\rho,$$

where the right-hand side is interpreted as a chain in $\alpha \times \bar{\rho}$. Because the map $\mathcal{C}(\alpha, \rho) \times_\rho S_\rho \rightarrow \bar{\rho}$ factors through the circle bundle projection

$$\mathcal{C}(\alpha, \rho) \times_\rho S_\rho \rightarrow \mathcal{C}(\alpha, \rho),$$

it follows that every stratum of this chain has small rank. So it is a degenerate chain, hence zero on the chain level, and our relation holds.

The final handful of cases follow by nearly identical arguments, occasionally recalling that $B(\rho, \beta)$ and $\mathcal{C}(\rho, \beta)$ define the same geometric chains in $\rho \times \beta$, as they are collapse-equivalent. \square

As mentioned above, there is a copy of the algebraic mapping cone

$$\text{Cone} = \text{Cone}(C_*(SO_3)[-1] \rightarrow C_*(S^2)[-2])$$

as a summand of the flow complex $CM_*(\mathcal{B}_\rho)$ (though not a subcomplex!). We think of the flow category \mathcal{B}_ρ as being the flow-categorical analogue of this mapping cone construction.

In fact, we will build a comparison map between these two flow categories based on that idea. There is a degree-zero chain map $S : C_*(S^2) \rightarrow \text{Cone}$, given by taking a chain $\phi : P \rightarrow S^2$ to the pullback $\phi \times_{S^2} SO(3)$ in $SO(3)$. If P is a point, it is sent to a circle; if ϕ is the fundamental class of S^2 it is sent to the fundamental class of $SO(3)$. This is a chain map because pushing $\phi \times_{S^2} SO(3)$ down to S^2 sends it to a degenerate chain (hence zero), so the mapping-cone portion of the differential is zero on $\phi \times_{S^2} SO(3)$.

It is easy to see by hand that the map above induces an isomorphism on homology $H_*(S^2) \rightarrow H_*(\text{Cone})$.

Next, we use this observation to define a bimodule $\mathcal{S}_\rho : \mathcal{C} \rightarrow \mathcal{B}_\rho$ which induces an isomorphism on the homology of their flow complexes.

Construction 5.19. Define the $SO(3)$ -equivariant bimodule \mathcal{S}_ρ as follows. Let $\mathcal{S}_\rho(\alpha, \alpha) = \alpha$ for all $\alpha \neq \rho$, while $\mathcal{S}_\rho(\rho, S_\rho) = S_\rho$ and $\mathcal{S}_\rho(\alpha, \beta) = 0$ in all other cases. That is,

$$\mathcal{S}_\rho = \left(\begin{array}{c|cc} & \alpha \neq \rho & \rho \\ \hline \beta \neq \rho & \delta_{\alpha=\beta}\alpha & 0 \\ S_\rho & 0 & S_\rho \\ \bar{\rho} & 0 & 0 \end{array} \right)$$

with

$$\tilde{i}_{\mathcal{S}}(\alpha, \beta) = \tilde{i}_{\mathcal{C}}(\alpha, \beta) - s(\beta),$$

where $s(\beta) = 0$ if $\beta \notin I$ while $s(S_\rho) = -1$ and $s(\bar{\rho}) = -2$.

Here $\delta_{\alpha=\beta}\alpha$ means this term is zero unless $\alpha = \beta$, in which case it is given by α . It is a straightforward definition-push to verify that \mathcal{S} is a bimodule, and the nontrivial parts of the argument are similar in spirit to parts of Theorem 5.18: the key point is that $B(\rho, \beta)$ and $\mathcal{C}(\rho, \beta)$ define the same geometric chain in $\rho \times \beta$. \diamond

Proposition 5.20. *For the $SO(3)$ -equivariant bimodule $S_\rho : \mathcal{C} \rightarrow \mathcal{B}_\rho(\mathcal{C})$ defined above, the induced chain map*

$$CM_*(\mathcal{S}) : CM_*(\mathcal{C}_\rho) \rightarrow CM_*(\mathcal{B})$$

is an isomorphism on homology.

Proof. The induced map on chains sends each $C_*(\alpha)$ identically onto $C_*(\alpha)$ whenever $\alpha \neq \rho$, and on the remaining factor is the map $S : C_*(\rho) \rightarrow \text{Cone}$ discussed before the construction above.

To see that $CM_*(\mathcal{S})$ induces an isomorphism on homology, we will use a spectral sequence argument using ‘periodic filtrations’; see Remark 4.10 for more details on these and their associated spectral sequences. We give $CM_*(\mathcal{C})$ and $CM_*(\mathcal{B})$ periodic filtrations, with filtration on the unrolled complex given by

$$F_k \widetilde{CM}_*(\mathcal{C}) = \bigoplus_{\tilde{i}_{\mathcal{C}}(\alpha) \leq k} C_*(\alpha) [\tilde{i}_{\mathcal{C}}(\alpha)] \subset \widetilde{CM}_*(\mathcal{C}),$$

$$F_k \widetilde{CM}_*(\mathcal{B}) = \bigoplus_{\substack{\alpha \in \text{Ob}(\mathcal{B}) \\ \tilde{i}_{\mathcal{C}}(\alpha) \leq k}} C_*(\alpha) [\tilde{i}_{\mathcal{B}}(\alpha)] \subset \widetilde{CM}_*(\mathcal{B}),$$

where in the second expression we filter by the \mathcal{C} -index of an object of \mathcal{B} ; that is, we filter by

$$\tilde{i}_{\mathcal{C}}(S_\rho) = \tilde{i}_{\mathcal{C}}(\bar{\rho}) = \tilde{i}_{\mathcal{C}}(\rho)$$

for $\rho \in I$ and by the usual $\tilde{i}_{\mathcal{C}}(\alpha)$ for $\alpha \neq \rho$.

The associated graded complex in the first case is $\bigoplus_{\alpha \in \text{Ob}(\mathcal{C})} C_*(\alpha) [\tilde{i}_{\mathcal{C}}(\alpha)]$ with differential the direct sum of the geometric chain complex differentials. The associated graded complex in the second case is

$$\bigoplus_{\alpha \neq \rho} C_*(\alpha) [\tilde{i}_{\mathcal{C}}(\alpha)] \oplus \text{Cone}[\tilde{i}_{\mathcal{C}}(\rho)].$$

Because the induced map on associated graded complexes is the identity on $C_*(\alpha)$ for $\alpha \neq \rho$, and on $C_*(\rho)$ is given by the map S described above, we see that S_ρ induces an isomorphism on the homology of the associated graded complexes.

By the discussion before the Construction, this map induces an isomorphism on homology between the associated graded complexes, and hence (because the spectral sequence associated to this periodic filtration is absolutely convergent) induces an isomorphism on HM_* . \square

Therefore, for the purposes of equivariant cohomology groups — which are preserved under $SO(3)$ -equivariant quasi-isomorphisms — we may freely pass between \mathcal{C} and \mathcal{B}_ρ . The corresponding flow complexes are not, however, equivariantly homotopy equivalent. This will be discussed further in Section 7 below, where it will be used to provide a wall-crossing formula of sorts for Floer’s non-equivariant instanton homology groups.

Any bimodule $\mathcal{B}_\rho(\mathcal{C}) \rightarrow \mathcal{C}'$ induces a bimodule $\mathcal{C} \rightarrow \mathcal{C}'$ by precomposition with \mathcal{S}_ρ . This leads us to ask: when does an $SO(3)$ -equivariant bimodule $\mathcal{W} : \mathcal{C} \rightarrow \mathcal{C}'$ enjoy a lift to an $SO(3)$ -equivariant bimodule $\mathcal{B}_\rho(\mathcal{W}) : \mathcal{B}_\rho\mathcal{C} \rightarrow \mathcal{C}'$?

This is not always the case (for instance, this is not possible for the identity bimodule). It turns out that this is true precisely when a certain index condition holds.

Construction 5.21. Suppose \mathcal{C} and \mathcal{C}' are flow categories satisfying the assumptions at the beginning of this section, and suppose we have chosen an orbit ρ of type S^2 in $\text{Ob}(\mathcal{C})$, an $SO(3)$ -equivariant complex line bundle $\mathcal{H}_\rho \rightarrow \rho$, as well as a compatible family of equivariant transverse sections $\psi_{\rho\alpha} : \mathcal{C}(\rho, \alpha) \rightarrow \mathcal{H}_\rho$ lifting the endpoint map to ρ .

Suppose $\mathcal{W} : \mathcal{C} \rightarrow \mathcal{C}'$ is an $SO(3)$ -equivariant bimodule, satisfying the following condition. For each $\alpha' \in \text{Ob}(\mathcal{C}')$ there is a stratified-smooth subspace $\mathcal{W}^{\text{red}}(\rho, \alpha')$ of those points with stabilizer conjugate to $SO(2)$.

We demand that at each $\Lambda \in \mathcal{W}^{\text{red}}(\rho, \alpha')$, the normal space to this locus is positive-dimension. It follows that the normal space has dimension at least two, as this normal space is an $SO(2)$ -space with weight one $SO(2)$ -action.

For each $\alpha' \in \text{Ob}(\mathcal{C}')$, choose an $SO(3)$ -equivariant transverse section

$$\psi_{\rho\alpha'}^W : \mathcal{W}(\rho, \alpha') \rightarrow \mathcal{H}_\rho.$$

On the boundary face $\mathcal{C}(\rho, \alpha) \times_\alpha \mathcal{W}(\alpha, \alpha')$ we demand this is given by the section $\psi_{\rho\alpha'}^W \times 0$ chosen above, while on the boundary face $\mathcal{W}(\rho, \beta') \times_{\beta'} \mathcal{C}'(\beta', \alpha')$ we demand this is given by $\psi_{\rho\beta'}^W \times 0$.

If there were no reducibles in these moduli spaces, these would be straightforward to construct by induction; at a reducible of type S^2 , the construction of such sections is possible by our assumption that each reducible has non-trivial normal space. This corresponds to the fact that there exists an $SO(2)$ -equivariant map $f : \mathbb{C}^n \rightarrow \mathbb{C}$ which is transverse to zero if and only if $n \geq 1$.

We write $B(\mathcal{W}; \rho, \alpha')$ and $Z(\mathcal{W}; \rho, \alpha')$ for the blowup and zero locus, respectively, of $\psi_{\rho\alpha'}^W$ over $\mathcal{W}(\rho, \alpha')$.

Given this, one may define a bimodule $\mathcal{B}_\rho(\mathcal{W}) : \mathcal{B}_\rho\mathcal{C} \rightarrow \mathcal{C}'$ by the formula

$$\mathcal{B}_\rho(\mathcal{W})(\alpha, \alpha') = \left(\begin{array}{c|ccc} & \alpha \neq S_\rho, \bar{\rho} & \alpha = S_\rho & \alpha = \bar{\rho} \\ \alpha' & \mathcal{W}(\alpha, \alpha') & B(\mathcal{W}; \rho, \alpha') & Z(\mathcal{W}; \rho, \alpha') \end{array} \right)$$


Here

$$\tilde{i}_{\mathcal{B}_\rho(\mathcal{W})}(\alpha, \alpha') = \tilde{i}_{\mathcal{W}}(\alpha, \alpha') - s(\alpha),$$

where the shift s is defined by the same formula as above. ◇

This defines a bimodule lifting \mathcal{W} .

Again, here is a diagram which may be helpful when verifying the bimodule relations. The left column depicts the flow category $\mathcal{B}_\rho\mathcal{C}$, while the right column depicts the flow category \mathcal{C}' , and the arrows running left to right depict the bimodule \mathcal{W} . As before, the bimodule relation for $\partial\mathcal{W}(\alpha, \alpha')$ amounts to saying that the boundary of any arrow running left to right is the signed sum over all two-step paths from α to α' .



Theorem 5.22. *The formulas above define a bimodule $\mathcal{B}_\rho(\mathcal{W}) : \mathcal{B}_\rho\mathcal{C} \rightarrow \mathcal{I}$. This bimodule satisfies $\mathcal{B}_\rho(\mathcal{W}) \circ \mathcal{S}_\rho = \mathcal{W}$.*

Proof. This follows by an analysis much like Theorem 5.18, though one needs to verify *fewer* relations. Once again, one repeatedly uses the fact that $S_\rho \times_{S_\rho} B(\mathcal{W}; \rho, \alpha')$ is collapse-equivalent to $\mathcal{W}(\rho, \alpha')$ so defines the same geometric chain, as well as the boundary formula for $B(\mathcal{W}; \rho, \alpha')$ given by Lemma A.37 and the boundary formula for $\mathcal{W}(\alpha, \alpha')$ given by the definition of bimodule.

The statement about the composite bimodule follows nearly immediately from the relevant definitions. The only term in this composite where this isn't tautological is

$$(\mathcal{B}_\rho(\mathcal{W}) \circ \mathcal{S}_\rho)(\rho, \alpha').$$

This reads

$$S_\rho \times_{S_\rho} B(\mathcal{W}; \rho, \alpha') = B(\mathcal{W}; \rho, \alpha') \text{ as a chain in } \rho \times \alpha',$$

which we know coincides with $\mathcal{W}(\rho, \alpha')$ at the level of geometric chains, as desired. \square

5.4 The bimodule of an obstructed cobordism

We will now show that the modified moduli spaces of Section 5.2 give rise to a bimodule *to* the blown-up flow category, opposite to the situation of Theorem 5.22.

Let $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ be a nearly unobstructed cobordism whose obstructed reducible we denote $\Lambda : \rho \rightarrow \rho'$. In Section 5.2, we constructed ‘modified moduli spaces’ $N_z^+(W, c; \alpha, \alpha')$. Our first goal in this section will be to use these to construct a bimodule

$$\mathcal{W}_- : \mathcal{I}(Y, w, \pi) \rightarrow \mathcal{B}_\rho \mathcal{I}(Y', w', \pi').$$

The minus sign indicates that $\tilde{i}(\Lambda) = -2$, so that index *decreases* across this obstructed reducible.

To make the notation more parseable, let us make some abbreviations. We write $M(\alpha, \beta) = \check{M}_\zeta^+(Y, w; \alpha, \beta)$, while $M'(\alpha', \beta')$ is the same for (Y', w') .

There are canonical equivariant diffeomorphisms $\rho \xleftarrow{e_-^\Lambda} \mathcal{O}_\Lambda \xrightarrow{e_+^\Lambda} \rho'$. We use these diffeomorphisms to pull back \mathcal{H}_Λ and S_Λ to bundles over ρ' , and we refer to the map $e_+^\rho : \rho \rightarrow \rho'$ for the map $e_+^\rho = e_+^\Lambda \circ (e_-^\Lambda)^{-1}$ obtained by inverting the first arrow in the above zig-zag. Every appearance of $M \times_\rho \mathcal{O}_\Lambda$ we replace with M in a canonical way, and similarly we replace $\mathcal{O}_\Lambda \times_{\rho'} M'$ with M' .

As in section 5.2 and the previous section, we choose a compatible family of transverse equivariant sections

$$\psi_{\alpha\rho} : M(\alpha, \rho) \rightarrow \mathcal{H}_{\rho'}, \quad \psi'_{\rho'\alpha'} : M'(\rho', \alpha') \rightarrow \mathcal{H}_{\rho'}.$$

We write $B(\alpha, \rho)$ and $B'(\rho', \beta')$ for the blowups along these sections, and $Z(\alpha, \rho)$ and $Z'(\rho', \alpha')$ for the zero loci of these sections; the family of sections described here plays the role of $\psi'_{\rho'\alpha'}$ in Section 5.3. Both $B(\alpha, \rho)$ and $B'(\rho', \beta')$ come with canonical maps to $S_{\rho'}$.

In Section 5.2, these were the sections $\psi_{\zeta\Lambda}^{\text{priv}}$, where ζ is the unique path with $-1 \leq i(\zeta) \leq 1$.

When defining the blown-up flow category, we will use $\psi'_{\rho'\alpha'} = -\psi_{\zeta\Lambda}^{\text{priv}}$. This minus sign makes no difference at the level of zero sets, but does change the map $B(\psi) \rightarrow S_{\rho'}$; it will be clear momentarily why this minus sign is important.

Finally, for the modified moduli spaces $N_\zeta^+(W, c; \alpha, \alpha')$, we abbreviate this expression to $N(\alpha, \alpha')$. In indices below, we suppress the symbols w in (Y, w) and c in (W, c) for ease of expression.

Construction 5.23. We define the bimodule $\mathcal{W}_- : \mathcal{I} \rightarrow \mathcal{B}_{\rho'} \mathcal{I}'$ by the following formula.

$$\mathcal{W}_-(\alpha, \alpha') = \left(\begin{array}{c|cc} & \alpha \neq \rho & \alpha = \rho \\ \hline \alpha' \neq \rho' & N(\alpha, \alpha') & N(\rho, \alpha') \\ \alpha' = S_{\rho'} & N(\alpha, \rho') \times_{\rho'} S_{\rho'} - (-1)^{\tilde{i}(Y; \alpha, \rho)} B(\alpha, \rho) & 0 \\ \alpha' = \bar{\rho}' & 0 & \rho \end{array} \right)$$

Here we have

$$\tilde{i}_{\mathcal{W}_-}(\alpha, \alpha') = \tilde{i}(W; \alpha, \alpha') + s(\alpha'),$$

where $s(\bar{\rho}') = 2$ and $s(S_{\rho'}) = 1$, while $s(\alpha') = 0$ otherwise. The bottom-right entry is the correspondence $\rho \xleftarrow{1_\rho} \rho \xrightarrow{e_+^\rho} \rho'$ induced by the endpoint maps of \mathcal{O}_Λ which was described above. \diamond

Remark 5.24. Entry $(2, 1)$ in this bimodule may be unexpected. What we will see in the calculation below is that this particular form is forced upon us by demanding that $\mathcal{W}_-(\alpha, \alpha') = N(\alpha, \alpha')$ whenever possible.

One could place this expression in its proper context by proving a more general statement: if $\mathcal{N}(\alpha, \alpha')$ is a collection of moduli spaces satisfying the boundary relations obtained in Proposition 5.14, then it assembles to give a bimodule $\mathcal{C} \rightarrow \mathcal{B}_\rho \mathcal{C}'$ by formulas along the lines of the above, and conversely every bimodule $\mathcal{C} \rightarrow \mathcal{B}_\rho \mathcal{C}'$ arises in this way. We will not bother formulating or proving such a statement. \diamond

Before writing out the proof, let's recall some of the important boundary relations here. Proposition 5.14 can be interpreted as giving the following boundary formula, at the level of geometric chains; the passage to geometric chains allows us to cancel out faces with opposite orientation, and equate things which are blowup-equivalent.

When $\alpha \neq \rho$ and $\alpha' \neq \rho'$, we have the boundary relation

$$\begin{aligned} (-1)^{\dim \alpha} \partial N(\alpha, \alpha') &= \sum_{\beta} M(\alpha, \beta) \times_{\beta} N(\beta, \alpha') \\ &\quad + \sum_{\beta'} (-1)^{\tilde{i}(W; \beta, \alpha') + 1} N(\alpha, \beta') \times_{\beta'} M'(\beta', \alpha') \\ &\quad + B(\alpha, \rho) \times_{S_\rho} B'(\rho', \alpha'). \end{aligned}$$


The final term deserves some comment. It appears in Proposition 5.14 as a negatively-oriented copy of the zero set of the family $t\psi_{\alpha\rho} - (1-t)\psi'_{\rho'\alpha'}$ on the pullback of \mathcal{H}_Λ over

$$M(\alpha, \rho) \times [0, 1] \times_{\rho'} M'(\rho', \alpha').$$

Applying the orientation-reversing isomorphism $t \mapsto 1 - t$ (which flips the negative orientation) on the unit interval and negating the section (which does not change the orientation on the zero set), this section is sent to $(1-t)\psi_{\alpha\rho} - t\psi'_{\rho'\alpha'}$. We show in Lemma A.39 in the appendix that the zero set of this section is identical as a geometric chain to the fiber product $B(\alpha, \rho) \times_{S_\rho} B'(\rho', \alpha')$, giving the statement written above. (This is where it is important to use $\psi'_{\rho'\alpha'} = -\psi_{\Lambda\zeta'}^{\text{priv}}$; with the opposite sign, Lemma A.39 would not be applicable.)

If $\alpha \neq \rho$ and $\alpha' = \rho$, the last term is instead $Z(\rho, \alpha)$. If $\alpha = \rho$ and $\alpha' \neq \rho'$, the final term is instead $-Z'(\rho', \alpha')$.

In the diagram below, the left column corresponds to the flow category \mathcal{C} , while the right column corresponds to the flow category $\mathcal{B}_\rho \mathcal{C}'$. The bimodule relations for $\mathcal{W}_-(\alpha, \alpha')$ assert that the boundary of the arrow running from (say) α to α' is the signed sum of the fiber product over all ways to go from α to α' in two steps, one along \mathcal{W}_- and one inside either the initial or ending flow category.



Theorem 5.25. For (W, c) a cobordism as in the beginning of this section, the moduli spaces \mathcal{W}_- defined above give a bimodule $\mathcal{W}_- : \mathcal{I} \rightarrow \mathcal{B}_{\rho'} \mathcal{I}'$.

Proof. The statement that \mathcal{W}_- gives a bimodule amounts to six boundary formulas. We will only verify two of these; the rest follow by nearly identical arguments, and are furthermore similar in spirit to the verifications in Theorem 5.18.

- First (filling in the values of \mathcal{W}_-), the bimodule relations for $\mathcal{W}_-(\alpha, \alpha')$ for $\alpha \neq \rho$ and $\alpha' \neq \bar{\rho}'$, $S_{\rho'}$ stipulate that

$$\begin{aligned}
(-1)^{\dim \alpha} \partial N(\alpha, \alpha') &= \sum_{\beta \in \mathfrak{C}(Y, w)} M(\alpha, \beta) \times_{\beta} N(\beta, \alpha') \\
&+ \sum_{\substack{\beta' \in \mathfrak{C}(Y', w') \\ \beta' \neq \rho'}} (-1)^{\tilde{i}_{\mathcal{W}_-}(\alpha, \beta') + 1} N(\alpha, \beta') \times_{\beta'} M'(\beta', \alpha') \\
&+ (-1)^{\tilde{i}_{\mathcal{W}_-}(\alpha, S_{\rho'}) + 1} N(\alpha, \rho') \times_{\rho'} S_{\rho'} \times_{S_{\rho'}} (-B'(\rho', \alpha')) \\
&+ (-1)^{\tilde{i}_{\mathcal{W}_-}(\alpha, S_{\rho'}) + 1} (-1)^{\tilde{i}(Y; \alpha, \rho) + 1} B(\alpha, \rho) \times_{S_{\rho'}} (-B'(\rho', \alpha')).
\end{aligned}$$

The third and last line arise from $\mathcal{W}_-(\alpha, S_{\rho'}) \times_{S_{\rho'}} \mathcal{B}_{\rho'} \mathcal{I}'(S_{\rho'}, \alpha')$. Our first order of business is to clean up these last few lines. First, $\tilde{i}_{\mathcal{W}_-}(\rho, \bar{\rho}') = 0$. Next, $\tilde{i}_{\mathcal{W}_-}(\alpha, S_{\rho'}) = \tilde{i}(W; \alpha, \rho') + 1$ while

$$\begin{aligned}
N(W; \alpha, \rho') \times_{\rho'} S_{\rho'} \times_{S_{\rho'}} (-B'(\rho', \alpha')) &= -N(W; \alpha, \rho') \times_{\rho'} B'(\rho', \alpha') \\
&= -N(W; \alpha, \rho') \times_{\rho'} M'(\rho', \alpha'),
\end{aligned}$$

because by Lemma A.37 the chains $B'(\rho', \alpha')$ and $M'(\rho', \alpha')$ are identical when considered as geometric chains in $\rho' \times \alpha'$. Lastly, the final line may be simplified with the observation that

$$\tilde{i}_{\mathcal{W}_-}(\alpha, S_{\rho'}) + 1 - \tilde{i}(Y; \alpha, \rho) - 1 = \tilde{i}(W; \alpha, \rho') + 1 - \tilde{i}(Y; \alpha, \rho) = \tilde{i}(W; \rho, \rho') = -2$$

so that the sign out front is -1 ; pulling the factor of -1 from $-B'(\rho', \alpha')$ out front as well, this reads

$$\begin{aligned} (-1)^{\dim \alpha} \partial N(\alpha, \alpha') &= \sum_{\beta \in \mathfrak{C}(Y, w)} M(\alpha, \beta) \times_{\beta} N(\beta, \alpha') \\ &+ \sum_{\beta' \in \mathfrak{C}(Y', w')} (-1)^{\tilde{i}(W; \alpha, \beta') + 1} N(\alpha, \beta') \times_{\beta'} M'(\beta', \alpha') \\ &+ B(\alpha, \rho) \times_{S_{\rho'}} B'(\rho', \alpha'). \end{aligned}$$

But this is precisely the relation written above, coming from Proposition 5.14(iii).

- Let's verify the bimodule relation for the other interesting case, $\mathcal{W}_-(\alpha, S_{\rho'})$ where $\alpha \neq \rho$. To briefly describe what happens: $N(\alpha, \rho') \times_{\rho'} S_{\rho'}$ carries an extra boundary term $Z(\alpha, \rho') \times_{\rho'} S_{\rho'}$; the boundary relation from the moduli space $B(\alpha, \rho)$ which appears in $\mathcal{W}_-(\alpha, S_{\rho'})$ will 'cancel it out'. To be more precise, the bimodule relations stipulate that

$$\begin{aligned} (-1)^{\dim \alpha} \partial \mathcal{W}_-(\alpha, S_{\rho'}) &= \sum_{\substack{\beta \in \mathfrak{C}(Y, w) \\ \alpha \neq \beta \neq \rho}} \mathcal{I}(\alpha, \beta) \times_{\beta} \mathcal{W}_-(\beta, S_{\rho'}) \\ &+ \sum_{\substack{\alpha' \in \mathfrak{C}(Y', w') \\ \alpha' \neq \rho'}} (-1)^{\tilde{i}_{\mathcal{W}_-}(\alpha, \alpha') + 1} \mathcal{W}_-(\alpha, \alpha') \times_{\alpha'} \mathcal{B}_{\rho'} \mathcal{I}'(\alpha', S_{\rho'}). \end{aligned}$$

(The last sum does not include a factorization through $\mathcal{B}_{\rho'} \mathcal{I}'(\bar{\rho}', S_{\rho'})$ because this moduli space is empty.) The term $\mathcal{I}(\alpha, \beta)$ is $M(\alpha, \beta)$; the term $\mathcal{B}_{\rho'} \mathcal{I}'(\alpha', S_{\rho'})$ is $M'(\alpha', \rho') \times_{\rho'} S_{\rho'}$. Filling in the values of \mathcal{W}_- , the right-hand-side simplifies to

$$\begin{aligned} &\sum_{\beta \neq \rho} M(\alpha, \beta) \times_{\beta} N(\beta, \rho') \times_{\rho'} S_{\rho'} - (-1)^{\tilde{i}(Y; \beta, \rho)} M(\alpha, \beta) \times_{\beta} B(\beta, \rho) \\ &+ \sum_{\alpha' \neq \rho'} (-1)^{\tilde{i}(W; \alpha, \alpha') + 1} N(\alpha, \alpha') \times_{\alpha'} M'(\alpha', \rho') \times_{\rho'} S_{\rho'}. \end{aligned}$$

The left-hand side of the bimodule relation reads

$$(-1)^{\dim \alpha} \partial \left(N(\alpha, \rho') \times_{\rho'} S_{\rho'} - (-1)^{\tilde{i}(Y; \alpha, \rho)} B(\alpha, \rho) \right).$$

Now use the boundary relation for $(-1)^{\dim \alpha} \partial N(\alpha, \rho')$ stated before this Proposition; most every term in that boundary relation is represented in the sum above, except for the final term corresponding to the boundary face $Z(\alpha, \rho)$.

Subtracting $(-1)^{\dim \alpha} \partial N(\alpha, \rho') \times_{\rho'} S_{\rho'}$ from both sides, what we're left with trying to show is that

$$(-1)^{\dim \alpha + \tilde{i}(Y; \alpha, \rho) - 1} \partial B(\alpha, \rho) = -Z(\alpha, \rho) \times_{\rho} S_{\rho} + \sum_{\beta} (-1)^{\tilde{i}(Y; \beta, \rho)} M(\alpha, \beta) \times_{\beta} B(\beta, \rho) \times_{\rho} S_{\rho}.$$

Multiplying through by $(-1)^{\dim \alpha + \tilde{i}(Y; \alpha, \rho)} - 1$, what we assert is that

$$\partial B(\alpha, \rho) = (-1)^{\dim \alpha + \tilde{i}(Y; \alpha, \rho)} Z(\alpha, \rho) \times_{\rho} S_{\rho} + \sum_{\beta} (-1)^{\dim \alpha + \tilde{i}(Y; \alpha, \beta)} M(\alpha, \beta) \times_{\beta} B(\beta, \rho).$$

Finally, this relation is given in Lemma A.37: the final term is $B(\psi|_{\partial M(\alpha, \rho)})$, while the first term is $(-1)^{\dim Z(\alpha, \rho) + 1} Z(\alpha, \rho) \times_{\rho} S_{\rho}$, as $\dim Z(\alpha, \rho) = \tilde{i}(Y; \alpha, \rho) + \dim \alpha - 1$.

Carrying out a similar discussion in all the other cases is similar but strictly simpler. \square

5.5 Invariance of instanton homology

Recall from Section 3.1 that associated to each weakly admissible pair (Y, w) and each regular perturbation π is a function $\sigma_{\pi} : \mathfrak{A}(Y, w) \rightarrow 2\mathbb{Z}$, called its *signature data*.

Recall also from Section 4.3 that associated to each (Y, w, π) with basepoint b is an $SO(3)$ -equivariant flow category $\mathcal{I}(Y, w, \pi)$, whose associated flow complex we call $\tilde{C}(Y, w, \pi)$, and that for each unobstructed cobordism $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ with regular perturbation π_W , homology orientation of W , and path between the basepoints, we associate a bimodule $\mathcal{I}(W, c, \pi_W, \gamma)$ and thus an $SO(3)$ -equivariant chain map

$$\tilde{C}(W, c, \pi_W, \gamma) : \tilde{C}(Y, w, \pi) \rightarrow \tilde{C}(Y', w', \pi'),$$

independent of π_W up to equivariant chain-homotopy. Furthermore, these compose functorially (up to equivariant chain homotopy).

Taking $(W, c) = \mathbb{R} \times (Y, w)$ (with trivial homology orientation and path $\gamma = \mathbb{R} \times b$) and applying the index computation Proposition 3.24, this cobordism is unobstructed if and only if we have the inequality on signature data functions $\sigma_{\pi} \leq \sigma_{\pi'}$. In that case, we obtain an $SO(3)$ -equivariant continuation map

$$\varphi_{\pi \rightarrow \pi'} : \tilde{C}(Y, w, \pi) \rightarrow \tilde{C}(Y, w, \pi').$$

Our goal for the remainder of this section is to prove the following.

Theorem 5.26. *Whenever π_0, π_1 are two regular perturbations on (Y, w) with $\sigma_{\pi_0} \leq \sigma_{\pi_1}$, the continuation map*

$$\varphi_{\pi_0 \rightarrow \pi_1} : \tilde{I}(Y, w, \pi_0) \rightarrow \tilde{I}(Y, w, \pi_1)$$

is an isomorphism.

Proof of Theorem 5.26. When $\sigma_{\pi_0} = \sigma_{\pi_1}$ this follows from the machinery of Section 4.3: there are continuation maps $\varphi_{\pi_0 \rightarrow \pi_1}$ and $\varphi_{\pi_1 \rightarrow \pi_0}$ which are equivariantly homotopy inverse, and hence each is an equivariant homotopy equivalence. The theorem above is primary interesting when $\sigma_{\pi_0} < \sigma_{\pi_1}$.

We will only need to prove this in a special case. We say that π_0 and π_1 have *adjacent signature data*, meaning that there exists a reducible ρ with $\sigma_{\pi_1} = \sigma_{\pi_0} + 4\delta_{\rho}$; that is, σ_{π_0} is equal to σ_{π_1} for all reducibles except ρ , where σ_{π_1} is larger by the smallest possible amount (4).

To see why this is sufficient, recall that Proposition 3.4 shows that any signature data function is realized as the signature data of some regular perturbation on (Y, w) . It follows that whenever $\sigma_{\pi_0} < \sigma_{\pi_1}$, we may find a sequence of intermediate perturbations $\pi_0, \pi_{1/n}, \dots, \pi_{(n-1)/n}, \pi_1$ so that

$$\sigma_{\pi_0} < \dots < \sigma_{\pi_1},$$

where each $(\pi_{i/n}, \pi_{(i+1)/n})$ have adjacent signature data. Because we can write the continuation map $\varphi_{\pi_0 \rightarrow \pi_1}$ as a composite of the individual continuation maps $\varphi_{\pi_{i/n} \rightarrow \pi_{(i+1)/n}}$, if the continuation map is an isomorphism for perturbations with adjacent signature data, it is an isomorphism whenever $\sigma_{\pi_0} \leq \sigma_{\pi_1}$.

So suppose that π_0 and π_1 have adjacent signature data. Write W_+ for the cylinder $\mathbb{R} \times (Y, c)$, considered as a cobordism $(Y, w, \pi_0) \rightarrow (Y, w, \pi_1)$ equipped with a regular perturbation π_+ interpolating between these, and W_- for the same cylinder considered as a cobordism in the opposite direction $(Y, w, \pi_1) \rightarrow (Y, w, \pi_0)$ equipped with an obstructed-regular perturbation π_- interpolating between these.

The cobordism W_- is a nearly unobstructed cobordism (its reducibles correspond to the constant trajectories on the cylinder, so their normal index is given by $\frac{1}{2}(\sigma_{\pi_0}(\rho) - \sigma_{\pi_1}(\rho))$, and the one obstructed reducible $\Lambda : \rho \rightarrow \rho$ has normal index -2), and thus by the machinery of the previous section we obtain a bimodule

$$\mathcal{W}_- : \mathcal{I}(Y, w, \pi_1) \rightarrow \mathcal{B}_\rho \mathcal{I}(Y, w, \pi_0).$$

We write φ_- for the induced map on flow complexes

$$\varphi_+ : \tilde{\mathcal{C}}(Y, w, \pi_1) \rightarrow S_\rho \tilde{\mathcal{C}}(Y, w, \pi_0).$$

On the other hand, the cobordism W_+ satisfies the hypotheses of Construction 5.21: the only reducible running $\rho \rightarrow \rho$ has normal index 2. We thus obtain a bimodule \mathcal{W}_+ running in the opposite direction, and we write φ_+ for the induced map on flow complexes.

Abbreviating $\mathcal{I}_0 = \mathcal{I}(Y, w, \pi_0)$ while $\mathcal{I}_1 = \mathcal{I}(Y, w, \pi_1)$, we obtain the following diagram of flow categories and bimodules:

$$\begin{array}{ccc} \mathcal{I}_0 & \xrightarrow{\mathcal{W}_{\pi_0 \rightarrow \pi_1}} & \mathcal{I}_1 \\ \downarrow S_\rho & \swarrow \mathcal{W}_- & \nearrow \mathcal{W}_+ \\ \mathcal{B}_\rho \mathcal{I}_0 & & \end{array}$$

Here $\mathcal{W}_{\pi_0 \rightarrow \pi_1}$ is the bimodule obtained from the usual (unmodified) moduli spaces of instantons on the unobstructed cobordism $\mathbb{R} \times Y : (Y, w, \pi_0) \rightarrow (Y, w, \pi_1)$; its induced map on the complex $\tilde{\mathcal{C}}$ is the continuation map we named $\varphi_{\pi_0 \rightarrow \pi_1}$ above. We established in Theorem 5.22 that we have an equality of composites

$$\mathcal{W}_+ \circ S_\rho = \mathcal{W}_{\pi_0 \rightarrow \pi_1},$$

so that on flow complexes we have $\varphi_+ \circ S_\rho = \varphi_{\pi_0 \rightarrow \pi_1}$. We established in Proposition 5.20 that S_ρ induces an isomorphism on homology, so we are reduced to showing that φ_+ is an isomorphism on homology. Indeed, we will prove that φ_+ is an equivariant chain homotopy equivalence.

To see this, we assert that there exist $SO(3)$ -equivariant bimodule homotopies

$$\mathcal{H}_1 : \mathbf{1}_{\mathcal{I}_1} \implies \mathcal{W}_+ \circ \mathcal{W}_-, \quad \mathcal{H}_0 : \mathbf{1}_{\mathcal{B}_\rho \mathcal{I}_0} \implies \mathcal{W}_- \circ \mathcal{W}_+.$$

It follows that the induced maps φ_+ and φ_- are equivariantly homotopy inverse, as desired.

The construction of these bimodule homotopies is carried out in Section 5.5.2 below, based on the analytic setup of Section 5.5.1. The fact that they are indeed bimodule homotopies are verified in Propositions 5.40 and 5.38 in Section 5.5.2. \square

Now fix a PID R . Associated to the $C_*(SO(3); R)$ -module \tilde{C} (together with its periodic filtration by index \tilde{i}) are three equivariant homology groups H^+, H^-, H^∞ , which fit into an exact triangle; each of these are modules over $H^{-*}(BSO(3); R)$. A (periodic filtration-preserving) chain map $\tilde{C} \rightarrow \tilde{C}'$ induces a map on all equivariant homology groups, which commutes with the structure maps in the exact triangle. If the equivariant map $\tilde{C} \rightarrow \tilde{C}'$ is a quasi-isomorphism, it induces an isomorphism on all equivariant homology groups.

These constructions for periodically-filtered and periodically-graded dg-modules were first described in [Mil19, Appendix A.8]; essentially, one unwraps \tilde{C} into a \mathbb{Z} -graded complex \bar{C} , uses this to filter certain complexes $C^+(\bar{C}), C^-(\bar{C}), C^\infty(\bar{C})$ arising from \mathbb{Z} -graded equivariant homological algebra, and performs a completion with this filtration to extract well-behaved equivariant homology groups (well-behaved meaning they satisfy the properties above).

However, the filtration used there to define H^∞ fails to be exhaustive, and so the completion procedure used in [Mil19, Appendix A.8] does not produce a well-behaved homology group H^∞ (in particular, one fitting into an exact triangle and preserved by equivariant quasi-isomorphisms). One should instead use the completion procedure of [Hel22, Section 5.4], called *full completion*; in doing so, one recovers all the properties claimed in [Mil19, Theorem A.29], and this is explicitly stated as [Hel22, Proposition 5.40].

We write $I^\bullet(Y, w, \pi; R)$ to mean any the various equivariant instanton homology groups $\tilde{I}, I^+, I^-, I^\infty$ associated to the complex $\tilde{C}(Y, w, \pi; R)$.

Corollary 5.27 (Instanton homology is well-defined). *The equivariant instanton homology groups $I^\bullet(Y, w, \pi; R)$ are independent of regular perturbation π , as well as their module structures and the exact triangle relating them.*

Proof. Given any two regular perturbations π, π' with $\sigma_\pi \leq \sigma_{\pi'}$, we showed above that the natural map $\varphi_{\pi \rightarrow \pi'} : \tilde{C}(Y, w, \pi; R) \rightarrow \tilde{C}(Y, w, \pi'; R)$ induces an isomorphism on homology. By the discussion before this corollary, $\varphi_{\pi \rightarrow \pi'}$ also induces an isomorphism on all groups $I^\bullet(Y, w, \pi; R)$, which is a module map that commutes with the structure maps in the exact triangle.

Because any signature-data function is realized as the signature data of some regular perturbation, and the maximum of two signature data functions is again a signature data function, given any two

regular perturbations π_0, π_1 , there is a regular perturbation π_2 with $\sigma_{\pi_2} = \max(\sigma_{\pi_0}, \sigma_{\pi_1})$. Because the continuation maps $\varphi_{02} = \varphi_{\pi_0 \rightarrow \pi_2}$ and $\varphi_{12} = \varphi_{\pi_1 \rightarrow \pi_2}$ induce isomorphisms on all equivariant homology groups, this provides an isomorphism between

$$\varphi_{12}^{-1} \varphi_{02} : I^\bullet(Y, w, \pi_0) \rightarrow I^\bullet(Y, w, \pi_1).$$

This isomorphism is canonical (independent of the choice of π_2). If one picks a π_2' with $\sigma_{\pi_2} \leq \sigma_{\pi_2'}$, then by functoriality of induced maps of cobordisms, we have

$$\varphi_{22'} \varphi_{02} = \varphi_{02'}, \quad \varphi_{22'} \varphi_{12} = \varphi_{12'},$$

so that

$$\varphi_{12'}^{-1} \varphi_{02'} = \varphi_{12}^{-1} \varphi_{22'}^{-1} \varphi_{22'} \varphi_{02} = \varphi_{12}^{-1} \varphi_{02}.$$

For any two π_2 and π_2' exceeding π_0 and π_1 , there is another π_2'' exceeding all of these. Because $\varphi_{12}^{-1} \varphi_{02} = \varphi_{12''}^{-1} \varphi_{02''} = \varphi_{12'}^{-1} \varphi_{02'}$, we see that the isomorphism between $I^\bullet(Y, w, \pi_0)$ and $I^\bullet(Y, w, \pi_1)$ is independent of any choices in its construction. \square

We use the previous corollary to give a definition of the instanton homology groups which is transparently independent of π .

Definition 5.28. Let (Y, w) be a weakly admissible pair. We set

$$I^\bullet(Y, w; R) = \bigoplus_{\pi \text{ regular}} I^\bullet(Y, w, \pi; R) \Big/ \varphi_{\pi \rightarrow \pi'}.$$

Here we quotient by the additive relation generated by setting

$$(\varphi_{\pi \rightarrow \pi'}(x), \pi') \sim (x, \pi)$$

for any $x \in I^\bullet(Y, w, \pi; R)$ and any π' with $\sigma_\pi \leq \sigma_{\pi'}$. \diamond

Notice the continuation map $\varphi_{\pi \rightarrow \pi}$ is the identity, all continuation maps are isomorphisms, we have $\varphi_{\pi' \rightarrow \pi''} \varphi_{\pi \rightarrow \pi'} = \varphi_{\pi \rightarrow \pi''}$, and for any π, π' there is a perturbation π'' so that $\sigma_{\pi''} \leq \sigma_\pi, \sigma_{\pi'}$. It is a straightforward exercise in group theory to see that the natural map $I^\bullet(Y, w, \pi) \rightarrow I^\bullet(Y, w)$ is an isomorphism for all regular perturbations π . This is the essential content of the fact proven in the Corollary above: not only are all $I^\bullet(Y, w, \pi)$ isomorphic, they are *canonically* so.

5.5.1 Modified moduli spaces for homotopies

For the weakly admissible pair (Y, w) , fix two regular perturbations π_0 and π_1 with adjacent signature data. A typical element of $\mathfrak{C}_{\pi_0}(Y, w)$ is denoted by α and a typical element of $\mathfrak{C}_{\pi_1}(Y, w)$ is denoted by α' . Any reducible element σ of $\mathfrak{C}_{\pi_0}(Y, w)$ corresponds to a unique reducible element σ' of $\mathfrak{C}_{\pi_1}(Y, w)$. There is a unique such reducible $\rho \in \mathfrak{C}_{\pi_0}(Y, w)$ such that

$$\sigma_{\pi_1}(\rho') - \sigma_{\pi_0}(\rho) = 4,$$

and otherwise σ_{π_0} and σ_{π_1} are equal.

Let (W_+, c_+, π_+) be the product cobordism from (Y, w) to (Y, w) together with a perturbation π_+ from π_0 to π_1 ; the subscript $+$ is meant to indicate that the signature data increases across the cylinder. Using our assumption on σ_{π_0} and σ_{π_1} , we may assume that π_+ is regular and for any reducible $\sigma \in \mathfrak{C}_{\pi_0}(Y, w)$, there is a unique (unobstructed) reducible instanton associated to (W_+, c_+, π_+) which is asymptotic to σ and σ' on the two ends. In particular, there is a unique reducible instanton Λ_+ on W_+ from ρ to ρ' with $\tilde{i}(\Lambda_+) = +2$. The remaining reducible instantons on (W_+, c_+) have $\tilde{i}(\Lambda) = 0$ and are isolated.

Similarly, let (W_-, c_-, π_-) be the product cobordism from (Y, w) to (Y, w) together with a perturbation π_- from π_1 to π_0 , the subscript indicating that signature data decreases along the cylinder. In this case, we may only assume that π_- is obstructed-regular as in Theorem 3.31(ii). To be more precise, we may arrange for π_- such that for any reducible $\sigma \in \mathfrak{C}_{\pi_0}(Y, w)$, there is a unique reducible instanton associated to (W_+, c_+, π_+) , which is asymptotic to σ and σ' on the two ends. This reducible instanton is unobstructed, with $\tilde{i}(\Lambda) = 0$, unless $\sigma = \rho$. We will write Λ_- for the obstructed reducible instanton that also satisfies $\tilde{i}(\Lambda_-) = -2$, and $\lambda_- : \rho' \rightarrow \rho$ for the corresponding path of connections.

Composition of (W_+, c_+, π_+) and (W_-, c_-, π_-) produces the product cobordism (W_-^+, c_-^+) together with a family of perturbations $\pi_-^+(t)$, where t denotes the length of the neck, taking values in $[T_0, \infty)$. In particular, we may regard (W_-^+, c_-^+, π_-^+) as a cobordism from (Y, w, π_0) to (Y, w, π_0) . We extend the parametrizing set to $I = [T_0, \infty]$, by including the broken perturbation given by π_+ and π_- . For any $\alpha, \beta \in \mathfrak{C}_{\pi_0}(Y, w)$ and a path z from α to β with $i(z) \leq 5$, consider the moduli space

$$M_z^+(W_-^+, I; \alpha, \beta) := \bigsqcup_{t \in I} M_z^+(W_-^+, \pi_-^+(t); \alpha, \beta). \quad (5.29)$$

Here the disjoint union is meant in the set-theoretic sense, not the topological sense; topologically, if \mathbf{A}_n is a sequence of $\pi_-^+(t_n)$ -instantons, and $t_n \rightarrow t$ while \mathbf{A}_n chain-converges to \mathbf{A} , then this space is topologized so that $(t_n, \mathbf{A}_n) \rightarrow (t, \mathbf{A})$. This is completely analogous to the discussion in Section 3.6, and the definition of $\pi_-^+(\infty)$ -instantons is identical to the definition given there.

The following lemma can be established by the same arguments which give Theorem 3.31.

Lemma 5.30. *The family $\pi_-^+(t)$ can be chosen such that any element in the moduli space*

$$M_z^+(W_-^+, I; \alpha, \beta) \setminus M_z^+(W_-^+, \pi_-^+(\infty); \alpha, \beta)$$

is regular.

Similar in spirit to Section 5.1, we will divide non-regular solutions of $M_z^+(W_-^+, I; \alpha, \beta)$ into two types, all of which belong to the stratum $M_z^+(W_-^+, \pi_-^+(\infty); \alpha, \beta)$.

For any α and $\beta = \rho$, we say that an element of the subspace

$$M_{z_+}^+(W_+, \pi_+; \alpha, \rho') \times_{\rho'} \mathcal{O}_- \quad (5.31)$$

of $M_z^+(W_-^+, I; \alpha, \beta)$ is an obstructed solution of type I, where \mathcal{O}_- denotes the orbit of Λ_- in the framed moduli space. Here z_+ is the unique path with $z = z_+ * \lambda_-$.

For any α and $\beta \neq \rho$, we say that an obstructed solution of type II is an element of the subspace

$$X_{z_+\Lambda_-\zeta} := M_{z_+}^+(W_+, \pi_+; \rho, \rho') \times_{\rho'} \mathcal{O}_- \times_{\rho} \check{M}_{\zeta}^+(\pi_0; \rho, \beta). \quad (5.32)$$

of $M_z^+(W_-^+, I; \alpha, \beta)$. Here we have $z = z_+ * \lambda_- * \zeta$.

We have analogues of Proposition 5.7 for the obstructed solutions of any of the above four types that describe the behavior of the moduli space $M_z^+(W_-^+, I; \alpha, \beta)$ in a neighborhood of obstructed solutions. We start by analyzing the case of reducible solutions of type I.

First note that the complex line bundle $\mathcal{H}_{\Lambda_-}^+$ over \mathcal{O}_- pulls back to a line bundle \mathcal{H}_{z_+} over $X_{z_+\Lambda_-}$. There is an $SO(3)$ -equivariant section $\psi_{z_+\Lambda_-}^{\text{an}}$ of the pullback of $\mathcal{H}_{\alpha, \rho'}$ to the space

$$(0, \infty] \times M_{z_+}^+(W_+, \pi_+; \alpha, \rho') \times_{\rho'} \mathcal{O}_-.$$

This section vanishes over

$$\{\infty\} \times M_{z_+}^+(W_+, \pi_+; \alpha, \rho') \times_{\rho'} \mathcal{O}_- \quad (5.33)$$

and $\psi_{z_+\Lambda_-}^{\text{an}}$ is transverse to the zero section in the complement of this subspace.

As in part (ii) of Proposition 5.7, there is also an injective open continuous map

$$\Phi_{z_+\Lambda_-} : \left(\psi_{z_+\Lambda_-}^{\text{an}} \right)^{-1}(0) \rightarrow M_z(W_-^+, I; \alpha, \rho)$$

whose image is an open neighborhood of the subspace of obstructed solutions of type I, where $z = z_+ * \lambda_-$.

For any $[x, y] \in M_{z_+}^+(W_+, \pi_+; \alpha, \alpha') \times_{\alpha'} \check{M}_{\zeta'}^+(\pi_1; \alpha', \rho')$ with $z_+ = z'_+ * \zeta'$, we have

$$\psi_{z_+\Lambda_-}^{\text{an}}(t, [x, y]) = \psi_{\zeta'\Lambda_-}^{\text{an}}(t, y)$$

where $\psi_{\zeta'\Lambda_-}^{\text{an}}$ is the map to $\mathcal{H}_{\Lambda_-}^+$ defined in Section 5.1. Similarly for any $[x, y] \in \check{M}_{\zeta}^+(\pi_0; \alpha, \gamma) \times_{\gamma} M_{z_+}^+(W_+, \pi_+; \gamma, \rho')$ with $z_+ = \zeta * z'_+$, we have

$$\psi_{z_+\Lambda_-}^{\text{an}}(t, [x, y]) = \psi_{z'_+\Lambda_-}^{\text{an}}(t, y).$$

Just as in Section 5.1, there are further compatibility maps on the open embeddings Φ .

For the obstructed solutions of type II, first we define the complex line bundle \mathcal{H}_{z_+} over $X_{z_+\Lambda_-\zeta}$ as the pullback of $\mathcal{H}_{\Lambda_-}^+$ as usual. Then we have the obstruction section $\psi_{z_+\Lambda_-\zeta}^{\text{an}}$ defined over

$$(0, \infty] \times M_{z_+}^+(W_+, \pi_+; \alpha, \rho') \times_{\rho'} \mathcal{O}_- \times_{\rho} [-\infty, 0) \times \check{M}_{\zeta}^+(\pi_0; \rho, \beta).$$

This section vanishes on the strata lying over $\{(\infty, -\infty)\} \in (0, \infty] \times [-\infty, 0)$, and is transverse to the zero section in the complement of these strata. Again, there is a map $\Phi_{z_+\Lambda_-\zeta}$ from the zero set of ψ^{an} into an open neighborhood of the obstructed solutions of type II.

The restriction of this obstruction section to the remaining strata is determined by the obstruction sections defined on smaller moduli spaces. For any $[x, \Lambda_-, y] \in X_{z_+\Lambda_-\zeta}$, we have

$$\psi_{z_+\Lambda_-\zeta}^{\text{an}}(t, [x, \Lambda_-], -\infty, y) = \psi_{z_+\Lambda_-}^{\text{an}}(t, [x, \Lambda_-]), \quad \psi_{z_+\Lambda_-\zeta}^{\text{an}}(\infty, [x, \Lambda_-], t', y) = \psi_{\Lambda_-\zeta}^{\text{an}}(\Lambda_-, t', y).$$

For any $[x, y] \in M_{z_+}'(W_+, \pi_+; \alpha, \alpha') \times_{\alpha'} \check{M}_{\zeta'}^+(\pi_1; \alpha', \rho')$ with $z_+ = z'_+ * \zeta'$ and $u \in \check{M}_{\zeta}(\pi_0; \rho, \beta)$, we have

$$\psi_{z_+ \Lambda_- \zeta}^{\text{an}}(t, [x, y, \Lambda_-], t', u) = \psi_{\zeta' \Lambda_- \zeta}^{\text{an}}(t, [y, \Lambda_-], t', u),$$

and for any $[x, y] \in \check{M}_{\zeta}^+(\pi_0; \alpha, \gamma) \times_{\gamma} M_{z_+}'(W_+, \pi_+; \gamma, \rho')$ with $z_+ = \zeta * z'_+$ and $u \in \check{M}_{\zeta}(\pi_0; \rho, \beta)$, we have

$$\psi_{z_+ \Lambda_- \zeta}^{\text{an}}(t, [x, y, \Lambda_-], t', u) = \psi_{z'_+ \Lambda_- \zeta}^{\text{an}}(t, [y, \Lambda_-], t', u).$$

Finally if $x \in M_{z_+}^+(W_+, \pi_+; \alpha, \rho')$ and $[y, u] \in \check{M}_{\zeta_0}(\pi_0; \rho, \gamma) \times_{\gamma} M_{\zeta_1}(\pi_0; \gamma, \beta)$ with $\zeta = \zeta_0 * \zeta_1$, then

$$\psi_{z_+ \Lambda_- \zeta}^{\text{an}}(t, [x, \Lambda_-], t', [y, u]) = \psi_{z_+ \Lambda_- \zeta_0}^{\text{an}}(t, [x, \Lambda_-], t', y).$$

For all of these obstructed solutions, there is a map Φ with appropriate decoration that identifies the zero locus of the obstruction with a neighborhood of obstructed solutions in the moduli space $M_z^+(W_-^+, I; \alpha, \beta)$.

As in the case of the modified moduli spaces, we first remove a neighborhood of obstructed solutions in the moduli spaces $M_z^+(W_-^+, I; \alpha, \beta)$ and then add a new part, which is defined in terms of a family of privileged sections of the obstruction bundles and homotopies from the analytical obstruction sections to these sections.

To initiate this construction, we fix T , and let $\psi_{z_+ \Lambda_-}^T$ denote the restriction of $\psi_{z_+ \Lambda_-}^{\text{an}}$ to the space $\{T\} \times M_{z_+}^+(W_+, \pi_+; \alpha, \rho') \times_{\rho'} \mathcal{O}_-$. Similarly, let $\psi_{z_+ \Lambda_- \zeta}^T$ be the restriction of $\psi_{z_+ \Lambda_- \zeta}^{\text{an}}$ to the subspace of $X_{z_+ \Lambda_- \zeta}$ lying over $A_T \subset (0, \infty] \times [-\infty, 0)$; recall that this subspace A_T was defined in Section 5.2.

We then defined privileged sections $\psi_{z_+ \Lambda_-}^{\text{priv}}$ and $\psi_{z_+ \Lambda_- \zeta}^{\text{priv}}$ over these subspaces, which satisfy the same compatibility conditions that the analytic gluing sections do. We demand that the latter takes the form

$$\psi_{z_+ \Lambda_- \zeta}^{\text{priv}}(t, [x, y], t') = \frac{T}{t} \cdot \psi_{z_+ \Lambda_-}^{\text{priv}}(x) - \frac{T}{t'} \cdot \psi_{\Lambda_- \zeta}^{\text{priv}}(y);$$

we assume that the two previously-chosen sections ψ^{priv} (one chosen here, the other chosen in Section 5.2) are chosen so that this is a transverse section.

We then choose homotopies of sections $\psi_{z_+ \Lambda_-}^{\text{htpy}}$ and $\psi_{z_+ \Lambda_- \zeta}^{\text{htpy}}$ once again satisfying the same compatibility properties, and interpolating between the analytic gluing sections ψ^T and the privileged sections ψ^{priv} .

Finally, the modified moduli spaces $N_z^+(W_-^+, I; \alpha, \beta)$ are defined by exactly the same procedure given in Section 5.2, with the same orientation conventions: they are oriented as zero sets of sections of complex line bundles, where the domains of these sections are oriented as products given the product structures written above. In the case of $\psi_{z_+ \Lambda_- \zeta}^{\text{priv}}$, we orient the zero set as the zero set of

$$s \psi_{z_+ \Lambda_-}^{\text{priv}}(x) + (1 - s) \cdot \psi_{\Lambda_- \zeta}^{\text{priv}}(y)$$

over the space

$$M_{z_+}(W_+, \pi_+; \alpha, \rho') \times_{\rho'} \mathcal{O}_- \times_{\rho} \left([0, 1] \times \check{M}_{\zeta}(Y; \rho, \beta) \right).$$

That these are stratified-smooth spaces follows by the same argument as in Section 5.2. Below, we record the boundary of these stratified-smooth spaces, together with their orientations; unlike Section

5.2, we use the language of geometric chains to avoid talking about faces which appear with opposite orientations.

Suppose $z = z_+ * \lambda_-$, where $z_+ : \alpha \rightarrow \rho'$ is a path of connections on W_+ .

Proposition 5.34. *The boundary of the stratified smooth space $N_z^+(W_-^+, I; \alpha, \beta)$, considered as a geometric chain in $\alpha \times \beta$, is given as follows.*

(i) *If $\beta = \rho$, then $\partial N_z^+(W_-^+, I; \alpha, \rho)$ is given by*

$$\begin{aligned} (-1)^{\dim \alpha} \partial N_z^+(W_-^+, I; \alpha, \rho) = & \sum_{z=z_+ * z_-} M_{z_+}^+(W_+, \pi_+; \alpha, \alpha') \times_{\alpha'} N_{z_-}^+(W_-, \pi_-; \alpha', \rho) \\ & - N_z^+(W_-^+, \pi_-^+(T_0); \alpha, \rho) \\ & + \sum_{z=\zeta * z'} (-1)^{\tilde{i}(\zeta)} \check{M}_{\zeta}^+(Y, \pi_0; \alpha, \gamma) \times_{\gamma} N_{z'}^+(W_-^+, I; \gamma, \rho) \\ & + \sum_{z=z' * \zeta'} (-1)^{\tilde{i}(z')} N_{z'}^+(W_-^+, I; \alpha, \gamma) \times_{\gamma} \check{M}_{\zeta'}^+(Y, \pi_0; \gamma, \rho) \\ & + \left(\psi_{z_+ \Lambda_-}^{\text{priv}} \right)^{-1}(0). \end{aligned}$$

*In the final term, we have $z = z_+ * \lambda_-$, where $z_+ : \alpha \rightarrow \rho'$ is a path of connections along W_+ .*

(ii) *If $\beta \neq \rho$, then $\partial N_z^+(W_-^+, I; \alpha, \beta)$ is given by*

$$\begin{aligned} (-1)^{\dim \alpha} \partial N_z^+(W_-^+, I; \alpha, \beta) = & \sum_{z=z_+ * z_-} M_{z_+}^+(W_+, \pi_+; \alpha, \alpha') \times_{\alpha'} N_{z_-}^+(W_-, \pi_-; \alpha', \beta) \\ & - N_z^+(W_-^+, \pi_-^+(T_0); \alpha, \beta) \\ & + \sum_{z=\zeta * z'} (-1)^{\tilde{i}(\zeta)} \check{M}_{\zeta}^+(Y, \pi_0; \alpha, \gamma) \times_{\gamma} N_{z'}^+(W_-^+, I; \gamma, \beta) \\ & + \sum_{z=z' * \zeta'} (-1)^{\tilde{i}(z')} N_{z'}^+(W_-^+, I; \alpha, \gamma) \times_{\gamma} \check{M}_{\zeta'}^+(Y, \pi_0; \gamma, \beta) \\ & - \left(\psi_{z_+ \Lambda_- \zeta}^{\text{priv}} \right)^{-1}(0). \end{aligned}$$

*where in the final term we have $z = z_+ * \lambda_- * \zeta$.*

Remark 5.35. The computation of boundary faces is analogous to the work concluding Section 5.2, including the discussion of signs. For the final term, the domain of this section is oriented as

$$\alpha \times (T_0, \infty) \times M_{z_+}^{\text{Fiber}}(W_+; \alpha, \rho') \times (-\infty, 0) \times \check{M}_{\zeta}^{\text{Fiber}}(Y, \pi_0; \rho, \beta);$$

here $t \in (T_0, \infty)$ corresponds to the parameter in $\pi_-^+(t)$, while $(-\infty, 0)$ corresponds to how far we translate the instanton on Y to the left before gluing.

To compute the orientation on the zero set, one shuffles $(T_0, \infty) \times (-\infty, 0)$ to the front incurring a sign of $(-1)^{\tilde{i}(z_+)}$, passes to the locus $A_T \cong [0, 1]$ where $T > T_0$ is large (which incurs a sign of -1 because

the outward-normal-first convention is here negative the orientation of the regular zero set convention), and shuffles that back to

$$(-1)^{\dim \alpha - 1} \alpha \times M_{z_+}^{\text{Fiber}}(W_+; \alpha, \rho') \times [0, 1] \times \check{M}_{\zeta}^{\text{Fiber}}(Y, \pi_0; \rho, \beta). \quad \diamond$$

Moving on to the other composite, one may construct a family of perturbations $\pi_+^-(t)$ on the cylinder which are equal to π_1 on the ends, so that $\pi_+^-(\infty)$ is the broken perturbation $\pi_- \sqcup \pi_+ : (Y, \pi_1) \rightarrow (Y, \pi_0) \rightarrow (Y, \pi_1)$. Running through the same construction, we obtain modified moduli spaces $N_z^+(W_+, I)$, and we enumerate the boundary strata (with signs) below. Here when doing gluing on both sides of Λ_- our open set is indexed by $(T_0, \infty) \times (0, \infty)$ (The first term corresponding to the perturbation parameter, the second corresponding to how far left we translate the instanton on $(\mathbb{R} \times Y, \pi_1)$ before gluing), and we instead use the function $L(t, t') = -1/t - 1/t'$ to truncate, with locus $A'_T = \{(t, t') \mid 1/t + 1/t' = 1/T\}$ oriented as the regular preimage $L^{-1}(-1/T)$. The isomorphism $A'_T \cong [0, 1]$ given by $t \mapsto 1 - T/t = T/t'$ is orientation-preserving, and sends the section $\psi_{\zeta' \Lambda_- z_+}^{\text{priv}}$ to

$$s \psi_{\zeta' \Lambda_-}^{\text{priv}}(x) + (1 - s) \cdot \psi_{\Lambda_- z_+}^{\text{priv}}(y).$$

Proposition 5.36. *The boundary of the stratified smooth space $N_z^+(W_+, I; \alpha', \beta')$ as a geometric chain is given as follows.*

(i) *If $\alpha' = \rho'$, then $\partial N_z^+(W_+, I; \rho', \beta')$ is given by*

$$\begin{aligned} (-1)^{\dim \rho'} \partial N_z^+(W_+, I; \rho', \beta') &= \sum_{z=z_- * z_+} N_{z_-}^+(W_-, \pi_-; \rho', \alpha) \times_{\alpha} M_{z_+}^+(W_+, \pi_+; \alpha, \beta') \\ &\quad - N_z^+(W_+, \pi_+(T_0); \rho', \beta') \\ &\quad + \sum_{z=\zeta' * z'} (-1)^{\tilde{i}(\zeta')} \check{M}_{\zeta'}^+(Y, \pi_1; \rho', \gamma') \times_{\gamma'} N_{z'}^+(W_+, I; \gamma', \beta') \\ &\quad + \sum_{z=z' * \zeta} (-1)^{\tilde{i}(z')} N_{z'}^+(W_+, I; \rho', \gamma') \times_{\gamma'} \check{M}_{\zeta}^+(Y, \pi_1; \gamma', \beta') \\ &\quad + \left(\psi_{\Lambda_- z_+}^{\text{priv}} \right)^{-1}(0). \end{aligned}$$

*In the final term, we have $z = \lambda_- * z_+$, where $z_+ : \rho \rightarrow \beta'$ is a path of connections along W_- . Of course, the sign $(-1)^{\dim \rho'}$ is simply 1.*

(ii) *If $\alpha' \neq \rho'$, then $\partial N_z^+(W_+, I; \alpha, \beta)$ is given by*

$$\begin{aligned} (-1)^{\dim \alpha'} \partial N_z^+(W_+, I; \alpha', \beta') &= \sum_{z=z_- * z_+} N_{z_-}^+(W_-, \pi_-; \alpha', \alpha) \times_{\alpha} N_{z_+}^+(W_+, \pi_+; \alpha, \beta') \\ &\quad - N_z^+(W_+, \pi_+(T_0); \alpha', \beta') \\ &\quad + \sum_{z=\zeta' * z'} (-1)^{\tilde{i}(\zeta')} \check{M}_{\zeta'}^+(Y, \pi_1; \alpha', \gamma') \times_{\gamma'} N_{z'}^+(W_+, I; \gamma', \beta') \\ &\quad + \sum_{z=z' * \zeta} (-1)^{\tilde{i}(z')} N_{z'}^+(W_+, I; \alpha', \gamma') \times_{\gamma'} \check{M}_{\zeta}^+(Y, \pi_1; \gamma', \beta') \\ &\quad - (-1)^{\tilde{i}(\zeta')} \left(\psi_{\zeta' \Lambda_- z_+}^{\text{priv}} \right)^{-1}(0), \end{aligned}$$

where in the final term we have $z = \zeta' * \lambda_- * z_+$.

Remark 5.37. For the final term, the domain of this section is oriented as

$$\alpha' \times (T_0, \infty) \times (0, \infty) \times \check{M}_{\zeta'}^{\text{Fiber}}(Y, \pi_1; \alpha', \rho') \times M_{z_+}^{\text{Fiber}}(W_+; \rho, \beta');$$

one shuffles $(T_0, \infty) \times (0, \infty)$ to the front incurring no additional sign, passes to the locus $A'_T \cong [0, 1]$ (which incurs a sign of -1 because the outward-normal-first convention is here negative the orientation of the regular zero set convention), and shuffles that back to

$$(-1)^{\dim \alpha' + \tilde{i}(\zeta') - 1} \alpha' \times M_{\zeta'}^{\text{Fiber}}(Y, \pi_1; \alpha', \rho') \times [0, 1] \times \check{M}_{z_+}^{\text{Fiber}}(W_+; \rho, \beta'). \quad \diamond$$

5.5.2 Homotopies of bimodules from modified moduli spaces

We can now use the constructions of the preceding section to define the desired bimodule homotopies. To check these easily, have on-hand the definitions of \mathcal{W}_+ from Construction 5.21 and \mathcal{W}_- from Construction 5.23, as well as the definition of bimodule homotopy from Definition 4.18. We begin with the easier of the two cases.

For compactness of notation, write $\mathcal{H}_+^-(\alpha, \beta)$ for the modified moduli spaces $N^+(W, \pi_+^-(t); \alpha, \beta)$ from the previous section. Let $\pi_+^-(T_0)$ be the constant perturbation at π_1 , so that $M^+(W, \pi_+^-(T_0))$ defines the identity bimodule on $\mathcal{I}_1 = \mathcal{I}(Y, \pi_1)$.

Proposition 5.38. *The modified moduli spaces above define a bimodule homotopy*

$$\mathcal{H}_+^- : \mathbf{1}_{\mathcal{I}_1} \implies \mathcal{W}_+ \circ \mathcal{W}_-.$$

Proof. We hope that the following diagram indicating the relevant bimodules is helpful when parsing the argument; the composite bimodule is given by summing over the fiber products of all two-step paths running from α' to β' , and we want to find a homotopy from the identity bimodule to this composite bimodule.

$$\begin{array}{ccccc}
 & & & \alpha & \\
 & & & \nearrow & \\
 & & & M(W_-; \alpha', \alpha) & \\
 & & & \curvearrowright & \\
 \alpha' & \xrightarrow{M(W_-; \alpha', \rho) \times_\rho S_\rho} & S_\rho & \xrightarrow{B(W_+; \rho, \beta')} & \beta' \\
 & \searrow^{(-1)^{\tilde{i}(\zeta')} B(\pi_1; \alpha', \rho')} & & & \\
 & & & \curvearrowleft & \\
 & & & Z(W_+; \rho, \beta') & \\
 & & & \searrow & \\
 & & & \bar{\rho} & \\
 & & & \nwarrow & \\
 & & & \delta_{\alpha' = \rho'} \mathcal{O}_- & \\
 & & & \curvearrowright & \\
 & & & M(W_+; \alpha, \beta') & \\
 & & & \searrow & \\
 & & & &
 \end{array}$$

The boundary relations enumerated in Proposition 5.36 almost immediately show that this gives a bimodule homotopy from the identity bimodule. The only interesting stratum is the last one, arising from

obstructed gluing theory; we will focus on the case $\beta' \neq \rho'$, as the case $\beta = \rho$ is simpler. In the homotopy relation, this boundary stratum comes from the composite term, as

$$\mathcal{W}_-(\alpha', \rho) \times_{\rho} \mathcal{W}_+(\rho, \beta')$$

has the extra term

$$-(-1)^{\tilde{i}(\zeta)} B_{\zeta}(\alpha', \rho') \times_{S_{\rho}} B(W; \rho, \beta');$$

because we use $-\psi_{\zeta' \Lambda_-}^{\text{priv}}$ in defining the blowups, Lemma A.39 identifies this with $-(-1)^{\tilde{i}(\zeta)}$ times the zero set of $(1-s)\psi_{\Lambda_- z_+}^{\text{priv}} - s\psi_{\zeta' \Lambda_-}^{\text{priv}}$, which is identified with the zero set of the section $\psi_{\zeta' \Lambda_- z_+}^{\text{priv}}$. So this additional term in the homotopy relation is precisely the same as the additional term in the boundary relation of Proposition 5.36, as desired. \square

Though the obstructed-gluing analysis of the two cases is nearly identical (except for one sign), the definition of the bimodule homotopy running in the other direction, and the verification that it is a bimodule homotopy, are significantly more intricate.

Construction 5.39. For compactness of notation, write $H_{\pm}^+(\alpha, \beta)$ for the modified moduli spaces $N^+(W, \pi_{\pm}^+(t); \alpha, \beta)$ from the previous section. Here we assume that $\pi_{\pm}^+(T_0)$ is the constant perturbation at π_0 , so that $N_z^+(W_{\pm}^+, \pi_{\pm}^+(T_0); \alpha, \beta)$ is either diffeomorphic to $[0, 1] \times \check{M}(\pi_0; \alpha, \beta)$, with endpoint maps determined by the latter factor), or when $\alpha = \beta$ equal to α itself. In the former case, this is a trivial chain (it supports the orientation-reversing isomorphism given by $t \mapsto 1 - t$ in the first factor), hence zero at the chain level, as in Remark 4.16.

Before moving on, observe that $H_{\pm}^+(\rho, \rho)^{\text{red}}$ is diffeomorphic to $[0, 1] \times \rho$, and (because the endpoint maps determine an equivariant map to $\rho \times \rho$, which has two isolated reducible orbits) the endpoint maps factor through the projection to ρ , so this is also a trivial chain.

Over each moduli space $H_{\pm}^+(\rho, \beta)$, we will define certain equivariant sections of the complex line bundle \mathcal{H}_{ρ} , written $\psi_{\rho\beta}^H$. We prescribe the boundary values on the various faces enumerated in Proposition 5.34: $\psi_{\rho\beta}^H$ should be equal to the privileged sections $\psi_{z_+ \Lambda_-}^{\text{priv}}$ on the first listed face; should agree with the privileged sections $\psi_{\zeta' \Lambda_-}^{\text{priv}}$ on the third face; should be pulled back from $\psi_{\rho\gamma}^H$ on the fourth enumerated face; and on the final face — diffeomorphic to $B(W_+; \rho, \rho') \times_{S_{\rho'}} B(\pi_0; \rho, \rho)$ — it should be pulled back from a section $\psi_{\rho\rho}^+$ over the first factor. On the part of the fourth face given by $H_{\pm}^+(\rho, \rho)^{\text{red}} \times_{\rho} \check{M}(\pi_0; \rho, \beta)$, we demand that the section $\psi_{\rho\beta}^H$ factor through the projection to $\check{M}(\pi_0; \rho, \beta)$. This guarantees that even after blowing up, this boundary component is a trivial chain.

Finally, in the case $\beta = \rho$, we only attempt to define a transverse section over the irreducible locus of $H_{\pm}^+(\rho, \rho)$; there is no transverse section over the reducible locus, but we will (mostly) be able to ignore it, as this is a trivial chain.

Define $B_{\pm}^+(\rho, \beta)$ to be the blowup of $\psi_{\rho\beta}^H$, so a geometric chain in $S_{\rho} \times \beta$; let $Z_{\pm}^+(\rho, \beta)$ be the zero set of $\psi_{\rho\beta}^H$, so still a geometric chain in $\rho \times \beta$.

In the special case $\beta = \rho$, we set $\widehat{B}_{\pm}^+(\rho, \rho)$ to be the blowup of the section $\psi_{\rho\rho}^H \sqcup \psi_{\rho\rho}^+$ over

$$H_{\pm}^+(\rho, \rho)^{\text{irr}} \times_{\rho} S_{\rho} - B(W_+; \rho, \rho');$$

transversality dictates that the zero set $\widehat{Z}_-^+(\rho, \rho)$ is empty.

Then the bimodule homotopy


$$\mathcal{H}_-^+ : \mathbf{1}_{\mathcal{B}_\rho \mathcal{I}_0} \implies \mathcal{W}_- \circ \mathcal{W}_+$$

is defined by the formula


$$\mathcal{H}_-^+(\alpha, \beta) = \left(\begin{array}{c|ccc} & \alpha & S_\rho & \bar{\rho} \\ \hline \beta & H_-^+(\alpha, \beta) & -B_-^+(\rho, \beta) & Z_-^+(\rho, \beta) \\ S_\rho & H_-^+(\alpha, \rho) \times_\rho S_\rho - (-1)^{\tilde{i}(z_+)} B_{z_+}(W_+; \alpha, \rho) & -\widehat{B}_H(\rho, \rho) & 0 \\ \bar{\rho} & 0 & 0 & 0 \end{array} \right). \quad \diamond$$

Proposition 5.40. *The formula above defines a bimodule homotopy.*

To aid with the verification, we include a diagram indicating both of the bimodules we are composing.



Next, we depict the relevant flow categories and the bimodule homotopies.



Then the relation we are attempting to prove asserts that for every left-to-right arrow in the second diagram (from, say, α to β), its boundary is a sum over all two-step paths in the second diagram which start at α and end at β , plus the corresponding fiber product in the first diagram, minus $\delta_{\alpha=\beta}\alpha$.

Sketch of proof. Because this computation is intricate and delicate, we will avoid discussing the signs (which would only serve to obfuscate the discussion); the careful reader should not have difficulty inserting the signs into the arguments below. That $\mathcal{H}_-^+(\alpha, \beta)$ satisfies the homotopy relation for $\alpha, \beta \neq \rho$ follows exactly as in the computation that \mathcal{W}_- is a bimodule.

The novel parts are those involving moduli spaces of the form $H_-^+(\rho, \rho)$. We will only discuss those parts of the bimodule relation which involve this term below, as all other parts of the discussion follow the same strategy as previous arguments; the reader should understand this as a guide to clarifying potentially confusing details when running through the verification.

- In entry (1, 2), the term $\mathcal{H}_-^+(S_\rho, S_\rho) \times_{S_\rho} \mathcal{B}_\rho \mathcal{C}(S_\rho, \beta)$ is given by

$$B(H_-^+(\rho, \rho)^{\text{irr}} \times_\rho S_\rho + B(W_+; \rho, \rho')) \times_{S_\rho} B_\zeta(\rho, \beta);$$

the left-hand blowup is determined entirely by the first factor, so we focus on this last. The latter term simplifies to give the $(\psi_{z_+ \Lambda_- \zeta}^{\text{priv}})^{-1}(0)$ term, blown up on the left-hand side (this is where we use the assumption that $\psi_{\rho\beta}^H$ is pulled back from a section $\psi_{\rho\rho}^H$ on the first factor here). The former term simplifies to give

$$B(H_-^+(\rho, \rho)^{\text{irr}} \times_\rho \check{M}_\zeta(\rho, \beta));$$

by the assumption that the section on this term is pulled back from the first factor, this simplifies to $B_-^+(\rho, \rho)^{\text{irr}} \times_\rho \check{M}_\zeta(\rho, \beta)$, which is one of the terms promised by the boundary relation Proposition 5.34. Finally, this leaves one term in the known boundary relation unaccounted for: the blowup of $H_-^+(\rho, \rho)^{\text{red}} \times_\rho \check{M}_\zeta(\rho, \beta)$. However, we assumed that the section used to define the blowup factors through a projection to the latter factor (and every stratum of the first factor has positive dimension), so that this chain has small rank on each positive-dimensional stratum. In particular, the corresponding blown-up chain is degenerate, so that this term is zero on the chain level. This squares away the unexpected factors in both the bimodule relation and the boundary relation Proposition 5.34.

Entry (1, 3) is similar but strictly simpler because $\widehat{Z}_-^+(\rho, \rho) = \emptyset$.

- In entry (2, 1), the novel terms arise from $\mathcal{B}_\rho \mathcal{C}(\alpha, S_\rho) \times_{S_\rho} \mathcal{H}_-^+(S_\rho, S_\rho)$, which may be expanded as

$$\check{M}(\alpha, \rho) \times_\rho S_\rho \times_{S_\rho} B(H_-^+(\rho, \rho)^{\text{irr}} \times_\rho S_\rho + B(W_+; \rho, \rho'));$$

cancelling out $S_\rho \times_{S_\rho}$ and replacing the blowup with the original chain, this simplifies to

$$\check{M}(\alpha, \rho) \times_\rho H_-^+(\rho, \rho)^{\text{irr}} \times_\rho S_\rho + \check{M}(\alpha, \rho) \times_\rho M(W_+; \rho, \rho'),$$

which contribute the terms corresponding to factorization through ρ to the boundary relations of both $H_-^+(\alpha, \rho) \times_\rho S_\rho$ and $B_{z_+}(W_+; \alpha, \rho)$. In the former case, this still leaves an absent term corresponding to $H_-^+(\rho, \rho)^{\text{red}}$; but this is a degenerate geometric chain, so zero at the chain level, so irrelevant in checking the boundary relation.

- In entry (3, 1), there are only two terms in the bimodule homotopy relation for $\partial\mathcal{H}_-^+(\alpha, \bar{\rho})$ which are not tautologically empty: the first is

$$\mathcal{M}_1(\alpha, \bar{\rho}) = \mathcal{W}_+(\alpha, \rho') \times_{\rho'} \mathcal{W}_-(\rho', \bar{\rho}) = M(W_+; \alpha, \rho') \times_{\rho'} \mathcal{O}_- \cong M(W_+; \alpha, \rho'),$$

and the second is

$$\mathcal{H}_-^+(\alpha, S_\rho) \times_{S_\rho} \mathcal{B}_\rho \mathcal{C}(S_\rho, \bar{\rho}) = (H_-^+(\alpha, \rho) \times_\rho S_\rho + B_{z_+}(W_+; \alpha, \rho)) \times_{S_\rho} S_\rho,$$

which simplifies to

$$H_-^+(\alpha, \rho) \times_\rho S_\rho + B_{z_+}(W_+; \alpha, \rho)$$

as a chain in $\alpha \times \rho$.

Now the first term is degenerate — it factors through the projection to $H_-^+(\alpha, \rho)$, so has small rank — and considering the second as a chain in $\alpha \times \rho$ instead of $\alpha \times S_\rho$ leaves it equivalent to the chain $M(W_+; \alpha, \rho)$. Thus this term cancels out with the previous.

- Entry (2, 2) is by far the most intricate; there is no avoiding talking about every component. First, let us point out that in

$$\widehat{B}_H(\rho, \rho) = B(H_-^+(\rho, \rho)^{\text{irr}} \times_\rho S_\rho + B(W_+; \rho, \rho')),$$

the first term has no boundary: this moduli space has fiber-dimension 1 over S_ρ , so that $H_-^+(\rho, \rho)^{\text{irr}}$ is 3-dimensional, so its quotient by $SO(3)$ is hence a 0-manifold, which has no boundary. Because the section over this space is non-vanishing, taking the left blowup is equivalent to choosing a lift of the map to ρ to a map to S_ρ ; no new boundary terms appear.

So we are engaged in a computation of the boundary of the second term. Let us first write the boundary predicted by the bimodule homotopy. This consists of four terms:

$$\begin{aligned} \sum_\gamma \mathcal{B}(S_\rho, \gamma) \times_\gamma \mathcal{H}(\gamma, S_\rho) & \quad \sum_\gamma \mathcal{H}(S_\rho, \gamma) \times_\gamma \mathcal{B}(\gamma, S_\rho) \\ \sum_{\gamma'} \mathcal{W}_+(S_\rho, \gamma') \times_{\gamma'} \mathcal{W}_-(\gamma', S_\rho) & \quad \mathcal{M}_0(S_\rho, S_\rho) = S_\rho. \end{aligned}$$

Spelling these out, this amounts to S_ρ and six additional boundary terms:

$$B(\pi_0; \rho, \gamma) \times_\gamma H_-^+(\gamma, \rho) \times_\rho S_\rho \tag{5.41}$$

$$B(\pi_0; \rho, \gamma) \times_\gamma B(W_+; \gamma, \rho') \tag{5.42}$$

$$B_-^+(\rho, \gamma) \times_\gamma \check{M}(\pi_0; \gamma, \rho) \times_\rho S_\rho \tag{5.43}$$

$$B(W_+; \rho, \gamma') \times_{\gamma'} M(W_-; \gamma', \rho) \times_\rho S_\rho \tag{5.44}$$

$$B(W_+; \rho, \gamma') \times_{\gamma'} B(\pi_1; \gamma', \rho') \tag{5.45}$$

The boundary has fiber-dimension zero, so we may discard all terms which necessarily have positive fiber-dimension: the first, third, and fourth do not appear. What we are left with is precisely the boundary relation for $B(W_+; \rho, \rho')$; the final term S_ρ corresponds to $Z(W_+; \rho, \rho') \times_\rho S_\rho$, as this is identifiable with $\mathcal{O}_+ \times_\rho S_\rho = S_\rho$, cancelling out $\mathcal{M}_0(S_\rho, S_\rho) = S_\rho$. So this is indeed the boundary of $\widehat{B}_H(\rho, \rho)$, as desired.

- In entry (3, 3), the only terms in the bimodule homotopy relation which aren't tautologically zero (or of fiber-dimension ≥ 8) are $\mathcal{M}_1(\bar{\rho}, \bar{\rho}) - \mathcal{M}_0(\bar{\rho}, \bar{\rho})$. The latter is a copy of $\bar{\rho}$ itself, with the identity map as its endpoint maps. The first is a sum over fiber products, but the only term which isn't tautologically zero is $Z(W_+; \bar{\rho}, S_\rho) \times_{\bar{\rho}} \mathcal{O}_-$. Both factors consist of a single reducible orbit, so that the fiber product is also a single reducible orbit; the reducible locus in $H_-^+(\rho, \rho)^{\text{red}}$, diffeomorphic to $[0, 1] \times \rho$, gives a homotopy between their endpoint maps. Because the reducible orbits in $\rho \times \rho$ are isolated, this means the endpoint maps can be identified, so that $\mathcal{M}_1(\bar{\rho}, \bar{\rho}) - \mathcal{M}_0(\bar{\rho}, \bar{\rho}) = \bar{\rho} - \bar{\rho} = 0$. So this term in the bimodule homotopy relation is satisfied.

We leave boundary relations (2, 3) and (3, 2) to the interested reader; the arguments are no more difficult. \square

This completes the proof of Theorem 5.26.

6 Functoriality of equivariant instanton homology

6.1 Cobordism maps in the pseudo-unobstructed case

The results we have thusfar are enough to prove that equivariant instanton homology is a well-defined invariant of rational homology sphere. We described in Definition 5.28 how to define a group which is transparently independent of perturbation, which we called $I^\bullet(Y, w)$: we identify the various $I^\bullet(Y, w, \pi)$ using the continuation maps $\varphi_{\pi \rightarrow \pi'}$ whenever $\sigma_\pi \leq \sigma_{\pi'}$, which we have established are isomorphisms on equivariant homology. We would like to upgrade $I^\bullet(Y, w)$ to a functorial invariant. The goal of this section is to establish the following result, whose proof is given at the end of the section.

Theorem 6.1. *If $(W, c) : (Y, w) \rightarrow (Y', w')$ is a cobordism for which either $b^+(W) = 0$ or (W, c) supports no central connections, there is a well-defined induced map $I^\bullet(W, c) : I^\bullet(Y, w) \rightarrow I^\bullet(Y', w')$ on equivariant instanton homology, which is a module map in the appropriate sense and respects the exact triangles.*

The statement that (W, c) supports no central connections is equivalent to the claim that $PD(c) \in H^2(W; \mathbb{Z})$ is not twice some other class, or equivalently, there exists a closed and oriented surface $\Sigma \subset W$ with $\#(\Sigma \cap c)$ odd.

We would like to argue as follows. Because $I^\bullet(Y, w)$ is independent of the choice of perturbation, we might hope to choose perturbations π, π' so that $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is an unobstructed cobordism. Then we could define the induced map on equivariant instanton homology $I^\bullet(Y, w, \pi) \rightarrow I^\bullet(Y', w', \pi')$ by applying the construction of Theorem 4.21 and then composing with the equivariant homology functors. This induced map $(W, c)_{\pi \rightarrow \pi'}$ descends to a well-defined map on $I^\bullet(Y, w)$; to see this, suppose (π_0, π'_0) is another pair of perturbations for which (W, c) is unobstructed and $\sigma_{\pi_0} \leq \sigma_\pi$ while $\sigma_{\pi'} \leq \sigma_{\pi'_0}$. Then the fact that unobstructed cobordism maps compose functorially implies that

$$\varphi_{\pi' \rightarrow \pi'_0} (W, c)_{\pi \rightarrow \pi'} \varphi_{\pi_0 \rightarrow \pi} = (W, c)_{\pi_0 \rightarrow \pi'_0}.$$

We will discuss the structure of this argument in more detail towards the end of this section.

Unfortunately, we usually cannot find perturbations π, π' so that (W, c) is unobstructed: the impish pseudocentral reducibles are unimpressed by our ability to change perturbations on the ends. Their normal index $N(\Lambda; \pi, \pi')$ is independent of π and π' , and this quantity may well be negative. By Lemma 3.33, the best we can do is find perturbations (π, π') so that $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is ‘pseudo-unobstructed’: all its reducibles are either unobstructed or pseudocentral.

To prove Theorem 6.1, we must construct cobordism maps $\tilde{C}(Y, w, \pi) \rightarrow \tilde{C}(Y', w', \pi')$ associated to pseudo-unobstructed cobordisms, which are well-defined up to equivariant homotopy and compose correctly with the continuation maps.

Now recall the pentachotomy of Proposition 3.32. In cases (i) and (v), the relevant moduli spaces are cut out transversely (even in families), so that we have already constructed the relevant cobordism maps and verified that they are well-defined up to equivariant homotopy, by Theorem 4.21. In cases (ii) and (iv), we have an induced map associated to a regular perturbation π_W , but one has to argue carefully to show

that the induced map is well-defined up to equivariant homotopy. In case (iii) one has to argue carefully to show that the induced map exists at all.

First we will handle case (iii) of pseudo-unobstructed cobordisms with $b_1(W) - b^+(W) = 0$, explaining both how to guarantee that an induced map exists and why it is well-defined.

Let $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ be a pseudo-unobstructed cobordism with $b_1(W) - b^+(W) = 0$; by Proposition 3.32(iii), we may choose a perturbation π_W on (W, c) so that all instantons are cut out transversely except these pseudocentral reducibles, each of which corresponds to a 2-dimensional orbit \mathcal{O}_Λ of framed instantons on W whose linearized ASD operator has trivial kernel and cokernel isomorphic to \mathbb{C} .

Given $\theta \in \mathfrak{Z}(Y, w)$ and $\theta' \in \mathfrak{Z}(Y', w')$, we will write $\mathcal{P}(W, c; \theta, \theta')$ for the space of pseudocentral instantons on (W, c) running from θ to θ' . The $SO(3)$ -space \mathcal{P} is a disjoint union of orbits \mathcal{O}_Λ , each of which is diffeomorphic to S^2 . There is an obstruction bundle $\mathcal{H}_P \rightarrow \mathcal{P}$, an $SO(3)$ -equivariant complex line bundle, assigning to each instanton $\Lambda \in \mathcal{P}$ the cokernel of the associated linearized ASD operator.

We will write $\mathfrak{P}(W, c; \theta, \theta') = \pi_0 \mathcal{P}(W, c; \theta, \theta')$ for a corresponding discrete space of components. When (W, c) or (θ, θ') are clear from context, we will suppress them from notation.

Let us briefly discuss orientations; the naturally oriented object is the 0-dimensional space $\mathfrak{P}(W, c; \theta, \theta')$. To see this, recall the conventions of Section 3.7. So long as one chooses an element of $\Lambda(\theta)$ and $\Lambda(\theta')$ and a homology orientation on the cobordism, the moduli space $\mathfrak{P}(W, c; \theta, \theta')$ is canonically oriented. More precisely, these choices induce an orientation on the determinant line of the complex $T_\Lambda \mathcal{P} \xrightarrow{0} \mathcal{H}_\Lambda$; perturbing this so that the map is a nonzero Γ_Λ -equivariant map, the determinant line corresponds to the orientation of a generic zero set of a Γ_Λ -equivariant section of \mathcal{H}_Λ .

Such a generic section has two zeroes, which are canonically oriented with the same orientation. We thus consider this an orientation on the set $\pi_0 \mathcal{P} = \mathfrak{P}$.

Remark 6.2. Each element of $\mathfrak{P}(W, c; \theta, \theta')$ is oriented with the same sign, because the corresponding deformation complexes are homotopic. This sign is the same as the sign on the set of central connections $\mathfrak{Z}(W, c; \theta, \theta')$, for the same reason. This is mostly interesting when $b_1(W) = b^+(W) = 0$; when $b^+(W) > 0$ our standing assumption is that (W, c) supports no central connections, and one can perturb away the abelian connections. \diamond

Fix an orbit $\mathcal{O}_\Lambda \subset \mathcal{P}(W, c; \theta, \theta')$, and write $\lambda : \theta \rightarrow \theta'$ for the corresponding path of connections along (W, c) . When carrying out the obstructed gluing theory of Section 5.1, one finds that the ends of the moduli spaces $M_z^+(W, c; \alpha, \alpha')$ corresponding to gluing along \mathcal{O}_Λ are modeled on the zero set of an appropriate section:

- (l) If $\alpha' = \theta'$ and $z = \zeta * \lambda$, the corresponding end is modeled on the zero set of an $SO(3)$ -equivariant section of the pullback bundle over $(0, \infty) \times \check{M}_\zeta^+(Y, w; \alpha, \theta) \times \mathcal{O}_\Lambda$.
- (r) If $\alpha = \theta$ and $z = \lambda * \zeta'$, the corresponding end is modeled on the zero set of an $SO(3)$ -equivariant section of the pullback bundle over $\mathcal{O}_\Lambda \times (-\infty, 0) \times \check{M}_{\zeta'}^+(Y', w'; \theta', \alpha')$.
- (c) If $\alpha \neq \theta$ and $\alpha' \neq \theta'$ and $z = \zeta * \lambda * \zeta'$, the corresponding end is modeled on the zero set of an

$SO(3)$ -equivariant section of the pullback bundle over

$$(0, \infty) \times \check{M}_\zeta^+(Y, w; \alpha, \theta) \times \mathcal{O}_\Lambda \times (-\infty, 0) \times \check{M}_{\zeta'}^+(Y', w'; \theta', \alpha').$$

Applying the modified moduli space ideas of Section 5.2, we interpolate between these sections and a pre-chosen section. In choosing these pre-chosen sections, we will want them to satisfy some special properties.

For each pseudocentral reducible Λ , we choose priveleged $SO(3)$ -equivariant sections

$$\begin{aligned} \psi_{\zeta\Lambda}^{\text{priv}} &: \check{M}_\zeta^+(Y, w; \alpha, \theta) \times \mathcal{O}_\Lambda \rightarrow \mathcal{H}_\Lambda \\ \psi_{\Lambda\zeta'}^{\text{priv}} &: \mathcal{O}_\Lambda \times \check{M}_{\zeta'}^+(Y', w'; \theta', \alpha') \rightarrow \mathcal{H}_\Lambda \\ \psi_{\zeta\Lambda\zeta'}^{\text{priv}} &: \check{M}_\zeta^+(Y, w; \alpha, \theta) \times \mathcal{O}_\Lambda \times [0, 1] \times \check{M}_{\zeta'}^+(Y', w'; \theta', \alpha') \rightarrow \mathcal{H}_\Lambda \end{aligned}$$

which satisfy conditions precisely analagous to those enumerated in Section 5.2. In small degrees, we can understand the zero sets of these sections explicitly.

- When $\tilde{i}(\zeta) = 1$, the first map has domain equivariantly diffeomorphic to a disjoint union of copies of $SO(3) \times S^2$; an $SO(3)$ -equivariant section of this bundle may be identified with a section of the complex line bundle over S^2 of Euler class 2. Equivariance implies that the projection map $\pi : \psi^{-1}(0) \rightarrow SO(3)$ is a local diffeomorphism, and $\psi^{-1}(0)$ consists of a number of copies of $SO(3)$. The Euler class computation shows that the signed count of these is two, so that $\pi : \psi^{-1}(0) \rightarrow SO(3)$ identifies $\psi^{-1}(0)$ as twice the fundamental class in $C_3^{gm} SO(3) = \mathbb{Z}$.
- When $\tilde{i}(\zeta') = \dim \alpha + 1$, the second map has domain equivariantly diffeomorphic to a disjoint union of copies of $S^2 \times SO(3)$, and a similar discussion applies: $\psi^{-1}(0)$ can be identified with twice the fundamental class of $SO(3)$.

Now in all of these cases, $(\psi^{\text{priv}})^{-1}(0)$ takes a particularly simple form. In what follows, $\{\Lambda\}$ is oriented as the corresponding element of $\mathfrak{P}(W, c; \theta, \theta')$.

- In the first case, $(\psi_{\zeta\Lambda}^{\text{priv}})^{-1}(0)$ has dimension $\tilde{i}(\zeta) + \dim \alpha - 1$. As a geometric chain in $C_*(\alpha \times \theta')$, this is zero unless $\tilde{i}(\zeta) = 1$, in which case this geometric chain coincides with $2\check{M}_\zeta \times \{\Lambda\}$, as discussed above.
- In the second case, $(\psi_{\Lambda\zeta'}^{\text{priv}})^{-1}(0)$ has dimension $\tilde{i}(\zeta') - 1$; this is always at least $\dim \alpha'$, and this chain is zero as an element of $C_*(\theta \times \alpha')$ unless $\tilde{i}(\zeta') = 1 + \dim \alpha$, in which case it coincides with $-2\{\Lambda\} \times \check{M}_{\zeta'}$ as above. The sign arises because $\{-\infty\}$ is a negatively-oriented boundary point in $[-\infty, 0)$.
- In the third case, $(\psi_{\zeta\Lambda\zeta'}^{\text{priv}})^{-1}(0)$ has dimension $\tilde{i}(\zeta) + \tilde{i}(\zeta') + \dim \alpha - 1$. The term $\tilde{i}(\zeta) + \dim \alpha - 1$ is the dimension of $\check{M}_\zeta(\alpha, \theta)$, which is always at least $\dim \alpha$; the term $\tilde{i}(\zeta') - 1$ is the dimension of $\check{M}_{\zeta'}(\theta', \alpha')$, which is always at least $\dim \alpha'$. It follows that

$$\tilde{i}(\zeta) + \tilde{i}(\zeta') + \dim \alpha - 1 = (\tilde{i}(\zeta) + \dim \alpha - 1) + (\tilde{i}(\zeta') - 1) + 1 \geq \dim \alpha + \dim \alpha' + 1,$$

and the chain complex $C_*(\alpha \times \alpha')$ is supported in degrees $* \leq \dim \alpha + \dim \alpha'$. So this chain is zero on the chain level.

All this being said, one may follow the recipe of Section 5.2 to produce *modified moduli spaces* associated to (W, c, π_W) : these will be moduli spaces $N_z^+(W, c; \alpha, \alpha')$ with the usual boundary terms together with, for each Λ , two oppositely-oriented boundary terms corresponding to the section of \mathcal{H}_Λ arising from obstructed gluing theory, and additional boundary terms corresponding to the zero set of $\psi_{\zeta\Lambda}^{\text{priv}}$ or $\psi_{\Lambda\zeta'}^{\text{priv}}$ or $\psi_{\zeta\Lambda\zeta'}^{\text{priv}}$ as appropriate.

Proposition 6.3. *Let $N_z^+(W, c; \alpha, \alpha')$ be the modified moduli spaces constructed above. As geometric chains in $\alpha \times \alpha'$, we have the following equalities.*

- If $\alpha' = \theta'$ is central, then

$$\begin{aligned} (-1)^{\dim \alpha} \partial N^+(W, c; \alpha, \theta') &= \sum_{\beta} \check{M}^+(Y, w; \alpha, \beta) \times_{\beta} N^+(W, c; \beta, \theta') \\ &\quad + \sum_{\beta'} (-1)^{\tilde{i}(W; \alpha, \beta')} N^+(W, c; \alpha, \beta') \times_{\beta'} \check{M}^+(Y', w'; \beta', \theta') \\ &\quad + \sum_{\theta} 2\check{M}^+(Y, w; \alpha, \theta) \times \mathfrak{P}(W, c; \theta, \theta'). \end{aligned}$$

- If $\alpha = \theta$ is central, then

$$\begin{aligned} \partial N^+(W, c; \theta, \alpha') &= \sum_{\beta} \check{M}^+(Y, w; \theta, \beta) \times_{\beta} N^+(W, c; \beta, \alpha') \\ &\quad + \sum_{\beta'} (-1)^{\tilde{i}(W; \theta, \beta')} N^+(W, c; \theta, \beta') \times_{\beta'} \check{M}^+(Y', w'; \beta', \alpha') \\ &\quad + \sum_{\theta'} -2\mathfrak{P}(W, c; \theta, \theta') \times \check{M}^+(Y', w'; \theta', \alpha'). \end{aligned}$$

- If neither α nor α' is central, then

$$\begin{aligned} (-1)^{\dim \alpha} \partial N^+(W, c; \alpha, \alpha') &= \sum_{\beta} \check{M}^+(Y, w; \alpha, \beta) \times_{\beta} N^+(W, c; \beta, \alpha') \\ &\quad + \sum_{\beta'} (-1)^{\tilde{i}(W; \alpha, \beta')} N^+(W, c; \alpha, \beta') \times_{\beta'} \check{M}^+(Y', w'; \beta', \alpha'). \end{aligned}$$

Given this, we may define a bimodule $\mathcal{I}(W, c)$ associated to the cobordism (W, c) , together with the necessary additional data of a homology orientation and a path between the basepoints on the boundary components. We define this by the formula

$$\mathcal{I}(W, c)(\alpha, \alpha') = N^+(W, c; \alpha, \alpha') + 2\mathfrak{P}(W, c; \alpha, \alpha'). \quad (6.4)$$

Here we interpret the final term $\mathfrak{P}(W, c; \alpha, \alpha')$ as a discrete set of points which is empty unless α, α' are both central, in which case it is the set of pseudocentral reducibles running from α to α' , with the

orientation discussed at the start of the section. Notice that if $\Lambda : \alpha \rightarrow \alpha'$ is pseudocentral, we have $\tilde{i}(W; \alpha, \alpha') = 0$, so this has the correct expected dimension.

This gives an induced map between instanton chain complexes for a chosen pseudo-regular perturbation π_W on (W, c) . When arguing that the induced map is well-defined, one may carry out essentially the same construction for the moduli spaces $M_z^+(W, c, I; \alpha, \alpha')$ for a 1-parameter family of metrics and perturbations: one needs to delete a neighborhood of the stratum $[0, 1] \times \mathcal{P}(W, c)$ from all the compactified moduli spaces for 1-parameter families of metric and perturbation.

Thus the boundary relations above (and the corresponding version in the case of 1-parameter families) establishes the following.

Corollary 6.5. *Whenever $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is a pseudo-unobstructed cobordism with $b_1(W) - b^+(W) = 0$, there is a natural map $\tilde{C}(Y, w, \pi) \rightarrow \tilde{C}(Y', w', \pi')$, independent of the choice of perturbation on (W, c) up to equivariant chain homotopy.*

Before moving on to the remaining cases, we take this opportunity to point out that the components of $\tilde{C}(W, c)$ are determined by explicit cohomological information. Write $r_W : H^2(W; \mathbb{Z})_{\text{tors}} \rightarrow H^2(Y; \mathbb{Z}) \oplus H^2(Y'; \mathbb{Z})$. Recall here from Lemma 2.1 that a central connection θ on (Y, w) corresponds to a cohomology class $x \in H^2(Y; \mathbb{Z})$ with $2x = PD(w)$, and similarly on (Y', w') . The statement of the following proposition uses the set of pairs

$$\tilde{Z}(W, c; \theta, \theta') = \{(y, y') \in H^2(W; \mathbb{Z}) \times H^2(W; \mathbb{Z}) \mid r_W(y) = r_W(y') = (x, x') \text{ and } y - y' \text{ is torsion}\}.$$

Proposition 6.6. *Suppose $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is a pseudo-unobstructed cobordism with $b_1(W) = b^+(W) = 0$. If $\theta \in \mathfrak{Z}(Y, w)$ and $\theta' \in \mathfrak{Z}(Y', w')$ and x, x' are the corresponding cohomology classes, the component of the map $\tilde{C}(W, c)$ running from $C_*(\theta) = \mathbb{Z}$ to $C_*(\theta') = \mathbb{Z}$ is multiplication by $\pm |\tilde{Z}(W, c; \theta, \theta')|$.*

Proof. The only component of $N^+(W, c; \theta, \theta')$ which is 0-dimensional is the component of central connections. By the definition of $\mathcal{I}(W, c)(\theta, \theta')$ in Equation 6.4, the induced map $\theta \rightarrow \theta'$ is given by the signed count of $\mathfrak{Z}(W, c; \theta, \theta') + 2\mathfrak{P}(W, c; \theta, \theta')$. Finally, the sign on every point is the same by Remark 6.2, so we just need to determine the number of points in this count.

Recall from Proposition 3.22 that a central connection on (W, c) extending θ and θ' corresponds to a cohomology class $y \in H^2(W; \mathbb{Z})$ with $2y = PD(c)$ and $r_W(y) = (x, x')$, where x and x' are the cohomology classes corresponding to θ and θ' . An abelian connection running between these is identified with a pair $\{y, y'\}$ with $y + y' = PD(c)$ and $y \neq y'$ for which $r_W(y) = r_W(y') = (x, x')$. Finally, Definition 3.28 says that such an abelian connection is called pseudocentral precisely when $(y' - y)^2 = 0$. Because $b^+(W) = 0$, this is true if and only if $y' - y$ is torsion.

There is a natural involution $\iota(y, y') = (y', y)$; the pseudocentral reducibles correspond to free orbits of ι , while the central reducibles correspond to fixed points of ι . It follows that

$$|\tilde{Z}(W, c; \theta, \theta')| = |\mathfrak{Z}(W, c; \theta, \theta')| + 2|\mathfrak{P}(W, c; \theta, \theta')|,$$

as desired. □

Remark 6.7. Consider the special case that $PD(c)$ is the sum of a torsion class and an even class which vanishes on the boundary: $PD(c) = 2a + t$ where $r_W(a) = 0$. Then the set $\tilde{Z}(W, c; \theta, \theta')$ is in bijection with $r_W^{-1}(x, x')$, via the map $z \mapsto (a + z, a + t - z)$. When w, w' , and c are furthermore empty (so that the corresponding $U(2)$ -bundle is trivial) and θ, θ' are taken to be the trivial connection, this set is in bijection with $\ker(r_W)$, and in particular guaranteed to be nonempty. \diamond

We move on to verifying that the induced map is well-defined in the remaining cases.

Proposition 6.8. *Suppose $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is a pseudo-unobstructed cobordism with $b_1(W) - b^+(W) = -1$. For a pseudo-regular path of perturbations π_W^t , the modified moduli spaces $N_z^+(W, c, I; \alpha, \alpha')$ define an $SO(3)$ -equivariant bimodule homotopy.*

Proof. Recall that being pseudo-regular means here that the ASD operator associated to Λ has trivial kernel and cokernel isomorphic to \mathbb{C}^2 , which assembles into an $SO(3)$ -equivariant rank 2 complex vector bundle \mathcal{H}_Λ over \mathcal{O}_Λ .

Given a pseudocentral $\Lambda : \theta \rightarrow \theta'$ in the family moduli space, consider a neighborhood of $\check{M}_\zeta(\alpha, \theta) \times \mathcal{O}_\Lambda$ in $M_z^+(W, c, I; \theta, \alpha')$; here $\tilde{i}(z) = \tilde{i}(\zeta) + \tilde{i}(\lambda) = \tilde{i}(\zeta) - 3$. Truncating and modifying as usual, we see that the modified moduli spaces will have an additional boundary component corresponding to the zero set of a priveleged section

$$\check{M}_\zeta(\alpha, \theta) \times \mathcal{O}_\Lambda \rightarrow \mathcal{H}_\Lambda.$$

The zero set is a free $SO(3)$ -space with dimension $\dim \alpha + \tilde{i}(\zeta) - 3$, so can only be nonempty if this quantity exceeds 3; it is also only non-trivial in $C_*(\alpha \times \theta')$ if this quantity is at most $\dim \alpha$. These are only simultaneously possible when α is irreducible and $\tilde{i}(\zeta) = 3$.

In this case, one may restrict this $SO(3)$ -equivariant section to a non-equivariant map $\check{M}_\zeta^{\text{Fiber}}(\alpha, \theta) \times \mathcal{O}_\Lambda \rightarrow \mathcal{H}_\Lambda$, and conversely induce an $SO(3)$ -equivariant section on the total space from a non-equivariant one on this subspace. Because this rank 2 complex vector bundle is pulled back from the second factor, there exists a nowhere-vanishing section (for instance, the composite of projection to the second factor and a nowhere-vanishing section $\mathcal{O}_\Lambda \rightarrow \mathcal{H}_\Lambda$, which exists for dimension reasons). Choose such a non-vanishing section for $\psi_{\zeta\Lambda}^{\text{priv}}$. Running the same argument on the other side, we may do the same for $\psi_{\Lambda\zeta'}$.

Then the boundary relations analogous to Proposition 6.3 have no factors corresponding to factorization through Λ in the dimension range where such factors could be nonzero, so the modified moduli spaces define a bimodule homotopy. \square

This establishes that the induced map is well-defined up to equivariant homotopy in the case $b_1(W) - b^+(W) = -1$. The only remaining case in the pentachotomy is $b_1(W) - b^+(W) = 1$. This case is in some sense more irritating than the others, because in 1-parameter families there may exist pseudocentral reducibles which have both non-trivial kernel and cokernel. The preceding gluing analysis does not apply here. Fortunately, this also doesn't matter: the obstructed reducibles only appear in moduli spaces of large dimension.

Lemma 6.9. *If $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is a pseudo-unobstructed cobordism with $b_1(W) - b^+(W) = 1$, and $M_z^+(W, c, I; \alpha, \alpha')$ is a family moduli space which contains a broken trajectory factoring*

through a pseudocentral reducible $\Lambda : \theta \rightarrow \theta'$, we have $\tilde{i}(z) + 1 \geq \dim \alpha' + 4$, so that this moduli space has dimension at least $\dim \alpha + \dim \alpha' + 4$.

Proof. Supposing this moduli space contains the broken instanton $(\mathbf{A}, \Lambda, \mathbf{A}')$, write $z = \zeta * \lambda * \zeta'$, where $\zeta : \alpha \rightarrow \theta$ and $\zeta' : \theta' \rightarrow \alpha'$ are the paths of connections along Y (resp Y') given by \mathbf{A} (resp \mathbf{A}'); here ζ and ζ' may be constant if, for instance, this moduli space runs from θ to α' or α to θ' . For any transversely cut out trajectory on a cylinder $\zeta : \alpha \rightarrow \theta$ we have $\tilde{i}(\zeta) \geq 0$ (with $\tilde{i}(\zeta) \geq 1$ if this is a nonconstant trajectory) and for $\zeta' : \theta' \rightarrow \alpha'$ we have $\tilde{i}(\zeta') \geq \dim \alpha'$ because the fiber over θ' is the moduli space itself, which maps submersively onto α' (with this framed index being larger than $\dim \alpha' + 1$ if the trajectory is nonconstant).

Now in general for a reducible $\Lambda : \theta \rightarrow \theta'$ we have

$$\tilde{i}(\lambda) = (b_1(W) - b^+(W) - \dim \Gamma_\Lambda) + N(\Lambda) + 3;$$

the first term is the index of the deformation complex internal to the $SO(2)$ -reducible locus, the second the index of the deformation complex normal to that locus, and the last term $\dim \Gamma_\theta$, as in Definition 3.10. If Λ is pseudocentral, $N(\Lambda) = 2(b_1(W) - b^+(W) - 1)$ and $\dim \Gamma_\Lambda = 1$. so that this reduces to $\tilde{i}(\lambda) = 3(b_1(W) - b^+(W)) = 3$ in our case of interest.

Thus in this case $\tilde{i}(\lambda) = 3$. By additivity, we have $\tilde{i}(z) = \tilde{i}(\zeta) + \tilde{i}(\lambda) + \tilde{i}(\zeta') \geq 3 + \dim \alpha'$, which gives the inequality stated in the Lemma. \square

Now a variant on Lemma 4.7 stated for bimodule homotopies instead of flow categories shows that to define the relevant bimodule homotopy, we only need those moduli spaces $M_z^+(W, c, I; \alpha, \alpha')$ of dimension $\leq \dim \alpha + \dim \alpha' + 1$. In this dimension range, all of our moduli spaces are smooth and satisfy the expected boundary relations. It follows that even in this case the induced map $\tilde{C}(W, c, \pi_W)$ is independent of π_W up to equivariant homotopy, which we have now established in all cases.

Proof of Theorem 6.1. Observe that when (W_1, c_1) and (W_2, c_2) are composable pseudo-unobstructed cobordisms, the composite is also pseudo-unobstructed, and the bimodules described above compose functorially up to equivariant homotopy (by a small variation on the standard argument).

For each regular perturbation π we have a well-defined equivariant instanton homology group $I^\bullet(Y, w; \pi)$. We write $\pi \leq \pi'$ to mean that their signature data satisfy $\sigma_\pi \leq \sigma_{\pi'}$. For each $\pi \leq \pi'$ we have a continuation map $\varphi_{\pi \rightarrow \pi'}$ between these, which is an isomorphism by Theorem 5.26, and these compose as you expect. The homology group $I^\bullet(Y, w)$ is defined by identifying all of these groups for varying $I^\bullet(Y, w; \sigma)$ using the continuation maps.

To rephrase this, we have a preordered set $\Pi(Y, w)$, the set of regular perturbations π on (Y, w) . When we have $\pi \leq \pi'$, we can identify the associated instanton homology groups via the continuation maps. Finally, for any two perturbations, there is always a larger one, so we can identify them both by passing to that larger one. We will use this same recipe for cobordism maps.

For a cobordism, there is a set $\Pi(W, c) \subset \Pi(Y, w) \times \Pi(Y', w')$ of pairs (π, π') so that $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is pseudo-unobstructed. This set is nonempty by Lemma 3.33. For each such, we

have defined a module map $I^\bullet(W, c)_{\pi \rightarrow \pi'} : I^\bullet(Y, w, \pi) \rightarrow I^\bullet(Y', w', \pi')$ respecting the exact triangles by applying the equivariant instanton homology functors to the map $\tilde{C}(W, c)$ described above, which we have now verified are well-defined up to equivariant homotopy. We abbreviate this to $(W, c)_{\pi \rightarrow \pi'}$ for convenience.

To see that this gives a well-defined map $I^\bullet(Y, w) \rightarrow I^\bullet(Y', w')$, we need to show that for any two pairs $(\pi, \pi') \in \Pi(W, c)$, the maps $(W, c)_{\pi \rightarrow \pi'}$ are related by the relevant continuation maps.

By the index formula, if $(\pi_1, \pi'_1) \in \Pi(W, c)$ and we have $\pi_0 \leq \pi_1$ and $\pi'_1 \leq \pi'_2$, then $(\pi_0, \pi'_2) \in \Pi(W, c)$ as well. By the functoriality mentioned at the beginning of this proof, we have

$$\varphi_{\pi'_1 \rightarrow \pi'_2}(W, c)_{\pi_1 \rightarrow \pi'_1} \varphi_{\pi_0 \rightarrow \pi_1} = (W, c)_{\pi_0 \rightarrow \pi'_2}.$$

Given any (π_0, π'_0) and (π_1, π'_1) in $\Pi(W, c)$, pick a perturbation π_2 with $\sigma_{\pi_2} = \min(\sigma_{\pi_0}, \sigma_{\pi_1})$ and a perturbation π'_2 with $\sigma_{\pi'_2} = \max(\sigma'_{\pi_0}, \sigma'_{\pi_1})$. We have $(\pi_2, \pi'_2) \in \Pi(W, c)$; the procedure above identifies both maps $(W, c)_{\pi_i \rightarrow \pi'_i}$ with $(W, c)_{\pi_2 \rightarrow \pi'_2}$, and hence with each other, using the various continuation maps. It is not much more work to show that the identification induced by any other pair (π_3, π'_3) with $\pi_3 \leq \pi_2$ and $\pi'_2 \leq \pi'_3$ agrees with the identification induced by (π_2, π'_2) .

Altogether, this shows that the cobordism maps given above assemble into a well-defined map between $I^\bullet(Y, w) \rightarrow I^\bullet(Y', w')$. \square

Remark 6.10. In practice, there are two ways to describe the induced map. If one feels free to choose the perturbations π, π' , then one can say the induced map $I^\bullet(W, c)$ is defined by choosing perturbations (π, π') so that (W, c) is pseudo-unobstructed, and defining this to be the map $I^\bullet(Y, w, \pi) \rightarrow I^\bullet(Y', w', \pi')$ given by applying the equivariant homology functor to the construction described earlier in this section.

On the other hand, if one wishes to define the induced map $I^\bullet(W, c)_{\pi_0 \rightarrow \pi'_0} : I^\bullet(Y, w, \pi_0) \rightarrow I^\bullet(Y', w', \pi'_0)$ between perturbations where (W, c) is not necessarily pseudo-unobstructed, instead one should choose new perturbations π_1, π'_1 with $(W, c) : (Y, w, \pi_1) \rightarrow (Y', w', \pi'_1)$ pseudo-unobstructed, and with $\sigma_{\pi_1} \leq \sigma_{\pi_0}$ and $\sigma_{\pi'_0} \leq \sigma_{\pi'_1}$; then the construction above defines

$$I^\bullet(W, c)_{\pi_0 \rightarrow \pi'_0} = \varphi_{\pi'_0 \rightarrow \pi'_1}^{-1} I^\bullet(W, c)_{\pi_1 \rightarrow \pi'_1} \varphi_{\pi_1 \rightarrow \pi_0}^{-1},$$

using that the continuation maps are known to be isomorphisms on equivariant instanton homology by Theorem 5.26. \diamond

6.2 Composites

Theorem 6.1 establishes that each cobordism $(W, c) : (Y, w) \rightarrow (Y', w')$ between admissible pairs induces a well-defined cobordism map. What remains is to establish that these compose functorially. If $(W_1, c_1) : (Y_0, w_0) \rightarrow (Y_1, w_1)$ and $(W_2, c_2) : (Y_1, w_1) \rightarrow (Y_2, w_2)$ are cobordisms, write $(W_{12}, c_{12}) : (Y_0, w_0) \rightarrow (Y_2, w_2)$ for their composite, obtained by pasting the two cobordisms along (Y_1, w_1) .

As mentioned at the beginning of the proof of Theorem 6.1, if π_i are regular perturbations on the respective admissible pairs (Y_i, w_i) for which (W_i, c_i) are pseudo-unobstructed, the composite bimodule

$\mathcal{I}(W_2, c_2) \circ \mathcal{I}(W_1, c_1)$ is homotopic to $\mathcal{I}(W_{12}, c_{12})$. Unfortunately, it is not always possible to choose π_i so that all three of these cobordisms are pseudo-unobstructed, even in favorable-seeming situations. For instance, suppose $b_1(W_i) = b^+(W_i) = 0$ for $i = 1, 2, 12$, and that W_{12} supports a pseudocentral reducible $\Lambda : \theta_0 \rightarrow \theta_2$ which restricts to (Y_1, w_1) as an abelian connection ρ_1 .

In this situation, the normal index is additive; writing Λ_i for the restriction of Λ to W_i , we see that $N(\Lambda_1) + N(\Lambda_2) = N(\Lambda) = -2$. It follows that *no matter what perturbations we choose on the Y_i , one of W_1 or W_2 must be obstructed*: neither Λ_i is a pseudocentral reducible, as it is abelian on one end, but one of them must have negative normal index.

The best we can prove is the following.

Lemma 6.11. *Suppose (W_i, c_i) are composable cobordisms with $b^+(W_i) - b_1(W_i) \geq 0$, and for which either $b^+(W_i) = 0$ or (W_i, c) supports no central connections. Then there exist perturbations π_0, π_1^\pm, π_2 on (Y_i, w_i) so that*

$$\begin{aligned} (W_1, c_1) &: (Y_0, w_0, \pi_0) \rightarrow (Y_1, w_1, \pi_1^+), \\ (W_2, c_2) &: (Y_1, w_1, \pi_1^-) \rightarrow (Y_2, w_2, \pi_2), \\ (W_{12}, c_{12}) &: (Y_0, w_0, \pi_0) \rightarrow (Y_2, w_2, \pi_2) \end{aligned}$$

are all pseudo-unobstructed, and $0 \leq \sigma_{\pi_1^+} - \sigma_{\pi_1^-} \leq 4$.

Proof. This is really a lemma about signature data functions, so we speak in terms of signature data functions σ_i instead of perturbations π_i ; choose one on each (Y_i, w_i) to start with. Write $F_{Y,w} = \text{Map}(\mathfrak{A}(Y, w), 2\mathbb{Z})$. Just as in the proof of Lemma 3.33, we will consider the normal index of a reducible as a function

$$N_i : \mathfrak{A}(W_i, c_i) \times F_{Y_{i-1}, w_{i-1}} \times F_{Y_i, w_i} \rightarrow 2\mathbb{Z}$$

with

$$N_i(\Lambda, f_{i-1}, f_i) = N_i(\Lambda, \sigma_{i-1}, \sigma_i) + \frac{f_i(r_i \Lambda) - f_{i-1}(r_{i-1} \Lambda)}{2}$$

for $i = 1, 2$, and similarly for $i = 12$. Here r_i is restriction to Y_i . For convenience below, ρ refers to an abelian flat connection on a 3-manifold, while θ to a central flat connection.

If $b_1(Y_1) > 0$ then Y_1 is admissible and the cobordisms (W_i, c_i) support no reducible instantons; the relevant unobstructedness conditions hold tautologically for both W_1 and W_2 . The composite W_{12} need not be unobstructed, but Lemma 3.33 shows that one may choose the signature data on the ends appropriately so that it is. A similar discussion applies for Y_0 and Y_2 ; for the remainder of the proof we assume all Y_i are rational homology spheres.

If N_{12} is the corresponding normal index function for the composite cobordism, we have an additivity property (thanks, in part, to the assumption that $b_1(Y_1) = 0$):

$$N_{12}(\Lambda, f_0, f_2) = N_1(\Lambda|_{W_1}, f_0, f_1) + N_2(\Lambda|_{W_2}, f_1, f_2) + c(\Lambda|_{Y_1}),$$

where $c(\alpha) = 0$ if α is abelian and $c(\alpha) = 2$ if α is central. Our goal is to choose f_0, f_1, f_2 so that $N_1 \geq b_1(W_1) - b^+(W_1)$ and $N_{12} \geq b_1(W_{12}) - b^+(W_{12})$ while $N_2 \geq b_1(W_2) - b^+(W_2) - 2$; then the desired signature data will be given by $\sigma_0 + f_0$ and $\sigma_2 + f_2$, while $\sigma_1^+ = \sigma_1 + f_1$ and $\sigma_1^- = \sigma_1 + f_1 - 2$.

To ensure that these inequalities hold for every reducible running $\Lambda_1 : \theta_0 \rightarrow \rho_1$ along the first cobordism and $\Lambda_2 : \rho_1 \rightarrow \theta_2$ along the second, we are obligated to choose f_1 first. In this case, we have

$$N_1(\Lambda_1, f_0, f_1) = N_1(\Lambda_1; \sigma_0, \sigma_1) + \frac{f_1(\rho_1)}{2} \quad \text{and} \quad N_2(\Lambda_2, f_1, f_2) = N_2(\Lambda_2; \sigma_1, \sigma_2) - \frac{f_1(\rho_1)}{2}.$$

The desired inequalities on N_1 and N_2 are equivalent to

$$2N_2(\Lambda_2; \sigma_1, \sigma_2) - 2(b_1(W_2) - b^+(W_2)) + 4 \geq f_1(\rho_1) \geq 2(b_1(W_1) - b^+(W_1)) - 2N_1(\Lambda_1; \sigma_0, \sigma_1)$$

for all abelian flat connections ρ_1 and all reducible instantons Λ_i which are asymptotic to ρ_1 at the appropriate end and a central connection on the other end.

Such a function exists if and only if

$$2N_1(\Lambda_1; \sigma_0, \sigma_1) + 2N_2(\Lambda_2; \sigma_1, \sigma_2) - 2(b_1(W_1) + b_1(W_2) - b^+(W_1) - b^+(W_2)) + 4 \geq 0$$

for all $\Lambda_1 : \theta_0 \rightarrow \rho_1$ and $\Lambda_2 : \rho_1 \rightarrow \theta_2$.

By additivity of index and the fact that Y_1 is a rational homology sphere, this is equivalent to the inequality

$$N_{12}(\Lambda_{12}; \sigma_0, \sigma_2) \geq b_1(W_{12}) - b^+(W_{12}) - 2.$$

Finally, the index formula Proposition 3.24 implies $N_{12} \geq 2b_1 - 2b^+ - 2$, and the assumption that $b_1 - b^+ \geq 0$ implies that this is at least $b_1 - b^+ - 2$. Therefore one may choose f_1 so that the given inequality holds. (If $b_1 - b^+ \geq 2$, it follows that in fact choose π_1, π_2, π_3 so that $N_i(\Lambda_i) \geq 0$ for all i .)

This establishes the desired inequality for all abelian instantons $\Lambda_1 : \theta_0 \rightarrow \rho_1$ and $\Lambda_2 : \rho_1 \rightarrow \theta_2$ with this choice of f_1 ; this choice is irrelevant to W_{12} .

This is the only difficult point; the rest of the argument follows as in Lemma 3.33. In fact, with these choices of signature data, the only obstructed non-pseudocentral reducibles Λ_2 on (W_2, c_2) run from an abelian connection to a central connection. \square

Because the situation in the preceding Lemma is quite subtle, so too is proving functoriality of the induced maps. As we will see in the proof of the following Theorem, establishing this requires slightly more general versions of our previous constructions. Instead of carrying out the necessary mild generalizations of previous constructions in detail we content ourselves to a sketch of the proof of functoriality.

Theorem 6.12. *Suppose $(W_1, c_1) : (Y_0, w_0) \rightarrow (Y_1, w_1)$ and $(W_2, c_2) : (Y_1, w_1) \rightarrow (Y_2, c_2)$ are cobordisms between admissible pairs which have either $b^+(W_i) = 0$ or that (W_i, c_i) supports no central connections. Then the same is true for the composite (W_{12}, c_{12}) , and the cobordism maps defined in Theorem 6.1 satisfy*

$$I^\bullet(W_2, c_2) \circ I^\bullet(W_1, c_1) = I^\bullet(W_{12}, c_{12}).$$

Sketch of proof. If (W_{12}, c_{12}) supports central connections, then by restriction so too do both (W_i, c_i) (and neither end is admissible, so all boundary components are rational homology spheres); in this case,

we must have $b^+(W_1) = b^+(W_2) = 0$ by assumption, but because Y_1 is a rational homology sphere this implies $b^+(W_{12}) = 0$ as well, proving the first claim about the composite.

As for the cobordism maps, to take advantage of Lemma 6.11 we want to reduce to the case $b_1(W_i) - b^+(W_i) \geq 0$ for both $i = 1, 2$. First we handle the case where at least one of these is negative. By an argument along the lines of Lemma 6.11, we may choose perturbations π_0, π_1, π_2 so that one of (W_i, c_i) is a pseudo-unobstructed cobordism and the other — while not pseudo-unobstructed — has $b_1(W_i) - b^+(W_i) < 0$. Because $b^+(W_i) < 0$, this implies that there exists a regular perturbation π_i^W on this cobordism with no reducible connections whatsoever; the same is necessarily true for $\pi_1^W \cup_L \pi_2^W$ when the length of the neck L is sufficiently large. The stated result follows in this case by a standard neck-stretching argument which implies that the induced map on the composite is the same as the map induced by $\pi_1^W \cup_L \pi_2^W$ for sufficiently large L .

So we assume $b_1(W_i) - b^+(W_i) \geq 0$ for both $i = 1, 2$, and suppose that we have chosen perturbations π_0, π_1^\pm, π_2 as in the statement of Lemma 6.11.

Now the situation is much more complicated. If we can establish the existence of the diagram below, as well as show that it is homotopy-commutative, we will have proved the desired theorem. In this diagram, we write for instance $\tilde{C}(\pi_0)$ in place of the full notation $\tilde{C}(Y_0, w_0, \pi_0; R)$ for the sake of space, and similarly we suppress the geometric representatives c . On the other hand, we include the perturbations on the ends as part of the notation for clarity.

$$\begin{array}{ccccc}
 & & \tilde{C}(\pi_1^+) & & \\
 & \nearrow^{W_1^{0 \rightarrow 1^+}} & \downarrow \widetilde{\text{Cyl}}^{1^+ \rightarrow 1^-} & & \\
 \tilde{C}(\pi_0) & \xrightarrow{\widetilde{W}_1^{0 \rightarrow 1^-}} & S\tilde{C}(\pi_1^-) & \xrightarrow{(SW_2)^{1^- \rightarrow 2}} & S\tilde{C}(\pi_2) \\
 & & \uparrow S_1 & & \uparrow S_2 \\
 & & \tilde{C}(\pi_1^-) & \xrightarrow{W_2^{1^- \rightarrow 2}} & \tilde{C}(\pi_2) \\
 & \searrow_{W_{12}^{0 \rightarrow 2}} & & &
 \end{array}$$

A handful of notational points need to be explained here. The existence of a handful of these $SO(3)$ -equivariant chain complexes and chain maps are slight generalizations of work carried out earlier in this article.

- Here $S\tilde{C}(\pi_1^-)$ and $S\tilde{C}(\pi_2)$ are defined by a variation on the blowup construction of Section 5.3, where we blow up simultaneously at *every* abelian reducible. This variation can be constructed by hand or by iteratively applying the version of the construction where we blow up at a single abelian reducible.
- The maps S_1, S_2 are defined by a variation on the map defined in Construction 5.19, with the difference being that $S_i(\rho, S_\rho) = S_\rho$ for *every* abelian reducible ρ . Again, these induce an

isomorphism on homology, and therefore equivariant homology.

- The maps $\widetilde{\text{Cyl}}^{1^+ \rightarrow 1^-}$ and $\widetilde{W}_1^{0 \rightarrow 1^-}$ are defined by a variation on the construction of Section 5.4 of the bimodule of an obstructed cobordism, applied to the cylinder $\mathbb{R} \times Y$ and the cobordism W_1 ; this variation allows for more than one obstructed abelian connection Λ , which may run from a central connection to an abelian connection. A careful reading shows that the construction and proofs of facts regarding these maps go through with minimal change; the only crucial points are that all obstructed abelian connections run $\alpha \rightarrow \rho$ where ρ is abelian (but α may be abelian or central), and each has normal index -2 .
- The bimodule inducing the map $(SW_2)^{1^- \rightarrow 2}$ is defined by applying Construction 5.21 to the bimodule $S_2 \circ W_2^{1^- \rightarrow 2}$; it can also be constructed by hand with little difficulty.

Now we assert that this diagram commutes up to homotopy. In particular, we assert the existence of equivariant homotopies between the following pairs of maps:

$$\begin{aligned} \widetilde{\text{Cyl}}^{1^+ \rightarrow 1^-} \circ \text{Cyl}^{1^- \rightarrow 1^+} &\sim S_1 \\ \widetilde{\text{Cyl}}^{1^+ \rightarrow 1^-} \circ W_1^{0 \rightarrow 1^+} &\sim \widetilde{W}_1^{0 \rightarrow 1^-} \\ (SW_2)^{1^- \rightarrow 2} \circ S_1 &\sim S_2 \circ W_2^{1^- \rightarrow 2} \\ (SW_2)^{1^- \rightarrow 2} \circ \widetilde{W}_1^{0 \rightarrow 1^-} &\sim S_2 \circ W_{12}^{0 \rightarrow 2}. \end{aligned}$$

The first statement was established in the proof of Theorem 5.26; here $\text{Cyl}^{1^- \rightarrow 1^+} : \tilde{\mathcal{C}}(\pi_1^-) \rightarrow \tilde{\mathcal{C}}(\pi_1^+)$ is the continuation map induced by the unobstructed cylinder.

The second statement follows by yet more obstructed gluing theory in the spirit of Section 5, now applied to the obstructed cobordism given by the composite of W_1 and the obstructed cylinder $\mathbb{R} \times Y_1 : (Y_1, \pi_1^+) \rightarrow (Y_1, \pi_1^-)$; the modified moduli space ideas apply with minimal change.

The third statement is tautological: these maps are in fact identically equal.

The final statement is more difficult, but is rather similar to the arguments carried out in Section 5.5. Again, this amounts to using obstructed gluing theory to construct the appropriate modified moduli spaces, and recognizing that their boundary relations establish a bimodule homotopy between the two composite bimodules represented by this relation.

Now recall our goal. Abusing notation, we write W also for the induced map on equivariant homology I^\bullet ; for consistency with previous sections we write $\varphi_{1^- \rightarrow 1^+}$ as the map induced by the unobstructed continuation map and $\tilde{\varphi}_{1^+ \rightarrow 1^-}$ for the map induced by the obstructed continuation map. As discussed in Remark 6.10, the composite

$$I^\bullet(W_2, c_2) \circ I^\bullet(W_1, c_1) : I^\bullet(\pi_0) \rightarrow I^\bullet(\pi_2)$$

is defined to be the composite

$$W_2^{1^- \rightarrow 2} \circ \varphi_{1^- \rightarrow 1^+}^{-1} \circ W_1^{0 \rightarrow 1^+},$$

and we want to show that this is equal to the map $W_{12}^{0 \rightarrow 2}$.

This follows purely formally using the identities above. Applying each of these in turn (occasionally rearranging using that S_i and φ are isomorphisms on equivariant homology), we find the desired equality:

$$\begin{aligned}
W_2^{1^- \rightarrow 2} \varphi_{1^- \rightarrow 1^+}^{-1} W_1^{0 \rightarrow 1^+} &= W_2^{1^- \rightarrow 2} S_1^{-1} \tilde{\varphi}_{1^+ \rightarrow 1^-} W_1^{0 \rightarrow 1^+} \\
&= W_2^{1^- \rightarrow 2} S_1^{-1} \tilde{W}_1^{0 \rightarrow 1^-} \\
&= S_2^{-1} (S W_2)^{1^- \rightarrow 2} \tilde{W}_1^{0 \rightarrow 1^-} \\
&= S_2^{-1} S_2 W_{12}^{0 \rightarrow 2} = W_{12}^{0 \rightarrow 2}. \quad \square
\end{aligned}$$

Altogether, we have proven Theorem 2.12(i): I^\bullet is functorial on a category whose objects are admissible pairs (Y, w) with a basepoint on Y , and whose morphisms are cobordisms (W, c) equipped with a homology orientation and a path between the basepoints, where either $b^+(W) = 0$ or (W, c) admits no central connections.

6.3 The case $b_2^+ > 0$ via the blowup trick

Thusfar we have established functoriality for negative-definite cobordisms and cobordisms which support no central connections. Our techniques do not apply to the case $b_2^+ > 0$ without change; now we have an obstructed *central* connection on the cobordism, while we have explained only how to resolve (mildly) obstructed abelian connections.

It is likely that a modification of the previous techniques (a blown-up flow category construction for central critical orbits instead) will allow us to extend functoriality to this setting, and this will be necessary for an investigation of the \mathbb{F}_2 -theory. We leave this to future work, instead opting to use a trick which requires inverting 2 in the coefficient ring. This is founded on an idea appearing in [MM92]: by taking the connected sum with a bundle over $\overline{\mathbb{C}\mathbb{P}^2}$ with odd first Chern class, we can remove central connections and thus define induced maps associated to all cobordisms.

We *very* briefly recall the relevant definitions and results. See [DK90, Section 5.2] for a careful discussion of the setup.

Let (W, c) be a cobordism, and $\Sigma \subset W$ a closed oriented surface embedded in W ; arbitrarily choose a spin structure on this surface. The index bundle of the dual Dirac operator $\tilde{\mathcal{D}}_{A, \Sigma}^*$ (paired to the restriction of $A \in \tilde{\mathcal{B}}(W, c)$ to Σ) induces an $SO(3)$ -equivariant complex line bundle \tilde{L}_Σ over the configuration space of framed connections on W . Given a reducible orbit $\mathcal{O}_\Lambda \subset \tilde{\mathcal{B}}(W, c)$, the stabilizer Γ_Λ acts on the fiber of \tilde{L}_Σ over Λ . If Λ is central, this action is necessarily trivial, as every homomorphism $SO(3) \rightarrow SO(2)$ is trivial. If Λ is abelian, respecting a parallel splitting $E_c \cong L_1 \oplus L_2$ into complex line bundles with first Chern classes $\{x, y\}$, the weight of the $SO(2) \cong \Gamma_\Lambda \subset SO(3)$ action is given by $\langle x - y, [\Sigma] \rangle$. The first Chern class of the associated complex line bundle over \mathcal{O}_Λ is $-2\langle x - y, [\Sigma] \rangle$.

In favorable situations (where there are few reducibles, or at least those reducibles that do exist appear with positive normal index) one may choose a generic $SO(3)$ -equivariant section of \tilde{L}_Σ which is

transverse to zero on all moduli spaces $M_z^+(W, c; \alpha, \alpha')$. In this situation, the cut-down moduli spaces again define a bimodule $\mathcal{I}(W, c, \Sigma)$ of relative degree two smaller than $\mathcal{I}(W, c)$.

We will apply this construction in a special case.

Construction 6.13. Suppose (W, c) is a cobordism equipped with a base-path. Choose a point $p \in W$ disjoint from this base-path and the surface c . Take the blowup at p , producing the connected-sum $(W_{\#}, c_{\#}) = (W_{\#} \# \overline{\mathbb{C}\mathbb{P}^2}, c \cup \overline{\mathbb{C}\mathbb{P}^1})$. If z is a component of connections on W , we write $z_{\#}$ for the corresponding component on $W_{\#}$ arising as connected sum with the component of connections on $(\overline{\mathbb{C}\mathbb{P}^2}, \overline{\mathbb{C}\mathbb{P}^1})$ of index -1 (framed index 2), so that $\tilde{i}(z_{\#}) = \tilde{i}(z) + 2$. We will be investigating the cut-down moduli spaces associated to the surface $\Sigma = \overline{\mathbb{C}\mathbb{P}^1}$.

The reducibles on $(W_{\#}, c_{\#})$ are labeled by pairs $\{x' = (x, i), y' = (y, j)\}$, where $x, y \in H^2(W; \mathbb{Z})$ and i, j are integers such that $x + y = PD(c)$ and $i + j = 1$. Because it is never possible that $x' = y'$ (as $i + j$ is odd), this cobordism supports no central reducibles; if (W, c) supports a central reducible, then the corresponding reducible of least index on $(W_{\#}, c_{\#})$ is pseudo-central. If Λ is an abelian reducible on W with cohomology classes $\{x, y\}$, the reducible Λ_i on $(W_{\#}, c_{\#})$ with cohomology classes $\{(x, i), (y, 1 - i)\}$ has normal index $N(\Lambda_i) = N(\Lambda) + 2(2i - 1)^2 \geq N(\Lambda) + 2$.

It follows that if (W, c) is pseudo-unobstructed, then $(W_{\#}, c_{\#})$ is unobstructed. When $b^+(W) > 0$, it follows that for a generic choice of perturbation $(W_{\#}, c_{\#})$ supports no reducibles whatsoever. Choose such a perturbation, and then a section of \tilde{L}_{Σ} so that the moduli spaces $M_{z_{\#}}^+(W_{\#}, c_{\#}; \alpha, \alpha')$ are always transverse to zero. We write $M_{z_{\#}}^+(W_{\#}, c_{\#}, \overline{\mathbb{C}\mathbb{P}^1}; \alpha, \alpha')$ for the moduli spaces cut-down by these sections. These define a bimodule $\mathcal{I}(W_{\#}, c_{\#}, \pi_{\#})$. This bimodule is transparently well-defined up to equivariant homotopy if $b^+(W) > 1$.

Well-definedness in the case $b^+(W) = 1$ is somewhat more subtle, but entirely analogous to Proposition 6.8; if Θ is a central connection on (W, c) , then in 1-parameter families of perturbations on $(W_{\#}, c_{\#})$, the abelian connection Θ_0 may appear as an isolated obstructed reducible of normal index -2 on $(W_{\#}, c_{\#})$. Over such a reducible there is no transverse section of $\tilde{L}_{\Sigma} \rightarrow \mathcal{O}_{\Theta_0}$; these play the role that the obstructed reducibles of normal index -4 played in Proposition 6.8, though the sign on the zero locus is opposite, as the associated complex line bundle has first Chern class -2 instead of $+2$.

The same modified moduli space construction from Proposition 6.8 implies that the resulting bimodule $\mathcal{I}(W_{\#}, c_{\#}, [\overline{\mathbb{C}\mathbb{P}^1}], \pi_{\#})$ is independent of $\pi_{\#}$ up to equivariant homotopy.

We write $I^{\bullet}(W_{\#}, c_{\#}, [\overline{\mathbb{C}\mathbb{P}^1}])$ for the induced map on equivariant instanton homology. ◇

Before moving on to discussing functoriality, we should compare this cobordism map to the ones we've already defined.

Proposition 6.14. *If (W, c) is a pseudo-unobstructed cobordism with $b^+(W) > 0$, and 2 is invertible in R , then $-2I^{\bullet}(W, c) = I^{\bullet}(W_{\#}, c_{\#}, [\overline{\mathbb{C}\mathbb{P}^1}])$.*

Sketch of proof. Choose a perturbation on (W, c) with no reducibles; because (W, c) is pseudo-unobstructed with $b^+(W) > 0$, it supports no central connections. Write (W_L, c_L) for the connected sum $(W, c) \# (\overline{\mathbb{C}\mathbb{P}^2}, \overline{\mathbb{C}\mathbb{P}^1})$

above with a neck of length L . We may perform this connected sum in such a way that when $L = \infty$, the resulting space is conformally equivalent to $W \setminus p \sqcup \overline{\mathbb{C}\mathbb{P}^2} \setminus \{q\}$, and therefore finite-energy on (W_∞, c_∞) instantons can be identified with pairs of instantons on (W, c) and instantons on $(\overline{\mathbb{C}\mathbb{P}^2}, \overline{\mathbb{C}\mathbb{P}^1})$. The framed moduli space on $(\overline{\mathbb{C}\mathbb{P}^2}, \overline{\mathbb{C}\mathbb{P}^1})$ has dimension 2 modulo 8, so the only relevant moduli space is the 2-dimensional one, which consists of a single abelian instanton.

It follows that

$$M_{z\#}^+(W_\infty, c_\infty; \alpha, \alpha') = M_z^+(W, c; \alpha, \alpha') \times S^2.$$

Furthermore, the complex line bundle \tilde{L}_Σ is determined entirely by the $\overline{\mathbb{C}\mathbb{P}^2}$ factor, so is pulled back from (a negatively oriented) TS^2 over the latter factor. *If we ignore equivariance*, we could choose our section to be a section $S^2 \rightarrow TS^2$ with exactly two negatively-oriented zeroes; this would identify $M_{z\#}^+(W_\infty, c_\infty, \overline{\mathbb{C}\mathbb{P}^1}; \alpha, \alpha')$ with -2 copies of $M_z^+(W, c; \alpha, \alpha')$. This observation, first appearing (with different orientation conventions) as [Kot91, Theorem 6.9] and also appearing as [MM92, Proposition 3.2], says that we have not produced anything new.

Unfortunately, this does not lead to an identification of the bimodules $-2\mathcal{I}(W, c)$ and $\mathcal{I}(W_\#, c_\#, [\overline{\mathbb{C}\mathbb{P}^1}])$. The issue is that it is not at all obvious that one may choose an equivariant section of \tilde{L}_Σ whose zero set is oriented diffeomorphic to $-2M_z^+(W, c; \alpha, \alpha')$ and for which this diffeomorphism preserves (or at least nearly preserves) the endpoint maps. This is not an issue for the authors above; instead of ‘cutting down’ geometrically by taking a zero set of this complex line bundle, these authors take the cap product against a degree-2 cohomology class, where one does not have these equivariance concerns.

To solve this issue, we can pass to the finite-dimensional ‘Donaldson model’ of [Mil19, Section 7.2], which is defined under the hypothesis that $2 \in R$ is invertible. This finite-dimensional model is obtained by first replacing every irreducible orbit α in $\tilde{C}(Y, w)$ with a copy $\tilde{\alpha}$ of S^3 , and replacing the moduli spaces running between them with appropriate double or four-fold covers \tilde{M} , weighted by $1/2$ or $1/4$ as appropriate. The chain maps $\tilde{C}(W, c)$ are defined by a matrix, whose entries count intersections of the endpoint maps $\tilde{M}_z^+(W, c; \tilde{\alpha}, \tilde{\alpha}') \rightarrow \tilde{\alpha} \times \tilde{\alpha}'$ with one of the four submanifolds $p_\alpha \times q_{\alpha'}$ or $\tilde{\alpha} \times q_{\alpha'}$ or $p_\alpha \times \tilde{\alpha}'$ or $\tilde{\alpha} \times \tilde{\alpha}'$.

In this model, where the differential is defined by point-counts, the point-counts of the intersections of $\tilde{M}_{z\#}^+(W_\infty, c_\infty, \overline{\mathbb{C}\mathbb{P}^1}; \tilde{\alpha}, \tilde{\alpha}')$ with these submanifolds are precisely equal to -2 times the corresponding point counts of $\tilde{M}_z^+(W, c; \tilde{\alpha}, \tilde{\alpha}')$. To be precise, if $\tilde{M}_z^+(W, c; \tilde{\alpha}, \tilde{\alpha}') \rightarrow \tilde{\alpha} \times \tilde{\alpha}'$ has intersection number k with $S \subset \tilde{\alpha} \times \tilde{\alpha}'$, then $\tilde{M}_{z\#}^+(W_\#, c_\#; \tilde{\alpha}, \tilde{\alpha}') \cap S$ is identified with k copies of S^2 (counting with sign). Taking the zero set of the transverse section of \tilde{L}_Σ gives rise to $-2k$ points.

It follows that $-2I^\bullet(W, c) = I^\bullet(W_\#, c_\#, [\overline{\mathbb{C}\mathbb{P}^1}])$. □

Remark 6.15. The statement above extends to the case $b^+(W) = 0$, but one needs to explain more clearly how to define the right-hand-side in the case that $b^+ = 0$ (using a setup analogous to that of Section 6.1); we do not see substantial benefit to doing so.

Further, this statement should hold with respect to all coefficient rings, not simply those in which 2 is invertible. (In particular, over $R = \mathbb{F}_2$, the induced map $I^\bullet(W_\#, c_\#, [\overline{\mathbb{C}\mathbb{P}^1}])$). One proof would define a

Donaldson model without passing to a double cover of $SO(3)$, using a more intricate CW decomposition of $SO(3)$. This is not difficult, but we will not carry it out here.

Instead of passing to the Donaldson model, another way to resolve the issue is to quotient by a larger equivalence relation in the definition of the geometric chain complex, so that the zero set of an equivariant transverse section over $M \times S^2$ is identified with 2 copies of M . \diamond

This in hand, we can prove Theorem 2.12(ii) from the guide:

Theorem 6.16. *When R is a PID with $1/2 \in R$, then the functors $I^\bullet(Y, w; R)$ extend to the category Cob_{3+1}^w of weakly admissible pairs, allowing for indefinite cobordisms.*

Proof. Let $(W, c) : (Y, w, \pi) \rightarrow (Y, w, \pi')$ be a cobordism (with the necessary additional data). We define

$$I^\bullet(W, c) = \frac{1}{2} I^\bullet_{\#}(W, c).$$

We have already verified in Proposition 6.14 that this agrees with the definition of $I^\bullet(W, c)$ from Theorem 6.1. That this map is functorial follows from the following chain of equalities. Write $W^{12} = W^1 \cup_Y W^2$, where $W^1 : Y_1 \rightarrow Y$ and $W^2 : Y \rightarrow Y_2$ are cobordisms (with suitable additional data, which is suppressed from notation for convenience). Then

$$\begin{aligned} \frac{1}{2} I^\bullet(W_{\#}^{12}, \mathbf{E}_{\#}^{12}, \overline{\mathbb{C}\mathbb{P}}^1; R) &= \frac{1}{4} I^\bullet(W_{\#}^{12}, (\mathbf{E}_{\#}^{12})_{\#}, \overline{\mathbb{C}\mathbb{P}}_1^1 \sqcup \overline{\mathbb{C}\mathbb{P}}_2^1; R) \\ &= \frac{1}{4} I^\bullet(W_{\#}^1, \mathbf{E}_{\#}^1, \overline{\mathbb{C}\mathbb{P}}_1^1) \circ I^\bullet(W_{\#}^2, \mathbf{E}_{\#}^2, \overline{\mathbb{C}\mathbb{P}}_2^1) \\ &= \left(\frac{1}{2} I^\bullet(W_{\#}^1, \mathbf{E}_{\#}^1, \overline{\mathbb{C}\mathbb{P}}_1^1) \right) \circ \left(\frac{1}{2} I^\bullet(W_{\#}^2, \mathbf{E}_{\#}^2, \overline{\mathbb{C}\mathbb{P}}_2^1) \right). \end{aligned}$$

Here the subscripts on $\overline{\mathbb{C}\mathbb{P}}^1$ indicate which of the two blowups they correspond to. The second-to-last equality follows from the fact that instanton homology is functorial for unobstructed cobordisms by the usual neck-stretching argument. \square

7 Irreducible Floer homology and the Casson-Walker invariant

In Floer's original definition of the instanton homology of integer homology spheres, the invariant $I_*(Y)$ was obtained by ignoring the central connection and constructing a chain complex generated by the irreducible connections.

One may define such a chain complex for rational homology spheres, but it is no longer independent of the choice of perturbation. Instead, we will see it depends on a choice of signature data function σ , and use the results of Section 5.5 to see how it changes as we pass between adjacent signature data.

7.1 Floer's complex from algebra

Given an $SO(3)$ -equivariant flow category \mathcal{C} as in Section 4.2, we can extract a chain complex called the *irreducible flow complex*. By default, it is relatively graded; if \mathcal{C} is given an absolute grading, then the irreducible flow complex inherits an absolute grading as well.

We say an object α of \mathcal{C} is *irreducible* if it is diffeomorphic to $SO(3)$. We define a complex

$$C_*^{\text{irr}}(\mathcal{C}) = \bigoplus_{\substack{\alpha \in \text{Ob}(\mathcal{C}) \\ \alpha \text{ irreducible}}} \mathbb{Z}[\tilde{i}(\alpha)] \times_{\mathbb{Z}/2} \Lambda(\alpha),$$

where as in Section 4.2 the two-element set $\Lambda(\alpha)$ corresponds to orientation data, while $\tilde{i}(\alpha)$ is the absolute grading if one exists, but otherwise is meant to indicate that the relative grading between the α and β summands is $\tilde{i}(\alpha, \beta)$.

When $\tilde{i}(\alpha, \beta) = 1$, the moduli space $\mathcal{C}(\alpha, \beta)$ has dimension $3 = \dim \alpha$, and the map $\mathcal{C}(\alpha, \beta) \rightarrow \alpha$ is an $SO(3)$ -equivariant local diffeomorphism. The fiber above any point in α is canonically identified with any other, and we write $\mathcal{C}^{\text{irr}}(\alpha, \beta)$ for this fiber. When one chooses an element of $\Lambda(\alpha)$ and an element of $\Lambda(\beta)$, this fiber is given an orientation, which reverses upon changing either of those choices. To suppress discussion of these orientation elements, let us suppose we have prechosen such elements for each irreducible orbit.

We may thus define a differential on $C_*^{\text{irr}}(\mathcal{C})$ via

$$d^{\text{irr}}(\alpha) = \sum_{\beta} \#\mathcal{C}^{\text{irr}}(\alpha, \beta)\beta.$$

Instead of verifying that this squares to zero (which is straightforward), it is useful to identify the construction above with a purely algebraic avatar. Write $u = [SO(3)] \in C_3SO(3)$ for the representative of the fundamental class. Consider the image of $u : CM_*(\mathcal{C}) \rightarrow CM_{*+3}(\mathcal{C})$, which we write as $uCM_*(\mathcal{C})$. This is a chain complex with the differential \tilde{d} , because the Leibniz rule gives $\tilde{d}(ux) = (du)x - u\tilde{d}(x) = u(-\tilde{d}(x))$ as $d(u) = 0$. Next, we will show that this complex is closely related to $C_*^{\text{irr}}(\mathcal{C})$.

Proposition 7.1. *The complex $(C_*^{\text{irr}}(\mathcal{C}), d^{\text{irr}})$ is isomorphic to the complex $(uCM_*(\mathcal{C})[-3], -\tilde{d})$.*

Proof. First, let us recall some structural properties of the complex $CM_*(\mathcal{C})$, and in particular of geometric chains. This complex has underlying relatively graded abelian group given by

$$CM_*(\mathcal{C}) = \bigoplus_{\alpha \in \text{Ob}(\mathcal{C})} C_*(\alpha)[\tilde{i}(\alpha)] \otimes_{\mathbb{Z}/2} \mathbb{Z}[\Lambda(\alpha)].$$

Because we use the truncated geometric chain complex defined in Appendix A.5, and this complex has $C_*(\alpha)$ supported in degrees $[0, \dim \alpha]$, we see that if α is diffeomorphic to a point or S^2 , we must have $uC_*(\alpha) = 0$, while if $\alpha \cong SO(3)$, we have a noncanonical isomorphism $uC_*(\alpha) \cong \mathbb{Z}[3]$, with the fundamental class with respect to some orientation $[\alpha]$ being a cyclic generator.

By choosing the class u at the start we fixed an orientation on $SO(3)$; orient each irreducible α so that the map $SO(3) \rightarrow \alpha$ given by acting on a point is an orientation-preserving diffeomorphism. This makes the above isomorphism canonical, taking the cyclic generator to be the fundamental class with respect to the orientation described above; equivalently, by sending $u \cdot \text{pt}$ to $1 \in \mathbb{Z}$.

Thus, at the level of graded abelian groups, we have an identification

$$C_*^{\text{irr}}(\mathcal{C}) = uCM_*(\mathcal{C})[-3]$$

sending α to $[\alpha]$. We will see in fact their differentials coincide, so long as one interprets the differential on the right-hand-side as $-\tilde{d}$. To see this, observe that

$$\tilde{d}[\alpha] = \sum_{\beta} (-1)^{\dim \alpha} [\alpha] \times_{\alpha} \mathcal{C}(\alpha, \beta) = \sum_{\beta \in \text{Ob}^{\text{irr}}(\mathcal{C})} (-1)^{\dim \alpha} \mathcal{C}(\alpha, \beta);$$

if $\tilde{i}(\alpha, \beta) > 1$, the space $\mathcal{C}(\alpha, \beta)$ has dimension at least 4, so is zero in $C_*(\beta)$. Using also that $\dim \alpha = 3$, we may thus write this sum as

$$\tilde{d}[\alpha] = \sum_{\substack{\beta \in \text{Ob}^{\text{irr}}(\mathcal{C}) \\ \tilde{i}(\alpha, \beta) = 1}} -\mathcal{C}(\alpha, \beta).$$

Now each $\mathcal{C}(\alpha, \beta)$ here is three-dimensional, so each component projects diffeomorphically onto β ; these components are in one-to-one bijection with $\mathcal{C}^{\text{irr}}(\alpha, \beta)$, the points in a fiber above some $x \in \alpha$.

To compare orientations, recall that $\mathcal{C}^{\text{irr}}(\alpha, \beta)$ is oriented as a fiber over $x \in \alpha$, so that

$$\mathcal{C}(\alpha, \beta) \cong SO(3) \times \mathcal{C}^{\text{irr}}(\alpha, \beta)$$

is an oriented diffeomorphism, and the endpoint map $SO(3) \rightarrow \beta$ is equivariant, hence an oriented diffeomorphism by our initial choice of orientations. It follows that

$$\tilde{d}[\alpha] = - \sum_{\beta} \# \mathcal{C}^{\text{irr}}(\alpha, \beta)[\beta],$$

as desired. □

Now observe that if $\tilde{\mathcal{C}}$ is a dg $C_*SO(3)$ -module, then $u\tilde{\mathcal{C}}$ is again a chain complex; if $f : \tilde{\mathcal{C}} \rightarrow \tilde{\mathcal{C}}'$ is a $C_*SO(3)$ -equivariant map, then because $f(ux) = uf(x)$ we see that f restricts to a chain map $u\tilde{\mathcal{C}} \rightarrow u\tilde{\mathcal{C}}'$. Finally, if $H : \tilde{\mathcal{C}} \rightarrow \tilde{\mathcal{C}}'[1]$ is a $C_*SO(3)$ -equivariant chain homotopy, this induces a chain homotopy $u\tilde{\mathcal{C}} \rightarrow u\tilde{\mathcal{C}}'[1]$.

Corollary 7.2. *Given a weakly admissible pair with choice of signature data (Y, w, σ) , there is a well-defined ‘irreducible instanton homology group’ $I_*(Y, w, \sigma)$. If $(W, c) : (Y, w, \sigma) \rightarrow (Y', w', \sigma')$ is nearly good, then there is a well-defined induced map*

$$I(W, c) : I_*(Y, w, \sigma) \rightarrow I_*(Y', w', \sigma'),$$

which composes functorially.

Proof. As a corollary of Theorem 4.21, whenever π_0 and π_1 are two regular perturbations on (Y, w) with the same signature data, there is a $C_*SO(3)$ -equivariant continuation map $\tilde{C}(Y, w, \pi_0) \rightarrow \tilde{C}(Y, w, \pi_1)$ which is an $SO(3)$ -equivariant homotopy equivalence. Passing to the homology of $u\tilde{C}[-3]$, one may construct a canonically defined $I_*(Y, w, \sigma)$ which depends only on σ by the mechanism of Definition 5.28.

When $(W, c) : (Y, w, \pi) \rightarrow (Y', w', \pi')$ is a nearly good cobordism whose only obstructed reducibles are bad reducibles, we showed in Section 6.1 that there is a corresponding $SO(3)$ -equivariant chain map $\tilde{C}(Y, w, \pi) \rightarrow \tilde{C}(Y', w', \pi')$ which is well-defined up to equivariant chain homotopy and which composes functorially up to equivariant chain homotopy. Applying $u\tilde{C}[-3]$ and passing to homology, we get the claimed functoriality results. \square

While the groups $I_*(Y, w, \sigma)$ do depend on the signature data σ , we can take advantage of the proof of Theorem 5.26. If $\sigma_0 < \sigma_1$ are adjacent signature data on (Y, w) which differ on $\rho \in \mathfrak{A}(Y, w)$, we proved that theorem by establishing the existence of an $SO(3)$ -equivariant homotopy equivalence

$$CM_*\mathcal{B}_\rho\mathcal{I}(Y, w, \sigma_0) \simeq CM_*\mathcal{I}(Y, w, \sigma_1).$$

Writing $S_\rho\tilde{C}(Y, w, \sigma_0)$ for the former complex and $\tilde{C}(Y, w, \sigma_1)$ for the latter, we thus see that

$$I_*(Y, w, \sigma_1) \cong u\tilde{C}(Y, w, \sigma_1)[-3] \cong uS_\rho\tilde{C}(Y, w, \sigma_0)[-3].$$

To understand how I_* changes in passing from σ_0 to σ_1 , we merely need to understand the latter complex.

Proposition 7.3. *Let \mathcal{C} be an $SO(3)$ -equivariant flow category of the type considered above (for instance, any framed instanton flow category). The Morse complex $C_*^{\text{irr}}(\mathcal{B}_\rho\mathcal{C})$ is obtained as an algebraic mapping cone. Precisely, there is an isomorphism of (relatively graded) abelian groups*

$$uCM_*(\mathcal{B}_\rho\mathcal{C})[-3] = uCM_*(\mathcal{C})[-3] \oplus \mathbb{Z}[\tilde{i}_{\mathcal{C}}(\rho) - 1]$$

so that

$$d_\rho^{\text{irr}} = \begin{pmatrix} d^{\text{irr}} & 0 \\ -V_2 & 0 \end{pmatrix},$$

where $\langle V_2(\rho), \beta \rangle$ is the signed count of isolated irreducible components of $\mathcal{C}(\rho, \beta)$.

Proof. We have

$$uCM_*(\mathcal{B}_\rho\mathcal{C}) = \left(\bigoplus_{\alpha \in \text{Ob}^{\text{irr}}(\mathcal{C})} \mathbb{Z}\langle \alpha \rangle[\tilde{i}(\alpha)] \right) \oplus \mathbb{Z}\langle S_\rho \rangle[\tilde{i}(\rho) - 1];$$

we obtain the additional generator from $u\mathcal{C}_*(S_\rho)[-3]$, as $\tilde{i}_{\mathcal{B}}(S_\rho) = \tilde{i}_{\mathcal{C}}(\rho) - 1$. This gives the stated isomorphism at the level of relatively graded abelian groups.

Now it remains to compute the differential, which we may write as

$$d^{\text{irr}} = \begin{pmatrix} d_{CC} & d_{C\rho} \\ d_{\rho C} & d_{\rho\rho} \end{pmatrix}.$$

Each component is given by computing $\#\mathcal{B}^{\text{irr}}(\alpha, \beta)$ where α, β are irreducible objects of $\mathcal{B}_\rho\mathcal{C}$; equivalently, this is the signed count of the number of components. To enumerate these, recall the definition of the moduli spaces $\mathcal{B}(\alpha, \beta)$ from Construction 5.16. When neither α nor β is S_ρ , we have an equality $\mathcal{B}(\alpha, \beta) = \mathcal{C}(\alpha, \beta)$, so that $d_{CC} = d^{\text{irr}}$ is the irreducible differential coming from \mathcal{C} itself.

Because $\mathcal{B}(\alpha, S_\rho) = \mathcal{C}(\alpha, \rho) \times_\rho S_\rho$, and $\mathcal{C}(\alpha, \rho)$ consists of irreducible trajectories, this space has dimension at least 4 and hence does not contribute to the differential; so $d_{C\rho} = 0$. Even more easily, $\mathcal{B}(S_\rho, S_\rho) = \emptyset$ by definition, so that $d_{\rho\rho} = 0$.

The only remaining interesting component is $\mathcal{B}(S_\rho, \beta) = -B(\rho, \beta)$, where $B(\rho, \beta)$ is the real blowup of the moduli space $\mathcal{C}(\rho, \beta)$ along a certain $SO(3)$ -equivariant section. In the case we're interested in, $\mathcal{C}(\rho, \beta)$ is 3-dimensional — a disjoint union of free orbits of $SO(3)$ — this section may be assumed nonzero everywhere (define it to be a nonzero value at a point in each component, and extend the definition everywhere by equivariance). The components of $-B(\rho, \beta)$ are thus oriented diffeomorphic to the components of $-\mathcal{C}(\rho, \beta)$, and we defined $V_2(\rho) = \sum \#\mathcal{C}^{\text{irr}}(\rho, \beta)$. This shows that $d_{\rho C} = V_2$, as desired. \square

Because $I_*(Y, w, \sigma_1)$ is isomorphic to the homology of the irreducible flow complex attached to $\mathcal{B}_\rho\mathcal{I}(Y, w, \sigma_1)$, the preceding mapping cone formula immediately implies the following.

Corollary 7.4. *Suppose $\sigma_0 < \sigma_1$ are adjacent signature data on (Y, w) , differing at ρ . Write the continuation map as $\varphi : I_*(Y, w, \sigma_0) \rightarrow I_*(Y, w, \sigma_1)$, and $V_2 : \mathbb{Z}[\tilde{i}(\rho)] \rightarrow I_*(Y, w, \sigma_0)$ for the map sending $1 \mapsto [\sum_\alpha \#M(\rho, \alpha)]$. Then there is an exact triangle*

$$\cdots \rightarrow \mathbb{Z}[\tilde{i}(\rho)] \rightarrow I_*(Y, w, \sigma_0) \rightarrow I_*(Y, w, \sigma_1) \rightarrow \cdots$$

Recall from Definition 3.12 that $I_*(Y, w, \sigma)$ always carries a canonical $\mathbb{Z}/2$ grading. When ρ is reducible, the absolute $\mathbb{Z}/2$ -grading $\tilde{i}(\rho)$ is always even; this can be seen by an application of the index formula to a 4-manifold with reducible connection bounding (Y, w, ρ) .

It follows that $I_*(Y, w, \sigma)$ has a well-defined Euler characteristic, and as the signature data increases, the Euler characteristic decreases:

$$\chi(I_*(Y, \sigma_0)) = \chi(I_*(Y, \sigma_1)) + 1.$$

In the final section, we will use this to define a Casson–Walker type invariant. Through the corollary below, it also shows that our technical passage to the blown-up flow complex $S_\rho\tilde{\mathcal{C}}_*(Y, w, \sigma)$ was in some sense necessary; the continuation map is merely a quasi-isomorphism, and we couldn't have constructed an equivariant homotopy inverse.

Corollary 7.5. *If $\sigma_0 \leq \sigma_1$ are signature data on (Y, w) , the continuation map*

$$\varphi_{\sigma_0 \rightarrow \sigma_1} : \tilde{C}(Y, w, \sigma_0) \rightarrow \tilde{C}(Y, w, \sigma_1)$$

is an $SO(3)$ -equivariant homotopy equivalence if and only if $\sigma_0 = \sigma_1$, even though it is always an equivariant quasi-isomorphism.

Proof. When $\sigma_0 = \sigma_1$ the cylinder is unobstructed running in either direction, and the corresponding cobordism maps give equivariant chain homotopy inverses.

If C is a dg-module over $C_*SO(3)$, and $f : C \rightarrow C'$ is an $SO(3)$ -homotopy equivalence, then the corresponding map $f_* : uC \rightarrow uC'$ is a homotopy equivalence. So by Proposition 7.1, it follows that if $\varphi_{\sigma_0 \rightarrow \sigma_1}$ is an equivariant homotopy equivalence, then $\varphi_* : I_*(Y, w, \sigma_0) \rightarrow I_*(Y, w, \sigma_1)$ is an isomorphism.

Now by Corollary 7.4, if $\sigma_0 \leq \sigma_1$, then

$$\chi(I_*(Y, \sigma_0)) - \chi(I_*(Y, \sigma_1)) = \frac{1}{4} \sum_{\alpha \in \mathfrak{A}(Y, w)} \sigma_1(\alpha) - \sigma_0(\alpha) \geq 0$$

with equality if and only if $\sigma_0 = \sigma_1$. So if the continuation map φ is an equivariant homotopy equivalence, we have $\sigma_0 = \sigma_1$. \square

7.2 An instanton Casson–Walker invariant

If $\alpha \in \mathfrak{A}(Y, w)$, there is a corresponding representation $\text{ad}_\alpha : \pi_1(Y) \rightarrow SO(2)$. Recall from Definition 3.26 that the perturbed rho invariant of α with respect to a regular perturbation π is the rational number

$$\rho_\pi(\alpha) = \sigma_\pi(\alpha) + \rho(\text{ad}_\alpha),$$

with ρ being the relative η invariant of Atiyah–Patodi–Singer, and σ_π is the signature data function of Definition 3.2.

Notice two things. First, this function depends only on the signature data function σ_π , and does not otherwise depend on the perturbation. Second, if $\sigma_0 < \sigma_1$ are adjacent signature data differing precisely at α , we have $\rho_{\sigma_1}(\alpha) = \rho_{\sigma_0}(\alpha) + 4$. We can thus use this as a correction term to counteract the change in Euler characteristic of I_* as we move between chambers and obtain a well-defined rational number.

Definition 7.6. Let (Y, w) be a weakly admissible pair. We say that the *instanton Casson–Walker invariant* is the number

$$\lambda_I(Y, w) = \frac{1}{|\text{Tors } H_1 Y|} \left(\chi(I_*(Y, w, \pi)) + \frac{1}{4} \sum_{\alpha \in \mathfrak{A}(Y)} \rho_\pi(\alpha) \right),$$

for any choice of regular perturbation π . \diamond

This is independent of π because the change in the displayed quantities cancel out as we move between adjacent signature data. We expect this to be an instanton-theoretic definition of the Casson–Walker invariant.

Conjecture 7.7. *For any rational homology sphere Y , the instanton and Casson–Walker invariants coincide; we have $\lambda_I(Y, w) = \lambda_{CW}(Y)$ for all w .*

In particular, this would mean that $\lambda_I(Y, w)$ is independent of w .

Remark 7.8. The Casson invariant and its relatives have a long history. Soon after Casson defined his invariant of integer homology spheres, Taubes [Tau90] gave a definition in terms of signed counts of flat connections modulo gauge; the signs amount to saying that the Casson invariant computes the Euler characteristic of Floer’s instanton homology. Around the same time, Walker published his generalization to rational homology spheres [Wal92]; a signed count of points no longer gives an invariant, and instead one must add canonical rational-valued correction terms (as we do above) to get a well-defined invariant.

The Seiberg–Witten analogue of the Casson–Walker invariant is defined in [Chen:SW]. There, even in the case of integer homology spheres (indeed, even in the case of S^3) one needs to introduce correction terms to get a well-defined invariant; these correction terms only fail to be integers for rational homology spheres. The resulting quantity is indeed equal to the Casson–Walker invariant, as proved in [Lim:SW] for integer homology spheres and [MW:SWCW] for rational homology spheres.

Finally, a version of the Casson–Walker invariant in terms of a precursor to Heegaard Floer homology was provided in [OS:casson]; here again one uses a topological correction term, and their Casson–Walker invariant is given as a sum over spin^c structures. \diamond

As yet we have neither a surgery exact triangle or a connected sum theorem, so we cannot say much presently outside of some specific examples; but we can at least verify it in those few examples. To set convention, the lens space $L(p, q)$ is $-p/q$ surgery on the unknot.

Theorem 7.9. *Conjecture 7.7 is true when Y is an integer homology sphere or a lens space $L(p, q)$.*

Proof. When Y is an integer homology sphere, the sum over abelian flat connections is tautologically zero and the choice of signature data is vacuous, so we simply have $\lambda_I(Y) = \chi I_*(Y)$. The main result of [Tau90] is that we have an equality $\chi I_*(Y) = 2\lambda_C(Y)$ with half of the Casson invariant of Y . Walker’s normalization conventions have $\lambda_{CW}(Y) = 2\lambda_C(Y)$, giving us the desired result.

For lens spaces, the trivial perturbation is a regular perturbation, so we may take $\sigma = 0$. Then the invariant $\lambda_I(L(p, q), w)$ is given as

$$\frac{1}{4p} \sum_{\alpha \in \mathfrak{A}(L(p, q), w)} \rho(\alpha).$$

Write $H^2(L(p, q)) \cong \mathbb{Z}/p$. In the case $w = \emptyset$, the abelian reducibles can be identified with those integers ℓ with $0 < \ell < p/2$, and have adjoint representation given by $\text{ad}_\ell(i) = e^{2\pi i(2\ell)/p}$. In the case p is even and $[w]$ is nonzero in $\mathbb{Z}/2$ -homology, the abelian reducibles can be identified with the odd integers $0 < 2\ell + 1 < p$, and $\text{ad}_\ell(i) = e^{2\pi i(2\ell+1)/p}$.

Let us handle the first case. Using the formula of [HK11, Proposition 2.18], we have

$$\rho(\text{ad}_\ell) = \frac{4}{p} \sum_{k=1}^{p-1} \cot\left(\frac{\pi k}{p}\right) \cot\left(\frac{\pi k q}{p}\right) \sin^2\left(\frac{2\pi k \ell}{p}\right).$$

Summing over $0 < \ell < p/2$ and multiplying by $1/4p$, we see that $\lambda_I(L(p, q), \emptyset)$ is

$$\frac{1}{p^2} \sum_{0 < \ell < p/2} \sum_{k=1}^{p-1} \cot\left(\frac{\pi k}{p}\right) \cot\left(\frac{\pi k q}{p}\right) \sin^2\left(\frac{2\pi k \ell}{p}\right).$$

Let us compute the identical sum

$$\frac{1}{2p^2} \sum_{\ell=0}^{p-1} \sum_{k=1}^{p-1} \cot\left(\frac{\pi k}{p}\right) \cot\left(\frac{\pi k q}{p}\right) \sin^2\left(\frac{2\pi k \ell}{p}\right);$$

this coincides with the sum above because the \sin^2 term is symmetric with respect to swapping ℓ and $p - \ell$, and when p is even the term corresponding to $\ell = p/2$ automatically vanishes.

When p is even, notice that every term corresponding to $k = p/2$ automatically vanishes because $\cot(\pi/2) = 0$. Further, observe that $\sum_{\ell=0}^{p-1} \sin^2\left(\frac{2\pi k \ell}{p}\right) = p/2$ whenever $0 < k < p$ and $k \neq p/2$: this follows by the double-angle formula $\sin^2(x) = \frac{1}{2} - \frac{1}{2} \cos(2x)$ and the fact that $\sum_{\ell=0}^{p-1} \cos(4\pi k \ell/p) = 0$, because this latter expression is the real part of the sum $0 = 1 + \zeta + \dots + \zeta^{p-1}$. Here $\zeta = e^{4\pi i k/p} \neq 1$ because $0 < k < p$ and $k \neq p/2$.

Combining these two facts, we may write

$$\lambda_I(L(p, q), \emptyset) = \frac{1}{4p} \sum_{k=1}^{p-1} \cot\left(\frac{\pi k}{p}\right) \cot\left(\frac{\pi k q}{p}\right),$$

which is a Dedekind sum traditionally denoted $s(q; p)$.

Finally, we have $\lambda_{CW}(L(p, q)) = s(q; p)$, by [Wal92, Proposition 6.3]; in comparing to the reference, note that Walker's $L_{p/q}$ is what we call $L(-p, q)$, and hence $\lambda_{CW}(L(p, q))$ is written in Walker's book as $\lambda(-L_{p/q}) = -\lambda(L_{p/q})$, as desired.

The computation for the case $[w] \neq 0 \in H_1(Y; \mathbb{F}_2)$ is similar, with the relevant ρ invariant being

$$\rho(\text{ad}_\ell) = \frac{4}{p} \sum_{k=1}^{p-1} \cot\left(\frac{\pi k}{p}\right) \cot\left(\frac{\pi k q}{p}\right) \sin^2\left(\frac{\pi k \ell}{p}\right),$$

for odd $0 < \ell < p$, so that

$$\lambda_I(L(p, q), w) = \frac{1}{p^2} \sum_{j=0}^{p/2-1} \sum_{k=1}^{p-1} \cot\left(\frac{\pi k}{p}\right) \cot\left(\frac{\pi k q}{p}\right) \sin^2\left(\frac{\pi k(2j+1)}{p}\right).$$

This simplifies to the same quantity $s(q; p)$. The argument is nearly identical: eliminate the terms with $k = p/2$, and sum over j to get rid of the final factor; here one uses that when ζ is a p 'th root of unity for p even we have $\zeta + \zeta^3 + \dots + \zeta^{2p-1} = 0$. \square

Remark 7.10. Using the computations in [Hel22], together with further computations of the rho invariants for abelian connections on spherical manifolds, one can also verify the conjecture for any 3-manifold of the form $SU(2)/\Gamma$ (with $w = \emptyset$), where Γ is a finite subgroup of $SU(2)$. We will not carry out this computation in detail here. \diamond

Appendix

A Stratified-smooth spaces

In this appendix, we construct the model for singular chains $C_*(X)$ we use in the main body of the text, where X is a smooth manifold, which we will call *geometric chains*; this complex should be functorial for smooth maps and have homology groups canonically isomorphic to singular homology. Let us briefly outline the desired construction.

This geometric chain complex should be free on a certain generating set, the *nondegenerate, nontrivial probes*, where these probes are something analogous to (P, ϕ) a compact oriented smooth manifold with corners P with a smooth map ϕ to X ; the boundary map should be analogous to the passage to ∂P in the case of smooth manifolds with corners.

For our purposes, these probes must be more general than smooth manifolds with corners: they must include the moduli spaces of instantons $\check{M}_\zeta^+(\alpha, \beta)$ discussed in Section 3.4; in Definition 3.17, we named the kind of object we obtain a *stratified-smooth manifold*. When proving the excision property, one needs to be able to take transverse intersections, which leads to the more general notion of *stratified-smooth spaces* below, roughly the kind of space obtained by taking transverse intersections of stratified-smooth manifolds.

The first two sections of this appendix develop the bare minimum of foundations and orientation theory of stratified-smooth manifolds needed to define geometric chains. In Section 5 of the main text, we also need a certain blowup construction on these stratified-smooth spaces, which we set up in Appendix A.3, and carry out some necessary computations. In Appendix A.4 we construct the geometric chain complex and verify an appropriate version of the Eilenberg-Steenrod axioms, and therefore verify that the resulting homology groups are canonically isomorphic to singular homology. In the final section, Appendix A.5, we perform a truncation procedure on the geometric chain complex; this is used in Section 4 to ensure we only need to construct moduli spaces within a certain dimension range.

A.1 Foundations of stratified-smooth spaces

Smooth manifolds with corners should be special cases of stratified-smooth spaces, so first we must decide on a preferred definition. There are a number of inequivalent definitions (see [Joy12] for a survey). The more general definitions have more complicated boundary operators, which cannot be defined solely in terms of the codimension-1 strata. This is harder to work with effectively when defining and computing the boundary of a stratified-smooth space; we prefer to be more restrictive.

With eyes towards the moduli spaces in the main text, we instead use a version of manifolds with corners and stratified-smooth spaces where one is stratified by a set of faces, which need not be connected, whose closures are all again manifolds with corners (resp. stratified-smooth spaces). To do so, we record the stratification with an appropriate kind of partially ordered set.

Definition A.1. A graded poset is a finite partially ordered set $(\Delta, <)$ equipped with a ‘dimension’

function $d : \Delta \rightarrow \mathbb{Z}_{\geq 0}$, which is order-preserving in the sense that $s < t$ implies $d(s) < d(t)$.

We say that the largest integer in the image of d is the dimension of Δ . \diamond

The simplest example of a graded poset, and one which will be very useful for us, is the cube poset $\{0, 1\}^n$ for some natural number n . Writing elements as functions $s : \{1, \dots, n\} \rightarrow \{0, 1\}$, the ordering has $s \leq t$ if $s(i) \leq t(i)$ for all i , and the dimension function is $d(s) = \sum s(i)$.

For the objects of interest, the graded poset Δ will have the property that the sub-poset

$$\Delta_{\geq s} = \{t \in \Delta \mid t \geq s\}$$

is always isomorphic to a cube poset $\{0, 1\}^{\dim X - d(s)}$, with dimension function shifted up by $d(s)$.

Partially ordered sets are relevant because the natural relation between the different strata of a stratified space is a partial ordering. Below, we use a definition of stratified space which explicitly records the relevant partially ordered set.

Definition A.2. A *stratified space* (X, Δ) is a pair of a poset Δ and a topological space X equipped with, for each $s \in \Delta$, a locally closed subset $X_s \subset X$; these subsets are required to satisfy the property $\overline{X_s} = \bigcup_{t \leq s} X_t$.

An isomorphism $(X, \Delta_X) \cong (Y, \Delta_Y)$ of stratified spaces is a pair (f, \bar{f}) , where $\bar{f} : \Delta_X \rightarrow \Delta_Y$ is a poset isomorphism and f is a homeomorphism with $f(X_s) = Y_{\bar{f}(s)}$. \diamond

Next we record the naive version of a smooth version of stratified spaces, where one merely equips each stratum with a smooth structure.

Definition A.3. A *naive stratified-smooth space* (X, Δ) of dimension n is a pair of a graded poset Δ and a second-countable Hausdorff space X stratified by Δ , where each stratum X_s is given the structure of a smooth manifold of dimension $d(s)$.

A *naive stratified-smooth map* $f : X \rightarrow M$ from X to a smooth manifold M is a continuous function for which each $f_s : X_s \rightarrow M$ is smooth; we say $m \in M$ is a regular value if it is a regular value for all maps $f_s : X_s \rightarrow M$.

An *isomorphism* (f, \bar{f}) between two naive stratified-smooth spaces (X, Δ_X) and (Y, Δ_Y) is an isomorphism of stratified spaces so that the induced maps between strata $f_s : X_s \rightarrow Y_{\bar{f}(s)}$ are all diffeomorphisms. \diamond

Let us discuss one standard example of a naive stratified-smooth space, as well as two important constructions which will be used in the definition of the local models below.

Example A.4. The space $[0, \infty)^n$ can be given the structure of a naive stratified-smooth space. The stratification is by the cube poset $\Delta = \{0, 1\}^n$; the faces are products of $\{0\}$ and $(0, \infty)$. More precisely, write

$$F(i) = \begin{cases} \{0\} & i = 0 \\ (0, \infty) & i = 1 \end{cases}$$

Then the face associated to the stratum $s \in \{0, 1\}^m$ is the subspace

$$(0, \infty)^s := \prod_{i=1}^n F(s(i));$$

that is, the i 'th factor is $\{0\}$ if $s(i) = 0$ and is $(0, \infty)$ if $s(i) = 1$. \diamond

Example A.5. Let (X, Δ) be a naive stratified-smooth space, and let $U \subset X$ be an open subset. Then U also inherits the structure of a stratified-smooth space, taking $\Delta(U) \subset \Delta$ to be those s for which $U \cap X_s \neq \emptyset$ with strata $U_s = U \cap X_s$. \diamond

Example A.6. Let (X, Δ) be a naive stratified-smooth space of dimension n , and let $\psi : X \rightarrow \mathbb{R}^\ell$ be a stratified-smooth map with zero as a regular value. Write $\Delta(\psi) \subset \Delta$ for the sub-poset

$$\Delta(\psi) = \{s \in \Delta \mid \psi^{-1}(0) \cap X_s \neq \emptyset\},$$

with grading function $d_\psi(s) = d(s) - \ell$. Then $(\psi^{-1}(0), \Delta(\psi))$ is a naive stratified-smooth space of dimension $n - \ell$. \diamond

We combine these into one type of model space; a stratified-smooth space will be a naive stratified-smooth space locally isomorphic to one of the model spaces.

Definition A.7. Let $V \subset [0, \infty)^m \times \mathbb{R}^n$ be an open subset containing zero, and let $\psi : V \rightarrow \mathbb{R}^\ell$ be a naive stratified-smooth map with zero as a regular value. Then $Z(\psi) = \psi^{-1}(0)$ carries the natural structure of a naive stratified-smooth space of dimension $m + n - \ell$. We call such a naive stratified-space $Z(\psi)$ a *cutout locus*. \diamond

Notice that the assumption $0 \in V$ implies that $\Delta(V) = \{0, 1\}^m$; every face has a point near 0 and hence a point in V . Thus the stratification poset only changes when passing to $\Delta(\psi)$.

Before moving on to the definition of stratified-smooth space, we make a brief observation on the combinatorics of this poset $\Delta(\psi)$. The following argument is inspired by [SS10, Section 2c], which we follow again when discussing orientations later.

Lemma A.8. *The poset $\Delta(\psi) \subset \{0, 1\}^m$ is an upward-closed sub-poset; that is, if $s \in \Delta(\psi)$ and $t \in \{0, 1\}^m$ has $t > s$, then $t \in \Delta(\psi)$ as well.*

Proof. Suppose that $Z(\psi)_s \subset (0, \infty)^s \times \mathbb{R}^n$ is nonempty.¹¹ Because $Z(\psi)_s$ is the preimage of a regular value, it is a smooth manifold of codimension ℓ inside $(0, \infty)^s$. Choose a small sphere $c : S^{\ell-1} \rightarrow ((0, \infty)^s \times \mathbb{R}^n) \setminus Z(\psi)_s$ which has linking number one with $Z(\psi)_s$, so that in particular $\psi c : S^{\ell-1} \rightarrow \mathbb{R}^\ell - \{0\}$ is not null-homotopic.

It follows that in fact $c : S^{\ell-1} \rightarrow ([0, \infty)^m \times \mathbb{R}^n) \setminus Z(\psi)$ is not null-homotopic (for otherwise ψc would be null-homotopic, which it is not).

¹¹Recall that here $s \in \{0, 1\}^m$ and the exponent means the i 'th factor is either $\{0\}$ or $(0, \infty)$, depending on whether $s(i)$ is 0 or 1.

Now for $s < t$, we may homotop c into $c' : S^{\ell-1} \rightarrow (0, \infty)^t \times \mathbb{R}^n$ by an arbitrarily small homotopy; taking this homotopy sufficiently small, its image will not meet $Z(\psi)$, as $Z(\psi)$ is closed and the domain is compact. In particular, $\psi c'$ is homotopic to ψc as a map to $\mathbb{R}^n \setminus \{0\}$, and therefore $\psi c'$ is not null-homotopic; thus c' is not null-homotopic in $((0, \infty)^t \times \mathbb{R}^n) \setminus Z(\psi)_t$. In particular, $Z(\psi)_t$ is nonempty, which is what we wanted to show. \square

Remark A.9. It follows that $\Delta(\psi)$ is a *cubical poset*, which means that for each $s \in \Delta(\psi)$, we have $\Delta(\psi)_{\geq s} \cong \{0, 1\}_{\geq s}^m \cong \{0, 1\}^{m-d(s)}$; the poset of strata lying above a given stratum is isomorphic to the cube poset of appropriate dimension. Soon after we define the boundary operator, we will see that this observation underlies the relation $\partial^2 = 0$. \diamond

We now define the geometric objects of interest.

Definition A.10. A *stratified-smooth space* is a naive stratified-smooth space (X, Δ) so that, for each $x \in X$, there exists an open subset $x \in U \subset X$ and a cutout locus $Z(\psi)$ so that U is isomorphic to $Z(\psi)$ as naive stratified-smooth spaces. We call a choice of open set U and isomorphism $U \cong Z_\psi$ a *cutout chart* for x . \diamond

Note that in particular we do not have an ‘atlas’ of such charts in any sense; we merely demand that they exist.

Example A.11. In [Joy12], a sequence of increasingly more strict notions of manifolds with corners are defined, and $\langle N \rangle$ -*manifolds* defined there give examples of stratified-smooth spaces. Our Definition 3.17 of stratified-smooth manifold is very closely related to Joyce’s definition $\langle N \rangle$ -manifold, though ours is a weaker object, as the smooth structure on the top stratum of a stratified-smooth manifold does not extend up to the boundary.

In particular, the simplex Δ^k is a stratified-smooth space for all k , as are all products of simplices. \diamond


Example A.12. This is a more restrictive definition than it may seem at first, because of the combinatorial nature of an isomorphism between naive stratified-smooth spaces. This is best exemplified by the *teardrop*, defined in [Joy12, Remark 2.11] as an example of a naive topological manifold with corners which does not satisfy any of the more stringent definitions of manifold with corners. His teardrop is the set

$$\{(x, y) \mid x \geq 0, y^2 \leq x^2 - x^4\},$$

which has three strata: a corner stratum, one boundary stratum, and one open stratum. The stratification poset Δ can be identified with $\{0, 1, 2\}$ with the total order $0 < 1 < 2$.

This is not locally isomorphic as a naive stratified-smooth space to $[0, \infty)^2$ near the corner point. A neighborhood of the corner point in the teardrop has *one* (disconnected!) codimension-1 stratum, while a neighborhood of 0 in $[0, \infty)^2$ has two.

In fact, by Lemma A.8 above, the teardrop is not a stratified-smooth space: if it were, $\Delta_{\geq 0} = \{0, 1, 2\}$ would be isomorphic to a cubical poset, which it is not.



A picture of the teardrop with its three strata.

◇

Now that we have a definition of stratified-smooth space, we can define its boundary. Part of the reason we choose the combinatorial language in the definition of stratified-smooth space is that it makes the definition of boundary very simple, using the following lemma; notice that this lemma fails for the boundary stratum of the teardrop.

Lemma A.13. *Let X be a stratified-smooth space of dimension n , and let X_s be a stratum of X . Then $\overline{X_s}$ naturally carries the structure of a stratified-smooth space of dimension $d(s)$.*

Proof. The underlying topological space is of course the closure $\overline{X_s} \subset X$ with the subspace topology. Because $\overline{X_s} = \cup_{t \leq s} X_t$ we see that $\overline{X_s}$ carries the natural structure of a stratified space with smooth structure on each stratum, with face poset $\Delta_{\leq s}$.

We need to construct cutout charts for this subspace. If $x \in X_t \subset \overline{X_s}$, choose a cutout chart for X near x . To spell out all this data: choose an open set $x \in U \subset X$, an open subset $V \subset [0, \infty)^m \times \mathbb{R}^n$, a naive stratified-smooth map $\psi : V \rightarrow \mathbb{R}^\ell$ with zero as a regular value, an isomorphism of posets $\overline{\phi} : \Delta(U) \rightarrow \Delta(\psi)$, and $\phi : U \cong Z_\psi$ an isomorphism of naive stratified-smooth spaces lifting $\overline{\phi}$. Then $\overline{\phi}(s) \in \{0, 1\}^m$; we abbreviate this to \overline{s} .

Write $[0, \infty)^{\overline{s}}$ for the product of $\{0\}$ and $[0, \infty)$, with a 0 in the i 'th coordinate if $\overline{s}(i) = 0$ and $[0, \infty)$ otherwise. Write

$$V_{\leq \overline{s}} = V \cap ([0, \infty)^{\overline{s}} \times \mathbb{R}^n),$$

and $\psi_{\leq \overline{s}}$ for the restriction of ψ to this subspace.

Then ϕ and $\overline{\phi}$ restrict to give an isomorphism of naive stratified-smooth spaces between $U \cap \overline{X_s} = U_{\leq \overline{s}}$ and $Z(\psi_{\leq \overline{s}})$. This gives the desired cutout chart on $\overline{X_s}$. □

This is enough to define the boundary of a stratified-smooth space.

Definition A.14. Let X be a stratified-smooth space. The boundary of X is the stratified-smooth space

$$\partial X = \bigsqcup_{\substack{s \in \Delta \\ d(s) = \dim X - 1}} \overline{X}_s. \quad \diamond$$

Next, we will use Lemma A.8 to obtain an understanding of $\partial^2 X$.

Proposition A.15. *The boundary $\partial^2 X$ is isomorphic to $\{0, 1\} \times S$ for some stratified-smooth space S of dimension $\dim X - 2$.*

Proof. First we show that for all $s \in \Delta$ with $d(s) = \dim X - 2$, we have $\Delta_{\geq s} \cong \{0, 1\}^2$; in particular, for each such s , there are exactly two t with $d(t) = \dim X + 1$ and $s < t$. This follows by picking a point $x \in X_s$ and a cutout chart $x \in U \subset X$ with $U \cong Z_\psi$. Then we have

$$\Delta_{\geq s} = \Delta(U)_{\geq s} \cong \Delta(Z_\psi)_{\geq s} \cong \{0, 1\}^2;$$

the first isomorphism follows because if $s < t$ we have $X_t \cap U \neq \emptyset$ (because $X_s \subset \overline{X}_t$ and U is open), the second by the definition of isomorphism of stratified-smooth spaces, and the last by Lemma A.8 and the discussion in Remark A.9.

Next, observe that because

$$\partial \overline{X}_t = \bigsqcup_{\substack{s \in \Delta \\ s < t \\ d(s) = \dim X - 2}} \overline{X}_s,$$

we have

$$\partial^2 X = \bigsqcup_{\substack{t, s \in \Delta \\ s < t \\ d(s) + 1 = d(t) = \dim X - 1}} \overline{X}_s.$$

Because there are exactly two such t , this disjoint union is isomorphic to

$$\partial^2 X = \{0, 1\} \times \bigsqcup_{\substack{s \in \Delta \\ d(s) = \dim X - 2}} \overline{X}_s,$$

as desired. □

We move on to foundations of smooth maps and fiber products.

Definition A.16. Let (X, Δ) and (Y, Δ') be stratified-smooth spaces, and let B be a smooth manifold without boundary. We say $f : X \rightarrow B$ is stratified-smooth if f is continuous, smooth on each stratum, and satisfies the following property: for each $x \in X$ there exists a cutout chart $x \in U \cong Z(\psi) \subset V$ and an extension of $f \circ \psi : Z(\psi) \rightarrow B$ to a map $F : V \rightarrow B$ which is continuous and smooth on each stratum.

Given two stratified-smooth maps $f : X \rightarrow B$ and $g : Y \rightarrow B$ we say f and g are stratified-transverse if their restrictions to any pair of strata X_s and $Y_{s'}$ are transverse. □

Then as desired the fiber product of two stratified-smooth spaces is again stratified-smooth. This includes, for instance, preimages of regular values of stratified-smooth functions to smooth manifolds (or more generally, intersection with submanifolds, hence in particular a generic zero set of a section of a vector bundle).

Proposition A.17. *If f and g are stratified-transverse maps as in the above definition, there is a natural stratified-smooth space structure on the fiber product $X \times_B Y$, whose underlying topological space is the usual fiber product.*

Proof. The face poset of $X \times_B Y$ is a subset

$$\Delta_B \subset \Delta \times \Delta',$$

consisting of those pairs (s, s') so that $X_s \times_B Y_{s'}$ is nonempty, equipped with the grading $d_B(s, s') = d(s) + d'(s') - \dim B$. That the ordering induced from $\Delta \times \Delta'$ coincides with the closure-inclusion ordering amounts to the relation $\overline{X_s \times_B Y_{s'}} = \overline{X_s} \times_B \overline{Y_{s'}}$. The forward inclusion is obvious, while the reverse is not; still, it follows along lines analogous to the proof of Lemma A.8.

Pick $(x, y) \in X_s \times_B Y_{s'}$. Write $k = \dim X - d(s)$ and $k' = \dim Y - d'(s')$; we have open sets

$$V \subset [0, \infty)^k \times \mathbb{R}^{d(s)+m}, \quad V' \subset [0, \infty)^{k'} \times \mathbb{R}^{d'(s')+n}$$

and open neighborhoods

$$x \in U \subset X, \quad y \in U' \subset Y,$$

and maps $\psi : V \rightarrow \mathbb{R}^m$ and $\psi' : V' \rightarrow \mathbb{R}^n$, with isomorphisms of naive stratified-smooth spaces

$$\phi : Z(\psi) \cong V, \quad \phi' : Z(\psi') \rightarrow V'$$

sending x and y to zero, respectively.

Further, by the definition of stratified-smooth map, we may assume V and V' are equipped with extensions F and G of f and g , which are only defined on $Z(\psi)$ and $Z(\psi')$. We suppose these charts are chosen sufficiently small that the images of F and G lie in a single chart of B ; for convenience, we thus suppose $B = \mathbb{R}^b$.

Then consider the open set

$$\overline{U} = U \times U' \subset [0, \infty)^k \times \mathbb{R}^{d(s)+m} \times [0, \infty)^{k'} \times \mathbb{R}^{d'(s')+n},$$

and the map $\xi = (\psi, F - G, \psi') : \overline{U} \rightarrow \mathbb{R}^{m+b+n}$. Because the restriction of $F - G$ to the zero set of (ψ, ψ') is precisely $f - g$, the zero set of ξ is the fiber product $Z(\psi) \times_{\mathbb{R}^b} Z(\psi')$. Further, because ψ and ψ' have zero as a regular value, and $F - G$ has zero as a regular value on the zero locus of ψ and ψ' , it follows that ξ has zero as a regular value on each stratum as well.

Thus $U \times_B U' = Z(\xi)$, and (ϕ, ϕ') restricts to give a homeomorphism $Z(\xi) \cong V \times_B V'$ which sends strata to strata and is a diffeomorphism on each stratum. We have constructed the necessary cutout chart. \square

This fiber product operation is associative: if we have maps $f : X \rightarrow A$, $g : Y \rightarrow A \times B$ and $h : Z \rightarrow B$ satisfying the appropriate transversality requirements, then there is a canonical diffeomorphism of stratified-smooth spaces

$$X \times_A (Y \times_B Z) \cong (X \times_A Y) \times_B Z.$$

Further, using the description of its face poset in the previous argument, we may explicitly describe its boundary.

Proposition A.18. *The fiber product $X \times_B Y$ has boundary*

$$\partial(X \times_B Y) = (\partial X) \times_B Y \sqcup X \times_B \partial Y.$$

Proof. The left-hand side is

$$\bigsqcup_{\substack{s \in \Delta_X, t \in \Delta_Y \\ d(s) + d'(t) - \dim B = \dim X + \dim Y - \dim B - 1}} \overline{X}_s \times_B \overline{Y}_t;$$

now $d(s) + d'(t) = \dim X + \dim Y - 1$ if and only if one of $\dim X - d(s)$ and $\dim Y - d'(t)$ is one and the other is zero; this means this is precisely

$$\left(\bigsqcup_{\substack{s \in \Delta_X, t \in \Delta_Y \\ d(s) = \dim X \\ d(t) = \dim Y - 1}} \overline{X}_s \times_B \overline{Y}_t \right) \sqcup \left(\bigsqcup_{\substack{s \in \Delta_X, t \in \Delta_Y \\ d(s) = \dim X - 1 \\ d(t) = \dim Y}} \overline{X}_s \times_B \overline{Y}_t \right).$$

Because the disjoint union $\bigsqcup_{d(s) = \dim X} \overline{X}_s$ is X itself, while $\bigsqcup_{d(s) = \dim X - 1} \overline{X}_s = \partial X$, the above sum simplifies to

$$\partial X \times_B Y \sqcup X \times_B \partial Y,$$

as desired. \square

In addition to stratified-smooth spaces defined by equations (such as the general fiber products described above), we will also find need in Section 5.2 for stratified-smooth spaces defined by *inequalities*; we record the relevant statement below.

Proposition A.19. *If (X, Δ) is a stratified-smooth space and $L : X \rightarrow \mathbb{R}$ is stratified-smooth, then for a regular value c of L , the set $X^{\leq c} = L^{-1}(-\infty, c]$ is a stratified-smooth space with boundary*

$$\partial X^{\leq c} = L^{-1}(c) \sqcup (\partial X)^{\leq c}.$$

Proof. The stratification poset for $X^{\leq c}$ is denoted $\Delta^{\leq c}$, and is a sub-poset of $\Delta \times \{-1, 0\}$; we say $(s, -1) \in \Delta^{\leq c}$ if $L^{-1}(c) \cap X_s$ is nonempty, while $(s, 0) \in \Delta^{\leq c}$ if $L^{-1}(-\infty, c) \cap X_s \neq \emptyset$. Once we verify that $X^{\leq c}$ is indeed a stratified-smooth space, this gives the stated boundary formula.

The notation for $\Delta^{\leq c}$ chosen so that $d(s, 0) = d(s)$ and $d(s, -1) = d(s) - 1$. It is straightforward to see that this is a stratified topological space; the strata $(X^{\leq c})_{(s,0)}$ are open subsets of X_s , hence carry

the natural structure of smooth manifolds, whereas the strata $(X^{\leq c})_{(s,-1)}$ are given as regular values of a smooth map $L : X_s \rightarrow \mathbb{R}$, hence also in a natural way smooth manifolds. So $X^{\leq c}$ is at least a naive stratified-smooth space; the stated boundary relation follows immediately from the definition of $\Delta^{\leq c}$ and its dimension function.

The existence of cutout charts for $x \in X_{(s,0)}^{\leq c}$ follows immediately by restricting cutout charts for $x \in X_s$. The interesting case is $x \in X_{(s,-1)}^{\leq c}$. In this case, one starts by picking a cutout chart for $x \in X_s$; that is, there is a choice of open set $V \subset [0, \infty)^m \times \mathbb{R}^n$, a stratumwise-smooth map $\psi : V \rightarrow \mathbb{R}^\ell$, and an isomorphism $Z(\psi) \cong U$ of naive stratified-smooth spaces, where U is an open neighborhood of x . By the assumption that L is stratified-smooth, we may assume L extends to a function on V , which we denote by the same symbol.

We may extend this to a map $\hat{\psi} : V \times [0, \infty) \rightarrow \mathbb{R} \oplus \mathbb{R}^\ell$ defined by

$$\hat{\psi}(v, s) = (\psi(v), L(v) + s - c).$$

For (s, v) to be in the zero set of this map means precisely that $\psi(v) = 0$ and that $L(v) \leq c$. The map $\pi : Z(\hat{\psi}) \rightarrow Z(\psi)$ is a topological embedding, and a smooth embedding on each stratum; composing this with the map $\varphi : Z(\psi) \cong U$ gives us the desired stratified-smooth isomorphism from $Z(\hat{\psi})$ to $U_{\leq c}$, as desired. \square

A.2 Orientations on stratified-smooth spaces

In this section, we will define and analyze the notion of an *oriented* stratified-smooth space. Briefly, this is an orientation on the top strata which is consistent with the orientations induced by some cutout chart. We will also use cutout charts to define the boundary orientation.

The first of these requires that we first understand the orientation on a cutout chart; for the second, we must verify that the induced boundary orientation does not depend on the choice of cutout chart.

Construction A.20. To set conventions: suppose $f : X \rightarrow B$ is a smooth map between oriented smooth manifolds and $b \in B$ is a regular value. Given $x \in f^{-1}(b)$, we write $T_x Z$ for the tangent space to $f^{-1}(b)$; write $di : T_x Z \rightarrow T_x X$ for the differential of the inclusion map. If s is a splitting of the surjective map $df_x : T_x X \rightarrow T_b B$, then we orient $T_x Z$ by the assumption that $(di, s) : T_x Z \oplus T_b B \rightarrow T_x X$ is an orientation-preserving isomorphism.

Now let $\psi : [0, \infty) \times \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a continuous map which, on each stratum, is smooth and has zero as a regular value. Write ψ° for its restriction to $(0, \infty) \times \mathbb{R}^n$, and write ψ^∂ for its restriction to $\mathbb{R}^n \cong \{0\} \times \mathbb{R}^n$, equipped with the boundary orientation (here, negative the standard orientation).

Then the zero sets $Z(\psi^\circ)$ and $Z(\psi^\partial)$ are canonically oriented, being regular level sets of smooth maps between oriented smooth manifolds. \diamond


For convenience in what follows below, it is worth mildly generalizing the definition of cutout chart. Instead of frontloading the half-line coordinates and demanding that the cutout function ψ be defined on $V \subset [0, \infty)^m \times \mathbb{R}^n$ an open subset, we allow V to be an open subset of any product of copies of

$[0, \infty)$ and \mathbb{R} in any order. This is convenient primarily for easy discussion of products (where we would otherwise have to shuffle coordinates to move the half-line coordinates to the front, introducing a sign). When discussing the orientation on the boundary of $Z(\psi)$, one first passes to the boundary (with its standard boundary orientation) and then passes to the regular zero set of ψ^∂ on the corresponding oriented boundary strata.

This in hand, we can define an oriented stratified-smooth space.

Definition A.21. Let (X, Δ) be a stratified-smooth space. An *orientation* of (X, Δ) is a choice of orientation for each top stratum so that for each $x \in X$, there *exists* a cutout chart (U, V, ψ, ϕ) around x so that the induced diffeomorphism on the top stratum, $\phi : U^\circ \rightarrow Z(\psi)^\circ$, is orientation-preserving. \diamond

In understanding this definition, a useful example to keep in mind is the θ -graph: there is a continuous function on $[0, 1]^2$ (which, on each stratum, is smooth and has zero as a regular value) whose zero set on the boundary is $\{1/2\} \times \{0, 1\}$ but on the interior is three arcs running between these points. An orientation on this stratified-smooth space orients two of the arcs in one direction and one arc in the opposite direction. For a given cutout chart, only one of these six possible orientations will be oriented compatibly with that cutout chart; but the definition only asks that *some* cutout chart be oriented compatibly.



The six ways to make the θ -graph into an oriented stratified-smooth space.



The two orientations of the top strata which do not define an oriented stratified-smooth space.

There is no harm in the fact that this connected stratified-smooth space admits six distinct orientations. That the two ‘orientations’ in which all three arcs go the same direction are not compatible with any cutout chart is shown by the following lemma, which is the same argument as in [SS10, Section 2c]. This lemma will be the foundation of our proof that the boundary of a 1-dimensional stratified-smooth space is zero points (counted with sign), and that the boundary orientation is well-defined.

Lemma A.22. Suppose $\psi : [0, \infty) \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ is continuous and smooth with zero as a regular value on each stratum. Suppose furthermore that $(\psi^\partial)^{-1}(0) = \{0\}$, oriented negatively. Write $\psi_{t,r} : \{t\} \times D_r^n \rightarrow \mathbb{R}^n$ for the restriction to a slice at time t and the ball of radius $r > 0$ around the origin. For sufficiently small t, r , the map $\psi_{t,r}$ has zero as a regular value and $\psi_{t,r}^{-1}(0)$ is a finite collection of points whose orientations sum to one.

Proof. To briefly remark on the signs, remember that ψ^∂ has domain $\{0\} \times \mathbb{R}^n$ with the boundary orientation, which is negative the standard orientation; thus $\psi_{0,r}^{-1}(0)$ is a single point, oriented oppositely to $(\psi^\partial)^{-1}(0)$. In particular, $\psi_{0,r}^{-1}(0)$ is a single *positively-oriented* point.

By the Sard lemma, for a generic set of $t, r \in (0, \infty)^2$, the given map has zero as a regular value and has no zeroes on the boundary of D_r^n . Because D_r^n is compact, the desired count is finite. Further, for r, t sufficiently small, the subset $S_r^{n-1} \times [0, t]$ does not meet $\psi^{-1}(0)$ (by continuity of ψ and its behavior on the boundary). Write

$$\bar{\psi}_{t,r} = \psi i_{t,r} : S^{n-1} \rightarrow \mathbb{R}^n \setminus 0$$

for the composite of the inclusion $S_r^n \hookrightarrow \{t\} \times \mathbb{R}^n$ and the map ψ . Then the degree of $\bar{\psi}_{t,r}$ is precisely the number of intersection points of $\{t\} \times D_r^n$ with $\psi^{-1}(0)$.

By assumption, $\deg \bar{\psi}_{0,r} = 1$ for all r . Further, because $S_r^{n-1} \times [0, t]$ does not meet $\psi^{-1}(0)$, the maps $\bar{\psi}_{t,r}$ are homotopic for all sufficiently small $t \geq 0$ and $r > 0$. Thus the degree of $\bar{\psi}_{t,r}$ for small $t > 0$ is also equal to one, giving the desired result. \square

We may understand these counts intrinsically. Suppose X is an oriented 1-dimensional stratified-smooth space. We classify the connected components of the open stratum of X as three types: components which are equal to their own closure are called *isolated*. Otherwise, either $\bar{C} = C \cup \{x\}$ or $\bar{C} = C \cup \{x, y\}$. In either case, $C \cong \mathbb{R}$; if both ends of C accumulate (and accumulate to the same point) we say C is *heteroclinic*, whereas if exactly one end of C accumulates to x we say C is *homoclinic*. In the latter case, we say C *leaves* x or *approaches* x depending on whether or not x is an accumulation point of the negatively-oriented end of C or the positively-oriented end, respectively.

Lemma A.23. *In the situation of Lemma A.22, for sufficiently small $t, r > 0$ the count $\#\psi_{t,r}^{-1}(0)$ is the number of homoclinic components of $Z(\psi^\circ)$ entering 0 minus the number of homoclinic orbits of $Z(\psi^\circ)$ leaving 0.*

Proof. This amounts to counting the contribution of each component to the count $\psi_{t,r}^{-1}(0)$. Consider a neighborhood of $[0, \varepsilon] \times D_r^n$ of 0 with $\varepsilon, r > 0$ sufficiently small that

$$Z(\psi) \cap \partial([0, t] \times D_r^n) = Z(\psi) \cap (\{0, t\} \times D_r^n),$$

and so that this intersection is transverse. Write $Z'(\psi) = Z(\psi) \cap ([0, t] \times D_r^n)$; this is a compact 1-dimensional stratified-smooth space. Now consider what happens to the components of $Z(\psi)$ as we pass to $Z'(\psi)$:

- An isolated component of $Z(\psi)$ either intersects $Z'(\psi)$ trivially, or in an arc whose boundary lies on $\{t\} \times D_r^n$.
- A heteroclinic component intersects with $Z'(\psi)$ to give two arcs running from 0 to $\{t\} \times D_r^n$ (one oriented forward, the other oriented backward) and a collection of arcs running from $\{t\} \times D_r^n$ back to $\{t\} \times D_r^n$.

- A homoclinic component intersects with $Z'(\psi)$ to give one arc running from 0 to $\{t\} \times D_r^n$ (oriented forward if the component leaves x , and backward if it approaches x) and a handful of arcs running from $\{t\} \times D_r^n$ to $\{t\} \times D_r^n$.

Now $\#\psi_{t,r}^{-1}(0)$ is the same as the oriented sum of points in the intersection of $Z'(\psi)$ with the submanifold $\{t\} \times D_r^n$. The intersection number of this submanifold with any arc which runs from $\{t\} \times D_r^n$ back to itself is zero, while the intersection with an arc running from 0 to $\{t\} \times D_r^n$ is $+1$ and the intersection with an arc running $\{\varepsilon\} \times D_r^n$ to 0 is -1 . This means isolated components contribute nothing to the count; heteroclinic components contribute $1 - 1 = 0$ to the count; and homoclinic components contribute either 1 or -1 to the count, depending on whether they leave x or approach x . \square

In particular, this explains why only six of the eight apparent orientations of the θ -graph are compatible with a cutout chart. What's more, this description of the signed count is *intrinsic to the stratified space*: if $Z(\psi)$ is 1-dimensional, we may determine whether or not $\partial Z(\psi)$ is compatibly oriented with the orientation on the interior by investigating the orientation on each arc, without ever passing to a cutout chart.

Corollary A.24. *If $Z(\psi_i)$ are 1-dimensional and $\phi : Z(\psi_1) \rightarrow Z(\psi_2)$ is an isomorphism of stratified-smooth spaces which is orientation-preserving on the interior, then ϕ is orientation-preserving on the boundary, as well.*

We can now use this to confirm that the boundary orientation may be intrinsically defined in general (so is independent of the choice of a cutout chart).

Lemma A.25. *Suppose $\psi_i : [0, \infty) \times \mathbb{R}^{n_i} \rightarrow \mathbb{R}^{m_i}$ are a pair of continuous maps which are smooth and have zero as a regular value on each stratum; further suppose $\phi : Z(\psi_1) \rightarrow Z(\psi_2)$ is an isomorphism of stratified-smooth spaces which is orientation-preserving on the top stratum. Then ϕ is also orientation-preserving on the boundary strata.*

Proof. We reduce to the 1-dimensional case. If (X, Δ) is a naive stratified-smooth space, and $f : X \rightarrow \mathbb{R}^m$ is a naive stratified-smooth map, we say f is *special* if $f^{-1}(0)$ is a stratified-smooth space (that is, it admits cutout charts). Notice that if $\phi : X \rightarrow Y$ is an isomorphism of naive stratified-smooth spaces, and $f : Y \rightarrow \mathbb{R}^m$ is special, then $f\phi$ is also special; ϕ induces an isomorphism $f^{-1}(0) \cong (f\phi)^{-1}(0)$, and the property of being a stratified-smooth space — that is, admitting cutout charts — is preserved under isomorphisms of naive stratified-smooth spaces.

Now consider the situation in the statement of the lemma. Pick $x \in \partial Z(\psi_1)$; we aim to show that $d\phi_x$ is orientation preserving as a map $T_x \partial Z_1 \rightarrow T_{\phi(x)} \partial Z_2$. Passing to a smaller open neighborhood of x if necessary, choose a stratified-smooth function $f : Z(\psi_2) \rightarrow \mathbb{R}^{n_2 - m_2}$ so that f induces an isomorphism $T_x \partial Z_2 \rightarrow \mathbb{R}^{n_2 - m_2}$. Explicitly, one may choose f to be a linear map restricted to $Z \subset [0, \infty) \times \mathbb{R}^{n_2}$. Because this function extends to a cutout chart, it follows that $f^{-1}(0) \subset Z(\psi_2)$ again carries the structure of a stratified-smooth space, so f is special. Call this zero locus $Z(\psi_2, f)$.

We thus have a stratification-preserving homeomorphism which is a diffeomorphism on each stratum $\phi_f : Z(\psi_1, f\phi) \rightarrow Z(\psi_2, f)$. Further, because ϕ was orientation-preserving and the construction defining

orientations of zero sets is natural, ϕ_f is orientation-preserving on the top (1-dimensional) stratum and is orientation-preserving on the boundary stratum if and only if ϕ is orientation-preserving at x .

But by Corollary A.24, we see that ϕ_f is indeed orientation-preserving on the boundary; thus ϕ is orientation-preserving at an arbitrary boundary point, as desired. \square

Now we are prepared to define the boundary of oriented stratified-smooth spaces and check the basic properties.

Definition A.26. Let (X, Δ) be an oriented stratified-smooth space. Its *boundary* is the disjoint union

$$\partial X = \bigsqcup_{d(s)=\dim X-1} \overline{X}_s,$$

where each X_s is oriented with the *boundary orientation* induced by a choice of cutout chart at an arbitrary point $x \in X_s$. \diamond

Note that this boundary orientation is well-defined by the previous lemma: it is independent of the choice of cutout chart. Because an orientation on a stratified-smooth space is induced from an orientation on $[0, \infty)^m \times \mathbb{R}^n$ via a cutout chart, the boundary operation on stratified-smooth spaces shares many formal properties with the usual boundary operator on oriented smooth manifolds with corners.

Proposition A.27. Let (X, Δ) be an oriented stratified space. Then the double boundary $\partial^2(X, \Delta)$ supports an orientation-reversing involution. More precisely,

$$\partial^2(X, \Delta) \cong \bigsqcup_{d(s)=\dim X-2} \overline{X}_s \sqcup -\overline{X}_s.$$

Proof. We saw that

$$\partial^2 X = \bigsqcup_{\substack{s \in \Delta \\ d(s)=\dim X-2}} \overline{X}_s \sqcup \overline{X}_s$$

as unoriented stratified-smooth spaces in the proof of Proposition A.15. We need to verify that these two copies of \overline{X}_s appear with the opposite orientation.

Because the definition of boundary orientation is local, it suffices to verify this claim for a given codimension-2 face of $Z(\psi)$. But the (iterated) boundary orientation on a given stratum of $Z(\psi)$ is induced in a uniform way by the (iterated) boundary orientation on the corresponding stratum of $[0, \infty)^m \times \mathbb{R}^n$, and the stated property holds for Euclidean space with corners. \square

We have an analogue of the usual result that the boundary of a compact oriented 1-manifold consists of a finite number of points whose signs sum to zero. The proof is essentially identical to [Mil19, Lemma 6.3] and to the argument opening [SS10, Section 2c], and it is an immediate application of Lemma A.22.

Proposition A.28. Let (X, Δ) be a compact connected oriented 1-dimensional stratified-smooth space. Then $\partial(X, \Delta)$ is a finite set of points which are canonically oriented with a sign ± 1 ; the sum of these signs is zero, and hence $\partial(X, \Delta)$ supports an orientation-reversing diffeomorphism.

Proof. First, Δ is a finite set, so there are finitely many 0-dimensional strata. Each bottom stratum is a closed set, hence compact because X is compact. Thus the 0-dimensional strata in Δ form a compact smooth manifold of dimension 0.

Now each zero-dimensional stratum has a cutout chart $(U_s, V_s, \psi_s, \phi_s)$. If we delete each neighborhood U_s , what remains is a compact 1-dimensional manifold with boundary $X' \subset X$, and it is well-known that the boundary of a 1-dimensional manifold with boundary consists of zero points, counted with sign.

To be more precise, $\partial X' = \sqcup_s \partial U_s$. Then to conclude, one needs to verify that the signed count of points in ∂U_s agrees with the signed count of points in X_s ; but (choosing the U_s appropriately) this is precisely the content of Lemma A.22. \square

We conclude with a discussion of orientations on fiber products.

Proposition A.29. *If X and Y are oriented stratified-smooth spaces and $f : X \rightarrow B$ and $g : Y \rightarrow B$ are stratified-smooth maps to an oriented smooth manifold B without boundary, so that f and g are stratified-transverse, then $X \times_B Y$ has the natural structure of an oriented stratified-smooth space and*

$$\partial(X \times_B Y) = (\partial X) \times_B Y \sqcup (-1)^{\dim X + \dim B} X \times_B (\partial Y).$$

Proof. Recall that the fiber-product of smooth manifolds is oriented with ‘ B in the middle’. Precisely, at the level of vector spaces, suppose $f_i : V_i \rightarrow W$ are transverse. Write $I_j = \text{im}(f_j)$ and $K_j = \ker(f_j)$; additionally write $I_1 \cap I_2 = I_{12}$.

Split $I_1 \cong I'_1 \oplus I_{12}$ while $I_2 \cong I_{12} \oplus I'_2$. Further split

$$V_1 \cong K_1 \oplus I_1 \cong K_1 \oplus I'_1 \oplus I_{12}$$

and

$$V_2 \cong I_2 \oplus K_2 \cong I_{12} \oplus I'_2 \oplus K_2.$$

Lastly, we may now split

$$W \cong I'_1 \oplus I_{12} \oplus I'_2.$$

If V_1, V_2 , and W are oriented, choose orientations on $I'_1, I'_2, I_{12}, K_1, K_2$ arbitrarily so that the three above isomorphisms are orientation-preserving; there are four ways to do so, say

$$\begin{aligned} & (o_{I'_1}, o_{I'_2}, o_{I_{12}}, o_{K_1}, o_{K_2}), \\ & (-o_{I'_1}, o_{I'_2}, -o_{I_{12}}, o_{K_1}, -o_{K_2}), \\ & (o_{I'_1}, -o_{I'_2}, -o_{I_{12}}, -o_{K_1}, o_{K_2}), \\ & (-o_{I'_1}, -o_{I'_2}, o_{I_{12}}, -o_{K_1}, -o_{K_2}). \end{aligned}$$

Now orient $V_1 \times_W V_2 \cong K_1 \oplus I_{12} \oplus K_2$ so that this isomorphism is orientation-preserving; this defines the same orientation on the fiber product in all four cases above, so is well-defined. Applying this construction to the appropriate tangent spaces gives rise to our fiber product orientation on $X \times_B Y$.

The cutout charts constructed in the proof of Proposition A.17 are not orientation-preserving. However, we can analyze the difference from the expected orientation.

Suppose the cutout chart used for X in the above proposition is $\psi : V \rightarrow \mathbb{R}^\ell$, and write $b = \dim B, y = \dim Y$; a straightforward and tedious linear algebra argument shows that the orientation of the cutout chart in Proposition A.17 disagrees with the fiber product orientation by a factor of $(-1)^{\ell b + \ell y + by}$, and in particular this is consistent over the entire cutout chart.

It follows that the fiber product orientation on $X \times_B Y$ does indeed give the fiber product the structure of an oriented stratified-smooth space; one may simply flip the orientation on the cutout charts constructed before, if necessary, to make them orientation-preserving.

The desired boundary orientation follows from a tedious definition chase and the above sign computation; notice that it agrees with the expected boundary orientations on a fiber product when this is a transverse fiber product of smooth manifolds with boundary. \square

Remark A.30. Suppose the maps $X \rightarrow B$ and $Y \rightarrow B$ are stratumwise submersions. Then the coordinate-swapping diffeomorphism $\phi : X \times_B Y \rightarrow Y \times_B X$ has sign $(-1)^{(x-b)(y-b)}$, with $x = \dim X$ and so on. Because this is computed in the top stratum, this result follows from the corresponding fact for smooth manifolds, which is straightforward from the definition of fiber product orientation. \diamond

We conclude with a similar computation for the final proposition of Appendix A.1.

Proposition A.31. *Suppose (X, Δ) is an oriented stratified-smooth space and $L : X \rightarrow \mathbb{R}$ a stratified-smooth map, with c a regular value. Then $X^{\leq c}$ is also an oriented stratified-smooth space, with boundary*

$$\partial X^{\leq c} = (-1)^{\dim X - 1} L^{-1}(c) \sqcup (\partial X)^{\leq c},$$

where $L^{-1}(c)$ is oriented as a regular zero set.

Proof. The orientations on the top strata are inherited as an open subset of X . Recall that we orient zero sets $Z(\psi)$ for $\psi : V \rightarrow \mathbb{R}^\ell$ by the rule that $TZ \oplus \mathbb{R}^\ell \cong TV$.

It is straightforward linear algebra and definition-pushing to see that the cutout charts constructed in Proposition A.19 are orientation-preserving with respect to the existing orientation on the top stratum of X ; this amounts to saying that

$$(TZ \oplus \mathbb{R}^\ell) \oplus \mathbb{R} \cong TV \oplus \mathbb{R}$$

is orientation-preserving, where the map is upper-triangular and the identity on the bottom-right entry, and we know that the top-left entry is orientation-preserving.

In computing the sign on the boundary stratum $L^{-1}(c)$, the additional factor arises because boundaries are oriented with an outward-normal-first convention, while zero sets place the d/dt factor as the final basis element; comparing them requires commuting a factor of \mathbb{R} across $TL^{-1}(c)$, introducing the stated sign. \square

A.3 The real blowup of a stratified-smooth space

In Section 5 of the main text, we will find ourselves in the following situation.

Suppose we have a stratified-smooth space P , a vector bundle $q : E \rightarrow M$ over a smooth manifold, and a stratified-smooth map $\psi : P \rightarrow E$ transverse to the zero section. We will want to construct, in a canonical way, a stratified-smooth space P' with a map $\psi' : P' \rightarrow S(E)$, where $S(E)$ is the unit sphere bundle of E . Over the locus $P \setminus \psi^{-1}(0)$, the map $\psi/\|\psi\|$ gives a canonical such lift; however, this lift cannot be extended over the locus where $\psi = 0$.

To remedy this, we can replace our stratified-smooth space with its *real blowup* (or simply ‘blowup’) along $\psi = 0$; the definition of this space furnishes it with a canonical such lift.

In this section, we first define this blowup construction and verify that the result is again a stratified-smooth space. After that, we define an equivalence relation on stratified-smooth spaces with maps to M so that a stratified-smooth space is equivalent to any of its real blowups. We conclude by verifying some technical facts which we will make use of in Sections 5.3–5.5.

Construction A.32. Let P be an oriented stratified-smooth space. Let M be a closed smooth manifold equipped with an oriented vector bundle $q : E \rightarrow M$, and a stratified-smooth map $\psi : P \rightarrow E$ transverse to the zero section. Give E a fiberwise metric; its only purpose is to pick out a sphere bundle of E .

We define $B(\psi) \subset P \times [0, \infty) \times S(E)$ to be the subset given by

$$B(\psi) = \{(p, t, v) \mid \psi(p) = tv\}.$$

Both maps $\pi_1 : B(\psi) \rightarrow P$ and $\pi_3 : B(\psi) \rightarrow S(E)$, given by projection to first and last coordinate, will be important to us.

Let us give $B(\psi)$ the structure of a stratified-smooth space.

If P has face poset Δ , then call $B(\psi)$ ’s face poset $\Delta_B(\psi) \subset \Delta \times \{-1, 0\}$. To be explicit about this, we have a subset $\Delta_Z(\psi) \subset \Delta$ given by the faces on which $\psi^{-1}(0) \cap P_s \neq \emptyset$. Then $B(\psi)$ ’s face poset is

$$\Delta_B(\psi) = \Delta_Z(\psi) \times -1 \sqcup \Delta \times 0,$$

with grading function $d_\psi(s, -1) = d(s) - 1$ and $d_\psi(s, 0) = d(s)$, and with order induced from $\Delta \times \{-1, 0\}$ with the lexicographic order.

For any $s \in \Delta$, the strata $(s, 0)$ are given by

$$B(\psi)_{s,0} = \pi_1^{-1}(P_s \setminus \psi^{-1}(0)).$$

Being the inverse image of locally closed subsets by a continuous map, these are locally closed; a point in this stratum takes the form $(p, \|\psi(p)\|, \psi(p)/\|\psi(p)\|)$ for $p \in P_s$, and we give $B(\psi)_{s,0}$ the unique smooth structure so that $\pi_1 : B(\psi)_{s,0} \rightarrow P_s \setminus \psi^{-1}(0)$ is a diffeomorphism.

On the other hand, if $s \in \Delta_Z(\psi)$, then $B(\psi)_{s,-1} = \pi_1^{-1}(Z(\psi)_s)$; a point in this stratum takes the form $(p, 0, v)$ for $p \in \psi^{-1}(0) \cap P_s = Z(\psi)_s$ and any $v \in S(E_{q\psi(p)})$. In this case, π_3 gives $B(\psi)_{s,-1}$ the structure of a sphere bundle over $Z(\psi)_s$, canonically identified with the pullback of the sphere bundle

$E \rightarrow M$ under $f|_{Z(\psi)_s}$. We give $B(\psi)_{s,-1}$ the unique smooth structure so that this identification is a diffeomorphism.

Because

$$\dim Z(\psi)_s = d(s) - \text{rank}(E),$$

and the sphere bundle has fiber of dimension $\text{rank}(E) - 1$, we see that $\dim B(\psi)_{s,0} = d(s) - 1 = d_\psi(s, 0)$.

It is straightforward to confirm that

$$\begin{aligned} \overline{B(\psi)_{s,-1}} &= \bigcup_{\substack{t \in \Delta_Z \\ t \leq s}} B(\psi)_{t,-1}, \\ \overline{B(\psi)_{s,0}} &= \bigcup_{\substack{i \in \{-1,0\} \\ t \in \Delta_Z \\ t \leq s}} B(\psi)_{t,i}. \end{aligned}$$

Because we have an isomorphism of naive stratified-smooth spaces $\pi_1 : B(\psi)_{\geq s,0} \cong P_{\geq s}$, a cutout chart around any $(p, r, v) \in B(\psi)_{s,0}$ is induced by a choice of cutout chart around $p \in P_s$ (simply lift it along the isomorphism π_1).

Cutout charts around $(p, 0, v) \in B(\psi)_{s,-1}$ are more intricate. Suppose $(p, 0, v) \in B(\psi)_{s,-1}$. Pick a cutout chart

$$(U \subset P, V, f : V \rightarrow \mathbb{R}^k, \phi : U \rightarrow Z(f))$$

for $p \in P_s$. We may choose U small enough so that $\psi(U)$ lands inside a trivializing chart for the vector bundle E ; trivializing the bundle and setting $m = \dim M, e = \text{rank}(E)$, we may write ψ as a map

$$\psi : Z(f) \rightarrow \mathbb{R}^m \times \mathbb{R}^e.$$

By definition of stratified-smooth map, we may choose this cutout chart so that the section ψ extends to a map $\hat{\psi} = (\sigma, g) : V \rightarrow \mathbb{R}^m \times \mathbb{R}^e$.

Now choose a small chart $j : \mathbb{R}^{e-1} \rightarrow S^{e-1} \cong E_{q\psi(p)}$ around v . Recall here that $q : E \rightarrow M$ is the bundle projection.

Our cutout chart for $B(\psi)$ around p has domain

$$V' = V \times [0, \varepsilon) \times \mathbb{R}^{e-1}$$

and smooth map

$$h(x, t, v) = (f(x), g(x) - tj(v)) \in \mathbb{R}^k \times \mathbb{R}^e.$$

The map $\phi' : h^{-1}(0) \rightarrow B(\psi)$ is given by

$$(x, t, v) \mapsto (\phi^{-1}(x), t, j(v)) \in B(\psi) \subset P \times [0, \infty) \times S(E);$$

this is a homeomorphism onto its image, the set of $(t, x, v) \in B(\psi)$ with $\|\psi(x)\| < \varepsilon$ and $x \in U$, which is open. Because this map ϕ' also induces a bijection between the natural stratifications, it gives a cutout chart near $(p, 0, v)$.

Now orient $B(\psi)$ is oriented so that the map $\pi_1 : B(\psi) \rightarrow P$ is an orientation-preserving diffeomorphism off of Z . It is easy to see the cutout charts constructed above are orientation-preserving, so that this gives $B(\psi)$ the structure of an oriented stratified-smooth space. \diamond

In the situation above, (P, ϕ) and $(B(\psi), \phi\pi_1)$ are both stratified-smooth spaces equipped with stratified-smooth maps to M ; if P is connected, then $B(\psi)$ is as well. For the sake of our technical arguments in the main text, we wish to construct an equivalence relation on connected stratified-smooth spaces mapping to M so that these are considered equivalent. To get there — and to define the geometric chain complex in the next section — we need a handful of definitions.

Definition A.33. Let M be a smooth manifold. The set of *stratified-smooth probes* into M , written $\mathcal{P}(M)$, is the set of pairs (P, ϕ) , where P is a compact connected oriented stratified-smooth space, $\phi : P \rightarrow M$ is a stratified-smooth map, and we consider $(P, \phi) = (Q, \eta)$ if there exists an orientation-preserving isomorphism $f : P \rightarrow Q$ with $\eta = \phi f$. This set carries an involution, $\iota(P, \phi) = (-P, \phi)$, orientation-reversal of P . We call a probe which is fixed under this involution *trivial*.

We say that a probe (P, ϕ) is *degenerate* if $\dim P > 0$ and, for all strata P_s of dimension $d(s) > 0$ and all $p \in P_s$, we have

$$\text{rank}(d\phi_s : T_p P_s \rightarrow M) < d(s).$$

We write $\mathcal{D}(M) \subset \mathcal{P}(M)$ for the set of degenerate probes in M ; this subset is ι -invariant. \diamond

Remark A.34. There are two other notions of degenerate probe in the literature, defined using: the *small image* condition of [Lip14] and the *small rank* of [GF21]. The first says that a probe is small rank if its image can be covered by the image of a probe of smaller dimension. For the latter, the ‘small rank’ condition is that $d\phi$ has non-maximal rank on the top stratum. Both authors then say that a degenerate probe is a probe of small image/rank whose boundary is the union of a trivial probe and one of small image/rank.

Our definition of ‘degenerate probe’ includes fewer probes (for instance, the projection map $[0, 1]^2 \rightarrow [0, 1]$ is not degenerate according to our definition, whereas it is for both of these authors), and in particular our $C_*^{gm}(M)$ will usually be supported in infinitely many degrees; for these authors, it is supported in degrees $*$ $\in [0, \dim M + 1]$. However, when analyzing the notion of collapse-equivalence below, it is difficult to work with the condition that the boundary be a sum of a small rank probe and a trivial probe. The definition chosen above is better suited for studying collapse-equivalence. \diamond

We will now define *collapse-equivalence* of stratified-smooth spaces (though we will only apply it under the assumption that the given stratified-smooth spaces are connected), and then verify that (P, ϕ) and $(B(\psi), \phi\pi_1)$ are collapse-equivalent.

Definition A.35. Let (P, Δ) be a stratified-smooth space. We say a *stratified-smooth subspace of codimension k* is a stratified-smooth space (Z, Δ_Z) , an order-preserving injection $\bar{j} : \Delta_Z \rightarrow \Delta$ with $d(\bar{j}(s)) = d_Z(s) + k$, and a topological embedding $j : Z \rightarrow P$ so that j restricts on each stratum Z_s to a smooth embedding

$$j : Z_s \rightarrow P_{\bar{j}(s)}.$$

We identify Z with its image $j(Z)$ in the following discussion.

Let $(P, \phi) \in \mathcal{P}(M)$. We say that a closed subset $P_{SL} \subset P$ is a *small locus* if $P_{SL} = P_{\text{codim}>0} \cup P_D$, where $P_{\text{codim}>0}$ is the union of finitely many codimension $k > 0$ stratified-smooth subspaces of P (not necessarily compact), and $P_D \subset \partial P$ is the union of degenerate strata of positive codimension.

Suppose there exist small loci $P_{SL} \subset P, Q_{SL} \subset Q$, and an orientation-preserving diffeomorphism of stratified smooth spaces $f : P \setminus P_{SL} \rightarrow Q \setminus Q_{SL}$ so that $f\phi = \eta$. Then we say (P, ϕ) and (Q, η) are *collapse-equivalent*, and write this as $(P, \phi) \sim_C (Q, \eta)$. \diamond

This is indeed an equivalence relation, and plays well with the notion of degeneracy and boundary operators.

Lemma A.36. *Collapse-equivalence gives an equivalence relation \sim_C on $\mathcal{P}(M)$. If (P, ϕ) is degenerate, then so is (Q, η) . If $(P, \phi) \sim_C (Q, \eta)$ then $\iota(P, \phi) \sim_C \iota(Q, \eta)$ and $\partial(P, \phi) \sim_C \partial(Q, \eta) + (R, \sigma_D)$, where (R, σ_D) is degenerate.*

Proof. Reflexivity and symmetry are obvious, as is the claim about orientation-reversal.

For transitivity, observe that the union of two small loci is again a small locus; if $(P, \phi) \sim_C (Q, \eta) \sim_C (R, \psi)$ and we have small loci $P_{SL} \subset P, Q_{SL}^1, Q_{SL}^2 \subset Q$ and $R_{SL} \subset R$ so that

$$\begin{aligned} f : (P \setminus P_{SL}, \phi) &\cong (Q \setminus Q_{SL}^1, \eta) \\ g : (Q \setminus Q_{SL}^2, \eta) &\cong (R \setminus R_{SL}, \psi), \end{aligned}$$

Each of these small loci is written as a union $P_{SL} = P_{\text{codim}>0} \cup P_D$ of positive-codimension subspaces and degenerate strata. Write

$$P'_{\text{codim}>0} = P_{\text{codim}>0} \cup f^{-1}(Q_{\text{codim}>0}^2) \quad R'_{\text{codim}>0} = R_{\text{codim}>0} \cup g(Q_{\text{codim}>0}^1);$$

take the locus P'_D to be the union of P_D and the strata intersecting $f^{-1}(Q_D^2)$, similarly for R'_D . Notice that the components meeting $f^{-1}(Q_D^2)$ are degenerate: those which aren't already in P_D have small rank on the open dense set $(P \setminus P_{SL})_s$ of each positive-dimension stratum, so by semicontinuity of rank they have small rank on each positive-dimension stratum. So P'_D and similarly R'_D are both unions of degenerate strata of P and R .

Then $P'_{SL} = P'_{\text{codim}>0} \cup P'_D$ is a small locus, as is $R'_{SL} = R'_{\text{codim}>0} \cup R'_D$. Finally, fg gives an isomorphism $(P \setminus P'_{SL}, \phi) \cong (R \setminus R'_{SL}, \psi)$. This proves transitivity of \sim_C .

As implicitly mentioned above, degeneracy off of a small locus implies degeneracy on the whole probe, so collapse-equivalence preserves degeneracy.

The claim about boundaries is also clear: for all strata not contained in the small locus, the isomorphism $f : P \setminus P_{SL} \rightarrow Q \setminus Q_{SL}$ descends to an isomorphism $\overline{P}_s \setminus P_{SL} \rightarrow \overline{Q}_{f(s)} \setminus Q_{SL}$. Any strata contained in the small locus are degenerate, by definition of small locus. So f provides a collapse-equivalence between the non-degenerate parts of ∂P and the non-degenerate parts of ∂Q , and the remaining pieces are collapse-equivalent and form the chain (R, σ_D) mentioned in the statement of the proposition. \square

Henceforth, we write $\mathcal{Q}(M) = \mathcal{P}(M) / \sim_C$. We write $\mathcal{D}_Q(M) = \mathcal{D}(M) / \sim_C \subset \mathcal{Q}(M)$.

Lemma A.37. *Let $\phi : P \rightarrow M$ be an oriented stratified-smooth probe equipped with a lift $\psi : P \rightarrow E$ of ϕ to a section of a vector bundle $q : E \rightarrow M$, which is stratified-transverse to the zero section. Then the real blowup $B(\psi)$ satisfies the following properties:*

- (i) *$B(\psi)$ is an oriented stratified-smooth probe in $S(E)$ via the projection map $\pi_3 : B(\psi) \rightarrow S(E)$, with boundary relation*

$$\partial(B(\psi), \pi_3) = (B(\psi|_{\partial P}), \pi_3) - (-1)^{\dim Z_\psi} (Z_\psi \times_B S(E), \pi_2).$$

- (ii) *When composed with the projection map $q : S(E) \rightarrow M$, the blowup defines a probe $(B(\psi), q\pi_3) = (B(\psi), \phi\pi_1)$ in M which is collapse-equivalent to $(P, q\psi)$.*

Proof. The boundary relation may be readily off from the description of the cutout charts.

We defined collapse-equivalence so that this would be true. The projection map $\pi_1 : B(\psi) \rightarrow P$ is an isomorphism away from $Z = \psi^{-1}(0) \subset P$ (a small locus), and $(\pi^{-1}(Z), \phi\pi_1)$ is a subset of $\partial B(\psi)$ which factors through the fiber bundle projection $\pi_1^{-1}(Z) \rightarrow Z$ with positive-dimensional fibers. It follows that every face of $(\pi^{-1}(Z), \phi\pi_1)$ is small rank, and hence this is a degenerate probe. Thus $(B(\psi), \phi\pi_1)$ is collapse-equivalent to (P, σ) . \square

In the main text, we will need a variation on this which identifies two stratified-smooth probes related to this blowup construction. We include this construction and lemma in this section, as it's in the spirit of the work carried out above.

Construction A.38. Let (P, σ) and (Q, η) be transverse stratified-smooth probes mapping submersively to B , equipped with lifts $\psi : P \rightarrow E$ and $\phi : Q \rightarrow E$ lifting σ and η , respectively (where $E \rightarrow B$ is an oriented vector bundle of rank $\text{rk}(E) > 1$). Assume further that $(p, t, q) \mapsto (1 - t)\psi(p) - t\phi(q)$ is transverse over $P \times [0, 1] \times Q$ to the zero section.

We write $Z_{\psi\phi} \subset P \times_B ([0, 1] \times Q)$ for the zero set of $(p, t, q) \mapsto (1 - t)\psi(p) - t\phi(q)$; by the assumption that this is a transverse zero set, $Z_{\psi\phi}$ is an oriented stratified-smooth space. If P and Q are furthermore equipped with maps to A and C , respectively, $Z_{\psi\phi}$ defines a probe in $A \times C$.

There is another way to use these maps to construct a stratified-smooth probe in $A \times C$. We write $B(\psi)$ and $B(\phi)$ for the blowups, considered as probes in $A \times S(E)$ and $S(E) \times C$, and consider the probe $B(\psi) \times_{S(E)} B(\phi)$ in $A \times C$. This fiber product is clearly transverse over blowup strata (which map submersively onto the fibers of $S(E)$), and as will be seen in the proof below, on the open strata the fact that this is a transverse fiber product is equivalent to the statement that $(1 - t)\psi - t\phi$ is transverse to zero. \diamond

In Section 5, we will see stratified-smooth chains of the form $Z_{\psi\phi}$ naturally appear from analytic considerations, whereas stratified-smooth chains of the form $B(\psi) \times_{S(E)} B(\phi)$ appear in the same context from algebraic considerations. To square these two perspectives, we will need to see that these two stratified-smooth chains are collapse-equivalent.

Lemma A.39. *In the situation of the above construction, the fiber product $B(\psi) \times_{S(E)} B(\phi)$ is collapse-equivalent to $Z_{\psi\phi}$ as probes in $A \times C$. If B and E are both even-dimensional, then this collapse-equivalence is orientation-preserving (otherwise there is an overall sign).*

Proof. One should intuit this claim as follows. On the open stratum where $\psi, \phi \neq 0$, the space $B(\psi) \times_{S(E)} B(\phi)$ is diffeomorphic to the subset of the fiber product $P \times_B Q$ of those pairs (p, q) so that $\psi(p)$ and $\phi(q)$ are positive scalar multiples of one another. On the other hand, supposing $\psi(p), \phi(q)$ are both nonzero, there exists a unique t with $(1-t)\psi(p) = t\phi(q)$ if and only if the two vectors $\psi(p)$ and $\phi(q)$ are positive scalar multiples of one another.

Indeed, this suggests a reasonably simple map from one to the other. Points in the fiber product $B(\psi) \times_{S(E)} B(\phi)$ are sextuples (p, t, v, q, s, w) with

$$\sigma(p) = \eta(q), \quad \psi(p) = tv, \quad \phi(q) = sw, \quad v = w.$$

We would like to map $F : B(\psi) \times_{S(E)} B(\phi) \rightarrow Z_{\psi\phi}$ via

$$f(p, t, v, q, s, w) = \left(p, \frac{t}{s+t}, q\right).$$

Notice that because the maps to A and C are defined by projecting to $P \times Q$, and F preserves the projection to $P \times Q$, the map F preserves the map to $A \times C$.

However, F is not defined everywhere, and is not surjective.

The map F is ill-defined on is the set where $s = t = 0$; these are the points in the blowup corresponding to $\psi = \phi = 0$. This is a union of degenerate strata

$$(Z_\psi \times_B S(E)) \times_{S(E)} (S(E) \times_B Z_\phi) = Z_\psi \times_B S(E) \times_B Z_\phi$$

in the domain. This is degenerate because each stratum takes the form $Z_s \times_B S(E) \times_B Z_t$, which factors through the projection to $Z_s \times_B Z_t$ with positive-dimensional closed fibers, and hence has small rank on every stratum.

Next, F misses the codimension $\text{rank}(E)$ subset of $Z_{\psi\phi}$ defined by $t\psi - (1-t)\phi = 0$; this condition implies that $\psi = \phi = 0$, because the points in $Z_{\psi\phi}$ have $(1-t)\psi - t\phi = 0$ to begin with; this subset can be identified with $Z_\psi \times_B [0, 1] \times Z_\phi$.

Thus F defines an isomorphism away from small loci in the domain and codomain, which preserves the map to $A \times C$, and hence is a collapse-equivalence.

What remains is a discussion of orientations. We will freely use global symbols (such as M , or B) to mean small open subsets of these as appropriate; because everything maps submersively to B and B is even-dimensional, we may effectively ignore it by passing to fibers, so we suppress it below, writing E for an oriented vector space. We orient $S(E)$ with the outward-normal-first convention, so that the scaling map $\mathbb{R}_+ \times S(E) \cong E - 0$ is an orientation-preserving isomorphism.

An open locus in $B(\psi) \times_{S(E)} B(\phi)$ may be identified with

$$P^{\langle v \rangle} \times S(E) \times Q^{\langle v \rangle},$$

where the superscript indicates that we take a fiber above $\mathbb{R}^+v \subset E$, where v is fixed; in the first case, we orient this so that the natural locally-defined map $P^{\langle v \rangle} \times S(E) \cong P$ is an oriented diffeomorphism; in the second case $S(E) \times Q^{\langle v \rangle} \cong Q$, as is our convention for orientations on the fiber product. Then an open subset of the fiber product is identified with

$$B(\psi) \times_{S(E)} B(\phi) \sim P^{\langle v \rangle} \times S(E) \times Q^{\langle v \rangle}.$$

On the other side, there is a natural identification of $Z_{\psi\phi}$ with the zero locus of the map $(p, t, q) \mapsto (1-t)\psi(p) - t\phi(q)$ from $P \times [0, 1] \times Q$ to E . Using the identifications above near a point (p, t, q) with $\psi(p)$ and $\phi(q)$ positive scalar multiples of $v \in S(E)$, we write this as the zero set of a map on $P^{\langle v \rangle} \times S(E) \times [0, 1] \times S(E) \times Q^{\langle v \rangle}$. Further, we identify a neighborhood of $v \in S(E)$ with a neighborhood of zero in $\langle v \rangle^\perp$, oriented so that $\langle v \rangle \oplus \langle v \rangle^\perp \cong E$. If $\psi(p) = v_p$ and $\psi(q) = v_q$ with $(1-t)v_p = t_0v_q$, we can identify of $Z_{\psi\phi}$ near (p, t_0, q) as the zero set of the map

$$P^{\langle v \rangle} \times \langle v \rangle^\perp \times (0, 1) \times \langle v \rangle^\perp \times Q^{\langle v \rangle} \rightarrow E \quad \text{given by} \quad (p, w, t, u, q) \mapsto (1-t)w - tu + (t-t_0)v_p + (t_0-t)v_q.$$

For now, we ignore the outside terms and focus on the zero set, considered as a function on $\langle v \rangle^\perp \times (0, 1) \times \langle v \rangle^\perp$.

The first part of the given expression lies in $\langle v \rangle^\perp$, while the latter lies in $\langle v \rangle$; they both must be zero. Because v_p and v_q are positive scalar multiples of one another, $-(t-t_0)(v_p + v_q)$ is only zero if $t = t_0$, and then the first equation gives $u = \frac{1-t_0}{t_0}w$, so by projection to the w -factor, this zero set is diffeomorphic to $\langle v \rangle^\perp$.

To justify that this projection is an oriented diffeomorphism, we check that for fixed w_0 , the map

$$(t, u) \mapsto (1-t)w_0 - tu - (t-t_0)(v_p + v_q)$$

is an oriented diffeomorphism $(0, 1) \times \langle v \rangle^\perp \rightarrow E$ between open subsets. To see this, take the derivative; if (t_0, u_0) is fixed and $\{u_2, \dots, u_n\}$ is an oriented basis for $\langle v \rangle^\perp$, at (t_0, u_0) the maps the oriented basis $\{t, u_2, \dots, u_n\}$ to $\{-u_0 - w_0 - v_p - v_q, -tu_2, \dots, -tu_n\}$; this is equivalent to the oriented basis $\{-v, -u_2, \dots, -u_n\}$, which is equivalent to $(-1)^{\text{rank}(E)}$ the standard basis. By the assumption that E has even rank, this proves the claim that $Z_{\psi\phi}$ is locally oriented diffeomorphic to $P^{\langle v \rangle} \times \langle v \rangle^\perp \times Q^{\langle v \rangle}$.

Finally, with respect to this description in the domain and codomain, the map F is isotopic to the identity. Thus F is orientation-preserving, as desired. \square

A.4 Geometric homology from stratified-smooth spaces

In this section, we use the set of stratified-smooth probes modulo collapse-equivalence to define a chain complex associated to smooth manifolds without boundary which is natural under smooth maps. Using a variation on the Eilenberg-Steenrod axioms, we will verify that the resulting homology theory is naturally isomorphic to singular homology on closed manifolds (or, more precisely, on pairs (X, A) which are the interiors of a handlebody-pair).

We first outline the construction of the underlying graded group of our chain complex.

Construction A.40. Write $Q_d(X) = P_d(X)/\sim_C$ for the set of d -dimensional probes in X modulo collapse-equivalence.

Set

$$K_d(X) = \mathbb{Z}[Q_d(X)]/\mathbb{Z}[Q_d(X)]^\iota.$$

The purpose of collapsing the ι -invariant subgroup is to set $\iota(P, \phi) = -(P, \phi)$, so that orientation-reversal corresponds to negation; however, quotienting by this more naive relation would introduce 2-torsion into the complex, corresponding to probes which support orientation-reversing isomorphisms. Collapsing the entire ι -invariant subgroup avoids this issue.

If one writes $T_d(X) = Q_d(X)^\iota$ for the ‘trivial probes’ (those which admit an orientation-reversing isomorphism), then $K_d(X)$ can be identified with the free abelian group on $(Q_d \setminus T_d)/\iota$.

Next, collapse the subgroup of $K_d(X)$ spanned by the image of the degenerate probes. The degenerate probes $D(X) \subset P(X)$ define a \sim_C -invariant subset; we define the quotient by this relation to be $D_Q(X) \subset Q(X)$. We set

$$C_d^{gm}(X) = K_d(X)/\langle D_Q(X) \rangle.$$

We call this the complex of *geometric chains*, and an element a geometric chain. Again, the group $C_d^{gm}(X)$ is free abelian: one may give it a basis in bijection with $(Q_d \setminus (T_d \cup (D_Q)_d))/\iota$.

Then $C_*^{gm}(X)$ is the underlying graded abelian group of what we will call the geometric chain complex. ◇

To justify the name ‘geometric chain complex’, we should construct a well-defined boundary operator which squares to zero.

Lemma A.41. *The operation $\partial(P, \phi) = \sum_{d(s)=\dim P-1} (\overline{P}_s, \phi)$ sending a stratified-smooth probe to its geometric boundary gives a well-defined map $\partial : C_d(X) \rightarrow C_{d-1}(X)$ with $\partial^2 = 0$.*

Proof. To be slightly more precise, recall that the faces P_s are not assumed to be connected, whereas probes are; when we write (\overline{P}_s, ϕ) above, we mean the (finite) sum over its connected components. The compactness of P guarantees that the stratification poset Δ is finite, and the compactness of each \overline{P}_s guarantees there are only finitely many such components, so this is indeed a finite sum. As for orientations, each \overline{P}_s is equipped with its boundary orientation.

It is clear that ∂ defines a map $\mathbb{Z}[P_d(X)] \rightarrow \mathbb{Z}[P_{d-1}(X)]$; because $\partial\iota(P, \phi) = \iota\partial(P, \phi)$, this descends to a map on $\mathbb{Z}[P]/\mathbb{Z}[P]^\iota$.

Before moving on, we explain why the boundary of a degenerate chain is degenerate, or rather, zero in

$$\mathbb{Z}[P_{d-1}]/\mathbb{Z}[P_{d-1}]^\iota + \mathbb{Z}[D].$$

For $d > 1$, this is obvious: a degenerate chain is defined to be a chain of dimension $d > 0$ so that each stratum P_s of dimension $d(s) > 0$ is degenerate. When $d = 1$, this is slightly more subtle.

First, observe that the assumption ‘ ϕ has small rank’ means that $d\phi = 0$ on each connected component of a 1-dimensional stratum, and hence ϕ is constant on each such stratum. It follows that ϕ is constant on

the closure of each connected component of the 1-dimensional stratum; because P is connected and the 0-dimensional strata lie in the closure of the 1-dimensional strata, it follows that ϕ is also constant on the 0-dimensional stratum and thus that the constant value ϕ takes on different components all coincide. So ϕ is constant. So by Proposition A.28, the boundary $\partial(X, \phi)$ is an even number of points all mapping to the same point, whose signed sum is zero. So $\partial(X, \phi) = 0$ in $\mathbb{Z}[\mathbb{P}_0]/\mathbb{Z}[\mathbb{P}_0]^t$.

Next, the boundary map mentioned above does *not* descend naively to a well-defined map on $\mathbb{Z}[\mathbb{Q}]$, because the boundary of collapse-equivalent chains need not be collapse-equivalent; by Lemma A.36 all we know is that if $\phi \sim_C \eta$, then $\partial\phi \sim \partial\eta + \sigma_D$, where σ_D is a degenerate probe.

Still, this is enough to see that ∂ descends to an operation on C_* . For each equivalence class of nondegenerate probe $[(P, \phi)] \in \mathbb{Q}$, we choose a representative (P, ϕ) , and send

$$\partial[(P, \phi)] = \sum_{d(s)=\dim P-1} [(\overline{P_s}, \phi)];$$

this is well-defined regardless of the choice of representative because degenerate chains are identically zero in C_* , and if $(P, \phi) \sim_C (Q, \eta)$, then $\partial(P, \phi) \sim \partial(Q, \eta) + (R, \sigma_D)$, where (R, σ_D) is degenerate.

Finally, the fact that this operation has $\partial^2 = 0$ is Proposition A.27. \square

So for each smooth manifold without boundary X , we obtain a chain complex $C_*^{gm}(X)$ and a corresponding homology group $H_*^{gm}(X)$. This defines a functor on the category of smooth manifolds: if $f : X \rightarrow Y$ is a smooth map, simply set $f_*(P, \phi) = (P, f\phi)$. It is obvious both that this is functorial and that it respects the boundary operator.

We would like to extend this to a functor on the category of manifold pairs, which enjoys a relative long exact sequence on homology. This requires a technical lemma.

Lemma A.42. *When $i : A \rightarrow X$ is an injective immersion, the map $C_*^{gm}(A) \rightarrow C_*^{gm}(X)$ is injective.*

Proof. First we check that this map is injective on \mathbb{P} , then on $\mathbb{P} \setminus \mathbb{D}$, then on $\mathbb{Q} \setminus \mathbb{D}_Q$.

If $(P, \phi), (Q, \eta) \in \mathbb{P}(A)$ and $f : P \rightarrow Q$ is an isomorphism so that $i\eta = i\phi f$. Because $i : A \rightarrow X$ is injective, we also have $\eta = \phi f$. Thus if $(P, i\phi) = (Q, i\eta)$ then in fact $(P, \phi) = (Q, \eta)$. So the map $\mathbb{P}(A) \rightarrow \mathbb{P}(X)$ is injective. An identical argument shows that the induced map $\mathbb{Q}(A) \rightarrow \mathbb{Q}(X)$ is injective as well.

Suppose $(P, \phi) \in \mathbb{P}(A)$ is such that $(P, i\phi)$ is degenerate. Because $i\phi$ has small rank on every stratum of P of positive dimension, and di is injective, ϕ also has small rank on every stratum of P . Thus $i_*\mathbb{P}(A) \cap \mathbb{D}(X) = i_*\mathbb{D}(A)$, and thus the map i_* restricts to a well-defined injective map

$$\mathbb{P}(A) \setminus \mathbb{D}(A) \rightarrow \mathbb{P}(X) \setminus \mathbb{D}(X).$$

Finally, by Lemma A.36, we also have $i_*\mathbb{Q}(A) \cap \mathbb{D}_Q(X) = i_*\mathbb{D}_Q(A)$, and hence i_* gives a well-defined injection

$$\mathbb{Q}(A) \setminus \mathbb{D}_Q(A) \rightarrow \mathbb{Q}(X) \setminus \mathbb{D}_Q(X).$$

It remains to quotient by the involution, but again by injectivity of i , the probe (P, ϕ) in A will admit an orientation-reversing isomorphism if and only if $(P, i\phi)$ does. So the map

$$[\mathbb{Q}(A) \setminus (D_Q(A) \cup \mathbb{T}(A))] \rightarrow [\mathbb{Q}(X) \setminus (D_Q(X) \cup \mathbb{T}(X))]$$

is injective; the same is true after quotienting by the (now fixed-point-free) involution ι , and hence the map $C_*^{gm}(A) \rightarrow C_*^{gm}(X)$ is injective (it induces an injection on a basis). \square

Suppose $A \subset X$ is an open submanifold. We may thus define a homology group

$$C_*^{gm}(X, A) = C_*^{gm}(X) / i_* C_*^{gm}(A),$$

and this is obviously natural for maps of manifold pairs $(X, A) \rightarrow (Y, B)$. Further, because the map $C_*^{gm}(A) \rightarrow C_*^{gm}(X)$ is injective, the standard homological algebra shows that there exists a long exact sequence

$$\dots \xrightarrow{[-1]} H_*^{gm}(A) \rightarrow H_*^{gm}(X) \rightarrow H_*^{gm}(X, A) \xrightarrow{[-1]} \dots$$

While we're at it, we verify the two next-easiest Eilenberg-Steenrod axioms.

Lemma A.43. *For two pairs (X, A) and (Y, B) , the natural maps $i : (X, A) \rightarrow (X \sqcup Y, A \sqcup B)$ and $j : (Y, B) \rightarrow (X \sqcup Y, A \sqcup B)$ induce an isomorphism*

$$H_*^{gm}(X, A) \oplus H_*^{gm}(Y, B) \cong H_*^{gm}(X \sqcup Y, A \sqcup B).$$

Proof. Because probes (P, ϕ) are defined with P connected, a probe in $X \sqcup Y$ factors through either X or Y . Thus there are canonical bijections $\mathbb{Q}(X \sqcup Y) = \mathbb{Q}(X) \sqcup \mathbb{Q}(Y)$, and the boundary operator respects the obvious decomposition $C_*^{gm}(X \sqcup Y) \cong C_*^{gm}(X) \oplus C_*^{gm}(Y)$. The inverse to this decomposition is $i_* \oplus j_*$. The result follows. \square

Lemma A.44. *Let \star be the one-point manifold. There is a canonical isomorphism $C_*^{gm}(\star) \cong \mathbb{Z}[0]$, and hence a canonical isomorphism $H_*^{gm}(\star) \cong \mathbb{Z}[0]$.*

Proof. A probe (P, ϕ) to the point has ϕ constant; by definition, a probe is degenerate if $\dim P > 0$ and for all positive-dimensional strata, $d\phi$ has rank less than the dimension of the stratum. Because ϕ is constant, this is tautological, and thus all positive-dimensional probes are zero. Because there is only one zero-dimensional connected stratified-smooth space, $C_*^{gm}(\star)$ is free abelian on one degree-zero generator; further, the unique zero-dimensional stratified-smooth probe has a canonical orientation: the orientation $+1 \in \mathbb{R} = \det(\mathbb{R}^0)$. This gives a canonical generator of $C_0^{gm}(\star)$, and hence a canonical isomorphism to $\mathbb{Z}[0]$. \square

The homotopy axiom is only slightly more complicated. The usual proof follows by constructing a map to $C_*(X) \rightarrow C_*(I \times X)$ which satisfies an appropriate boundary formula, but we choose to only define $C_*^{gm}(X)$ for open manifolds. Instead, we construct an analogous map to $\mathbb{R} \times X$.

Lemma A.45. *There is a natural map*

$$\text{Cyl} : C_*^{gm}(X) \rightarrow C_{*+1}^{gm}(\mathbb{R} \times X);$$

writing $i^0 : X \rightarrow \mathbb{R} \times X$ for the inclusion of $\{0\} \times X$ and similarly for i^1 , we have the boundary formula

$$\partial \text{Cyl} + \text{Cyl} \partial = i_*^1 - i_*^0.$$

Given a smooth homotopy $H : I \times X \rightarrow Y$, one may extend this noncanonically to a smooth map $\hat{H} : \mathbb{R} \times X \rightarrow Y$. Then $\hat{H}_* \text{Cyl}$ gives a chain homotopy between H_0 and H_1 .

Proof. We set $\text{Cyl}(P, \phi) = (I \times P, \text{id} \times \phi)$, where $I = [0, 1]$. Because this operation respects orientation-reversal, collapse-equivalence, and sends degenerate chains to degenerate chains, it induces a map on the underlying chain complexes. The boundary formula here is a special case of the fiber-product boundary formula in Proposition A.29. \square

We have now verified almost all of the Eilenberg-Steenrod axioms. Unfortunately, the excision axiom is the most subtle, and an order of magnitude more difficult to prove.

Suppose (X, A) is a manifold pair, and B is a closed subset of X contained in A . We want to show that $H_*(X - B, A - B) \rightarrow H_*(X, A)$ is an isomorphism. Choose a separating codimension-1 closed submanifold C which is contained in A and so that B is contained on one side of C ; write $X = Y \cup_C Z$, with $B \subset Y$.

Given a stratified-smooth probe in X whose boundary lies in A , then ‘intersection with Z ’ should produce a stratified-smooth chain in $X - B$ whose boundary lies in $A - B$. In fact, if ϕ is the chain in X and ϕ_Y, ϕ_Z are intersections with Y and Z respectively, then ϕ should be homologous to $\phi_Y + \phi_Z$.

To see why, let us be more precise. Using the Urysohn lemma, choose a smooth function $f : X \rightarrow [-1, 1]$ with $B \subset f^{-1}(-1)$ and $A^c \subset f^{-1}(1)$; further suppose that 0 is a regular value. We write $Y = f^{-1}[-1, 0]$ and $Z = f^{-1}[0, 1]$, and $C = f^{-1}(0)$ for the codimension 1 submanifold arising as their common boundary. We write

$$W = \{(t, x) \in [-1, 1] \times X \mid t \leq 1 - |f(x)|\}$$

for the ‘creased cylinder’, or ‘silo’. It is a manifold with corners, whose boundary is given by $Y \sqcup Z \sqcup -X$, where $-X$ is X equipped with the opposite orientation.

Then intersection with W should provide the homology between ϕ and $\phi_Y + \phi_Z$. This, together with some formal manipulation, should prove the excision property. There are some additional difficulties involving transversality, which we will go through in the proof below.

Lemma A.46. *If (X, A) is a manifold pair and $B \subset X$ is a closed subset which is contained in A , then the map $i_* : H_*(X - B, A - B) \rightarrow H_*(X, A)$ induced by inclusion is an isomorphism.*

Proof. We will show that i_* is both injective and surjective.

To start, pick a homology class $[\phi] \in H_*(X, A)$; pick a representative

$$\phi = \sum [(P_i, \phi_i)] \in C_*^{gm}(X)$$

where (P_i, ϕ_i) are stratified-smooth probes, with $\partial\phi = \phi_A + \phi_{\text{cancel}}$ given by a sum of two terms: ϕ_A , a chain whose image lies in A , and ϕ_{cancel} , a chain which is zero in $C_*^{gm}(X)$. This means that ϕ_{cancel} is the disjoint union of components which are collapse-equivalent to each other with opposite orientations, and which thus cancel out after identifying collapse-equivalent chains.

Write $j_C : C \hookrightarrow X$ for the inclusion of the submanifold C constructed before this proposition. It need not be the case that the ϕ_i are transverse to j_C . However, by standard transversality arguments, one may construct a smooth homotopy $H : [0, 1] \times X \rightarrow X$ so that $H_0 = \text{Id}$ and $H_1\phi_i$ are all transverse to j_C ; furthermore, one may take H to be arbitrarily close to the constant homotopy. Using H , construct the chain $\phi_H = H(I \times \phi)$; this has boundary

$$\partial\phi_H = H_1\phi - \phi + H(I \times \partial\phi).$$

By the assumption that H is arbitrarily small and ϕ_A has image in A , we also have that $H(I \times \phi_A)$ has image in A as well as well; further, $H(I \times \phi_{\text{cancel}})$ still cancels out, so that $H(I \times \partial\phi) \in C_*^{gm}(A)$. If one thus considers ϕ_H as an element of $C_*^{gm}(X, A)$, this formula shows that $[H_1\phi] = [\phi]$. So we may assume that $[\phi]$ has a representative transverse to j_C , hence also to all of j_Y, j_Z, j_W .

Next, set $\phi_Z = Z \times_X \phi$ and similarly for ϕ_C, ϕ_Y, ϕ_W . Because

$$\partial\phi_Z = \phi_C + (\partial\phi)_Z,$$

and C is contained in A , and $\partial\phi = \phi_A + \phi_{\text{cancel}}$ with the first term contained in A , we see that $\partial\phi_Z$ lies in $C_*^{gm}(A)$; this requires the observation that $Z \times_X \phi_{\text{cancel}}$ still cancels out to give zero (which is true as j_Z is transverse to $\phi_{\text{cancel}} \subset \partial\phi$). One needs to do similar manipulations with ϕ_{cancel} a handful of other times in the course of the proof, but the discussion is nearly identical, so we suppress it from here on.

In particular, $[\phi_Z] \in H_*(X, A)$ is well-defined. Further, because $Z \subset X - B$, this in fact defines an element of $H_*(X - B, A - B)$. We wish to show that $[\phi_Z] = [\phi]$, so that the map i_* is surjective.

To see this, observe that

$$\partial(\phi_W) = \phi_Y + \phi_Z - \phi + (\partial\phi)_W.$$

Because ϕ_Y lies in $C_*^{gm}(A)$, and $(\partial\phi)$ lies in $C_*^{gm}(A)$, so too does $(\partial\phi)_W$; it follows that in $C_*^{gm}(X, A)$ we have $\partial\phi_W = \phi_Z - \phi$, and thus $[\phi] = [\phi_Z]$ in $H_*(X, A)$, as desired. So i_* is surjective.

Injectivity follows essentially the same lines, now using that we may construct H so small that chains which are already transverse to j_C remain transverse to j_C . \square

We are now ready to assemble this into the main theorem of this Appendix.

Theorem A.47. *The geometric chain complex $C_*^{gm}(X)$ is a functorial chain complex on the category of smooth manifolds and smooth maps. When X is the interior of a compact manifold, the resulting homology groups are canonically isomorphic to singular homology $H_*(X)$.*

Proof. We have now seen that $H_*^{gm}(X, A)$ defines a functor on the category MP of pairs (X, A) , where X is a smooth manifold without boundary and A is an open submanifold, and that this functor satisfies an appropriate version of the Eilenberg-Steenrod axioms for this category. Let us now explain how to modify the usual Eilenberg-Steenrod argument to this setting.

There is a closely related category, the category HP of *handlebody pairs* (Y, B) . Here, Y is a compact manifold (possibly with boundary) equipped with a handlebody decomposition, where the index of newly attached handles is monotonically increasing and B is a submanifold (possibly with boundary) obtained by only attaching some of these handles. We demand that Y is obtained by attaching handles of index no smaller than the largest index of a handle in B . Then (Y, B) carries a natural filtration $(Y^{(k)}, B^{(k)})$ of the submanifolds obtained by attaching only handles up to index k .

We say that a smooth map $(Y, B) \rightarrow (Y', B')$ is handle-preserving if it preserves this filtration. There is a functor $\text{HP} \rightarrow \text{MP}$ sending (Y, B) to (Y°, B°) .

Both MP and HP have a natural notion of homotopy; in the former category, this just means smooth homotopy; in the latter category, we say $f_0, f_1 : (Y, B) \rightarrow (Z, C)$ are homotopic if there exists a handle-preserving smooth map $(I \times Y, I \times B) \rightarrow (Z, C)$ interpolating between them. A cellular approximation-type lemma implies that the corresponding functor on homotopy categories $\text{hHP} \rightarrow \text{hMP}$ is an equivalence of categories.

Now by passing to MP, the geometric homology functor H_*^{gm} defines a functor on the category hHP. If we can show that H_*^{gm} is naturally isomorphic to singular homology on hHP, the result follows for hMP and hence MP as well. But what remains follows essentially the standard argument; one determines a canonical identification $H_*^{gm}(X^{(k)}, A^{(k)}) \cong H_*(X^{(k)}, A^{(k)})$ inductively on the index k , using the Eilenberg-Steenrod axioms. \square

Remark A.48. It seems probable that the assumption that X is the interior of a compact manifold could be dropped, but we will only actually use this theorem for the manifolds $X = SO(3), S^2, \star$, so it seems rather pointless to extend the result more generally. \diamond

A.5 Truncating the geometric chain complex

In some presentations of geometric homology theory, the chain complex $C_*^{gm}(M)$ is concentrated in degrees $[0, \dim M + 1]$. This is a desirable feature for two reasons. First, the periodic homological algebra of [Mil19, Appendix A] only works for bounded dg-algebras. Second, this phenomenon guarantees we only need the moduli spaces $\check{M}_\zeta^+(\alpha, \beta)$ up through a certain dimension range, which allows us to ignore the bubbling phenomenon.

This is not the case for our model $C_*^{gm}(M)$ whenever $\dim M > 0$. Pick any nondegenerate $\phi : P \rightarrow M$ of positive degree (for instance, take $(P, \phi) = (M, 1_M)$ to be the identity probe). Then $(P \times I^k, \phi\pi_1)$ is nondegenerate, as the positive-dimensional stratum $(P \times \{0\}, \phi)$ does not have small rank. So $C_*^{gm}(M)$ is nonzero in arbitrarily large degrees. This is because our rather strict notion of nondegeneracy includes fewer nondegenerate probes than other authors; see Remark A.34 for more details.

We briefly explain how to correct for this failure.

Let M be a smooth manifold. Set the *truncated geometric chain complex* to be

$$tC_k^{gm}(M) = \begin{cases} C_k^{gm}(M) & k < \dim M \\ C_k^{gm}(M)/\partial C_{k+1}^{gm}(M) & k = \dim M \\ 0 & k > \dim M \end{cases}.$$

That is, $tC_*^{gm}(M)$ is the complex obtained by collapsing the acyclic subcomplex

$$D_*(M) = \begin{cases} 0 & k < \dim M \\ \partial C_{k+1}^{gm}(M) & k = \dim M \\ C_k^{gm}(M) & k > \dim M. \end{cases}$$

Theorem A.49. *The complex $tC_*^{gm}(M)$ satisfies the following properties.*

- (i) *It is concentrated in degrees $[0, \dim M]$.*
- (ii) *Its homology groups are canonically isomorphic to singular homology.*
- (iii) *If $f : M \rightarrow N$ is a smooth map with $\dim M \geq \dim N$, then there is an induced chain map $f_* : tC_*^{gm}(M) \rightarrow tC_*^{gm}(N)$.*
- (iv) *Suppose $(W, f \times g) \in \mathcal{P}(M \times N)$ is a probe, and suppose $f : W \rightarrow M$ is a submersion. Then there is an induced map $\times_M W : tC_*^{gm}(M) \rightarrow tC_*^{gm}(N)$ of degree $\dim W - \dim M$, sending $(P, \phi) \mapsto (P \times_M W, \pi_2 g)$, and satisfying the boundary relation*

$$\partial(P \times_M W) = (\partial P) \times_M W + (-1)^{\dim P + \dim W} P \times_M (\partial W).$$

Proof. It is obvious that this complex is concentrated in degrees $[0, \dim M]$. Further, because $D_*(M)$ is an acyclic complex, the canonical map $C_*^{gm}(M) \rightarrow tC_*^{gm}(M)$ is a quasi-isomorphism. Because $H_*^{gm}(M)$ was already seen to be canonically isomorphic to singular homology, so too is $tH_*^{gm}(M)$.

The claimed maps are induced by maps on C_*^{gm} . We have already discussed the induced map f_* on C_*^{gm} ; because f_* is a chain map of degree zero, it obviously has $f_* D_*(M) \subset D_*(N)$, and thus induces a map on the truncated complexes.

To see that $\times_M W$ is well-defined at the level of C_*^{gm} , one needs to see that it respects all of: orientation-reversal of probes; degenerate probes; and collapse-equivalence of probes. It is easy to verify that it does. The stated boundary relation is given in Proposition A.29. We now need to verify it descends to tC .

Because $f : W \rightarrow M$ is a submersion, the map $\times_M W$ has non-negative degree. If $\deg(\times_M W) = \dim W - \dim M > 0$ it is obvious that $F(W)D_*(M) \subset D_*(N)$ for degree reasons. If $\deg F(W) = 0$, then $\dim W = \dim M$; for $f : W \rightarrow M$ is a submersion it must be on all boundary strata as well, but because boundary strata have dimension less than $\dim M$, it follows that $\partial W = \emptyset$. It follows that $F(W)$ is a chain map of degree zero and any such map descends to tC_*^{gm} as desired. \square

As a corollary to Theorem A.49(iii), if G is a Lie group, then $tC_*^{gm}(G)$ is a dg-algebra and if G acts on M then $tC_*^{gm}(M)$ is a dg-module.

References

- [AB95] D. M. Austin and P. J. Braam, *Morse-Bott theory and equivariant cohomology*, The Floer memorial volume, 1995, pp. 123–183. MR1362827 ↑8, 11, 46
- [AB96] David M. Austin and Peter J. Braam, *Equivariant Floer theory and gluing Donaldson polynomials*, *Topology* **35** (1996), no. 1, 167–200. MR1367280 ↑3, 6, 32
- [Dae20] Aliakbar Daemi, *Chern-Simons functional and the homology cobordism group*, *Duke Math. J.* **169** (2020), no. 15, 2827–2886. MR4158669 ↑7
- [DK90] Simon Donaldson and Peter Kronheimer, *The geometry of four-manifolds*, Oxford University Press, 1990. ↑6, 26, 55, 57, 104
- [Don02] S. K. Donaldson, *Floer homology groups in Yang-Mills theory*, Cambridge Tracts in Mathematics, vol. 147, Cambridge University Press, Cambridge, 2002. With the assistance of M. Furuta and D. Kotschick. MR1883043 ↑3, 5, 6, 29, 55, 56
- [Don90] ———, *Polynomial invariants for smooth four-manifolds*, *Topology* **29** (1990), no. 3, 257–315. MR1066174 ↑6
- [DS19] Aliakbar Daemi and Christopher Scaduto, *Equivariant aspects of singular instanton Floer homology*, arXiv e-prints (December 2019), arXiv:1912.08982, available at [1912.08982](https://arxiv.org/abs/1912.08982). ↑6
- [Flo88] Andreas Floer, *An instanton-invariant for 3-manifolds*, *Comm. Math. Phys.* **118** (1988), no. 2, 215–240. MR956166 ↑3, 10
- [Frø02] Kim Frøyshov, *Equivariant aspects of Yang-Mills Floer theory*, *Topology* **41** (2002), no. 3, 525–552. MR1910040 ↑5, 6
- [Frø96] ———, *The Seiberg-Witten equations and four-manifolds with boundary*, *Math. Res. Lett.* **3** (1996), no. 3, 373–390. MR1397685 ↑6
- [Fuk96] Kenji Fukaya, *Floer homology of connected sum of homology 3-spheres*, *Topology* **35** (1996), no. 1, 89–136. MR1367277 ↑40, 46
- [GF21] Dev Sinha Greg Friedman Anibal M. Medina-Mardones, *Flowing from intersection to cup product*, 2021. ↑133
- [Hel22] Gard Olav Helle, *Equivariant instanton Floer homology and calculations for the binary polyhedral spaces*, 2022. ↑4, 12, 79, 115
- [HK11] Matthew Hedden and Paul Kirk, *Chern-Simons invariants, SO(3) instantons, and $\mathbb{Z}/2$ homology cobordism*, *Chern-Simons gauge theory: 20 years after*, 2011, pp. 83–114. MR2809448 ↑114
- [Joy12] Dominic Joyce, *On manifolds with corners*, *Advances in geometric analysis*, 2012, pp. 225–258. MR3077259 ↑116, 119
- [KM07] Peter Kronheimer and Tomasz Mrowka, *Monopoles and three-manifolds*, New Mathematical Monographs, vol. 10, Cambridge University Press, Cambridge, 2007. MR2388043 ↑5, 6, 26, 28, 48
- [KM11] P. B. Kronheimer and T. S. Mrowka, *Knot homology groups from instantons*, *J. Topol.* **4** (2011), no. 4, 835–918. MR2860345 ↑3
- [KM95] ———, *Embedded surfaces and the structure of Donaldson’s polynomial invariants*, *J. Differential Geom.* **41** (1995), no. 3, 573–734. MR1338483 ↑15
- [Kot91] D. Kotschick, *SO(3)-invariants for 4-manifolds with $b_2^+ = 1$* , *Proc. London Math. Soc. (3)* **63** (1991), no. 2, 426–448. MR1114516 ↑106
- [Kro05] P. B. Kronheimer, *Four-manifold invariants from higher-rank bundles*, *J. Differential Geom.* **70** (2005), no. 1, 59–112. MR2192061 ↑15, 23, 39
- [Lin18] Francesco Lin, *A Morse-Bott approach to monopole Floer homology and the triangulation conjecture*, *Mem. Amer. Math. Soc.* **255** (2018), no. 1221, v+162. MR3827053 ↑28, 40
- [Lip14] Max Lipynskiy, *Geometric homology*, 2014. ↑40, 133
- [Mil19] S. Michael Miller, *Equivariant instanton homology*, 2019. ↑4, 5, 8, 10, 11, 12, 13, 18, 19, 20, 21, 23, 24, 25, 26, 28, 29, 30, 33, 36, 40, 47, 79, 106, 128, 143

- [MM92] John W. Morgan and Tomasz S. Mrowka, *A note on Donaldson's polynomial invariants*, Internat. Math. Res. Notices **10** (1992), 223–230. MR1191573 ↑15, 104, 106
- [MW02] Matilde Marcolli and Bai-Ling Wang, *Seiberg-Witten and Casson-Walker invariants for rational homology 3-spheres*, Proceedings of the Euroconference on Partial Differential Equations and their Applications to Geometry and Physics (Castelvecchio Pascoli, 2000), 2002, pp. 45–58. MR1919892 ↑
- [NST19] Yuta Nozaki, Kouki Sato, and Masaki Taniguchi, *Filtered instanton Floer homology and the homology cobordism group*, arXiv e-prints (May 2019), arXiv:1905.04001, available at 1905.04001. ↑7
- [OS03] Peter Ozsváth and Zoltán Szabó, *Absolutely graded Floer homologies and intersection forms for four-manifolds with boundary*, Adv. Math. **173** (2003), no. 2, 179–261. MR1957829 ↑6
- [OS04] ———, *Holomorphic disks and topological invariants for closed three-manifolds*, Ann. of Math. (2) **159** (2004), no. 3, 1027–1158. MR2113019 (2006b:57016) ↑4
- [Sca15] Christopher W. Scaduto, *Instantons and odd Khovanov homology*, J. Topol. **8** (2015), no. 3, 744–810. MR3394316 ↑3
- [SS10] Paul Seidel and Ivan Smith, *Localization for involutions in Floer cohomology*, Geom. Funct. Anal. **20** (2010), no. 6, 1464–1501. MR2739000 ↑118, 125, 128
- [Tau82] Clifford Henry Taubes, *Self-dual Yang-Mills connections on non-self-dual 4-manifolds*, J. Differential Geometry **17** (1982), no. 1, 139–170. MR658473 ↑54
- [Tau90] ———, *Casson's invariant and gauge theory*, J. Differential Geom. **31** (1990), no. 2, 547–599. MR1037415 ↑3, 4, 31, 113
- [Wal92] Kevin Walker, *An extension of Casson's invariant*, Annals of Mathematics Studies, vol. 126, Princeton University Press, Princeton, NJ, 1992. MR1154798 ↑4, 31, 113, 114
- [Zho20] Zhengyi Zhou, *Morse-bott cohomology from homological perturbation theory*, 2020. ↑44