An index relation for the quilted Atiyah-Floer conjecture

David L. Duncan

Abstract

Given a closed, connected, oriented 3-manifold with positive first Betti number, one can define an instanton Floer group as well as a quilted Lagrangian Floer group. Each of these is equipped with a relative grading. We show that the relative gradings agree.

Contents

1 Introduction **1**

2 Background **3**
2.1 Gauge theory 5
2.2 Symplectic geometry 7
2.3 Instanton Floer cohomology 16

3 The index relation **19**
3.1 Elliptic estimates on surfaces and cobordisms 20
3.2 Uniform elliptic estimates for $\mathcal{D}_{\epsilon,A}$ 25
3.3 Proof of the index relation 37

1 Introduction

Suppose Y is a closed, connected, oriented 3-manifold with positive first Betti number. The *quilted Atiyah-Floer conjecture* states that there is a natural isomorphism

$$HF_{\text{inst}}(Y) \cong HF_{\text{symp}}(Y)$$

between the instanton Floer group and the quilted Lagrangian Floer group associated to Y. These Floer groups are obtained from the homology of relatively $\mathbb{Z}_4$-graded chain complexes CF_{inst} and CF_{symp}, respectively. There is a natural *group* isomorphism $\Psi : CF_{\text{inst}} \to CF_{\text{symp}}$, and the main result of the present paper states that Ψ preserves the gradings; see Corollary 3.2. This is a part of the author's project to show that Ψ induces a graded isomorphism at the homology level.

To explain the result in more detail, we recall that the Floer chain complex CF_{inst} (resp. CF_{symp}) is defined using a certain moduli space of instantons (resp. holomorphic curves). The linearization of the equations defining the moduli space is a Fredholm operator. Then the relative grading for CF_{inst} (resp. CF_{symp}) can be defined as the index of this Fredholm operator; one reduces this index modulo 4 to get a quantity that is independent of the auxiliary choices (e.g., of the choice of instanton or holomorphic curve where the linearization takes place). We prove in Theorem 3.1 that the auxiliary choices can be made so that the two Fredholm indices agree; that Ψ preserves the gradings is then an immediate corollary.

Our approach to Theorem 3.1 is to first show the two Fredholm operators can be made surjective using the same set of auxiliary data. We then construct an isomorphism between the kernels. An upshot of this approach is that, along the way, we establish various uniform estimates on the linearized operators that are useful in their own right; see Theorems 3.11 and 3.13. Moreover, the approach we adopt generalizes to a large class of cylindrical end 4-manifolds that includes the cylinders $\mathbb{R} \times Y$ considered here; this more general set-up will appear elsewhere.

Remark 1.1. *(a) The instanton Floer group used here is a direct extension of the homology group defined by Floer in [12]; see also [3].*

(b) The idea for defining the quilted Lagrangian Floer group of Y is to appeal to Floer's Lagrangian intersection homology theory [13]. This method for obtaining 3-manifolds invariants has been in the math literature for some time, starting with Atiyah [1]. However, it was not known whether this type of procedure yields an invariant of 3-manifold with positive first Betti number until the work of K. Wehrheim and C. Woodward, who proved invariance using pseudoholomorphic quilts [26, 28].

(c) The quilted Atiyah-Floer conjecture is a cousin of the more standard Atiyah-Floer conjecture *for homology 3-spheres described by Atiyah in [1]. The main technical difference between the two conjectures is that the quilted variant uses the Betti number assumption on Y to avoid issues arising from reducible connections. See D. Salamon and K. Wehrheim [23] for an approach to the standard Atiyah-Floer conjecture, as well as for index considerations analogous to those appearing here.*

Our results extend work of S. Dostoglou and D. Salamon [5, 7] who, in the mid-90's, considered the special case where Y is a mapping torus. Since they were considering only mapping tori, Dostoglou-Salamon were able to work with holomorphic *cylinders* on the symplectic side. In order to handle 3-manifolds that are more general than mapping tori, we follow [26] and consider holomorphic *strips with Lagrangian boundary conditions*. A direct modification of Atiyah's neck-stretching ideas [1] tells us that we should expect instantons to approximate these strips. This, however, leads us to consider objects with only *approximate* Lagrangian boundary conditions. Therein lies the primary difficulty with our set-up: without Lagrangian boundary conditions on the nose, integration by parts produces boundary terms that are difficult to control analytically. In the present paper, we are only concerned with the linearized problem, but the same issue arises (see [9] for aspects of the nonlinear problem). In some sense, these boundary condition issues are addressed in Claim 2 appearing in the proof of Theorem 3.13; however, in another sense, concern for these issues permeates our approach altogether.

Acknowledgments: The author is grateful to his thesis advisor Chris Woodward for his insight and valuable suggestions. This work was partially supported by NSF grant DMS 1207194.

2 Background

Throughout this paper, Y will denote a fixed closed, connected, oriented 3-manifold equipped with a Morse function $f : Y \to S^1$ with non-empty, connected fibers. We will refer to the pair (Y, f) as a *broken circle fibration*. Given Y, such a function f exists if and only if Y has positive first Betti number; see [15]. Following [26] and [16], we will use f to decompose Y into a union of elementary cobordisms (a so-called *Cerf decomposition*) as follows. Assume f has been suitably homotoped so that its set of critical points is non-empty and in bijection with its set of critical values

$$\{c_{01}, c_{12}, \ldots, c_{(N-1)0}\} \subset S^1 = \mathbb{R}/C\mathbb{Z},$$

where $C > 0$ is some constant. It then follows that the number of critical values N must be even. Moreover, we can find some $\delta > 0$, and regular values $r_j \in S^1$ so that $c_{j(j+1)} \in [r_j + \delta, r_{j+1} - \delta]$. Here and below indices j should be taken modulo N, when appropriate. We may assume the circumference C is large enough to take $\delta = 1/2$. Define

$$\Sigma_j := f^{-1}(r_j - 1/2), \qquad Y_{j(j+1)} := f^{-1}\left([r_j + 1/2, r_{j+1} - 1/2]\right),$$

which are closed, connected, oriented surfaces and elementary cobordisms, respectively.

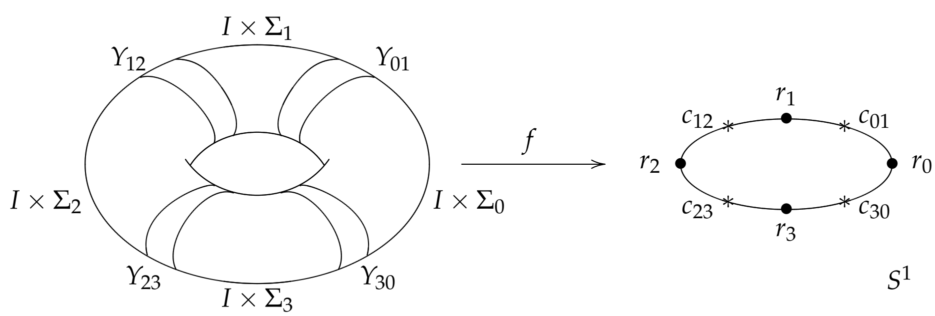

Figure 1: An illustration of a broken circle fibration.

Fix a metric g on Y. We refer to g, or its restriction to any submanifold of Y, as the *fixed metric*. Note that there are no critical values between $r_j - 1/2$ and $r_j + 1/2$, so $V := \nabla f / |\nabla f|$ is well-defined on $f^{-1}([r_j - 1/2, r_j + 1/2])$. The time-1 gradient flow of V provides an identification

$$f^{-1}([r_j - 1/2, r_j + 1/2]) \cong I \times \Sigma_j,$$

where we have set $I := [0,1]$. This also provides an identification of $f^{-1}(t)$ with Σ_j for $t \in \left[r_j - 1/2, r_j + 1/2\right]$. Then the function f, together with the metric g, allow us to view Y as the composition of cobordisms

$$Y_{01} \cup_{\Sigma_1} (I \times \Sigma_1) \cup_{\Sigma_1} Y_{12} \cup_{\Sigma_2} \ldots \cup_{\Sigma_{N-1}} Y_{(N-1)0} \cup_{\Sigma_0} (I \times \Sigma_0) \cup_{\Sigma_0} . \tag{1}$$

Note that (1) is *cyclic* in the sense that the cobordism $I \times \Sigma_0$ on the right is glued to the cobordism Y_{01} on the left. The case $N = 4$ is illustrated in Figure 1. We set

$$\Sigma_\bullet := \bigsqcup_j \Sigma_j \qquad \text{and} \qquad Y_\bullet := Y \setminus (I \times \Sigma_\bullet),$$

which allows us to write

$$Y = Y_\bullet \cup_{\partial Y_\bullet} (I \times \Sigma_\bullet). \tag{2}$$

We will refer to the connected components of the boundary $\partial Y_\bullet$ as the *seams*, and we will use the letter t to denote the coordinate variable on the interval I.

Remark 2.1. *Shortly we will consider the 4-manifold $\mathbb{R} \times Y$ and the obvious map to $\mathbb{R} \times S^1$ induced by f. The critical values of f induce parallel lines $\mathbb{R} \times c_{j(j+1)}$ in the cylinder $\mathbb{R} \times S^1$. In the language of [27], this gives $\mathbb{R} \times S^1$ the structure of a* quilted cylinder *with* seams $\mathbb{R} \times c_{j(j+1)}$. *This is effectively the source of the term 'quilted' in the present paper.*

Over $I \times \Sigma_\bullet$ the metric g has the form $dt^2 + g_\Sigma$, where g_Σ is a path of metrics on $\Sigma_\bullet$. To simplify the discussion, we assume that g has been chosen so that g_Σ is a *constant* path; this can always be achieved using the decomposition in (1) and a bump function. For $0 < \epsilon \leq 1$ define a new metric

$$g_\epsilon := \begin{cases} dt^2 + \epsilon^2 g_\Sigma & \text{on } I \times \Sigma_\bullet \\ \epsilon^2 g & \text{on } Y_\bullet. \end{cases}$$

Then $\epsilon^{-2} g_\epsilon$ is basically Atiyah's neck-stretching metric from [1].

Let $\mathcal{S}_1$ denote the smooth structure on Y (i.e., the smooth structure in which g and f are smooth). We will call this the *standard smooth structure*. It is important to note that when $\epsilon \neq 1$ the metric g_ϵ is *not* smooth in the standard smooth structure. However, there is a different smooth structure $\mathcal{S}_\epsilon = \mathcal{S}_\epsilon(Y)$ in which g_ϵ is smooth, and $Y^\epsilon := (Y, \mathcal{S}_\epsilon)$ is diffeomorphic to $(Y, \mathcal{S}_1)$; see [20, 8]. We call $\mathcal{S}_\epsilon$ the *ϵ-dependent smooth structure*, and say that a function, form, connection, etc. on Y is *ϵ-smooth* if it is smooth with respect to $\mathcal{S}_\epsilon$.

Much of our analysis will concern the product $\mathbb{R} \times Y$, and we will use s for the coordinate variable on $\mathbb{R}$. We equip $\mathbb{R} \times Y$ with the metric $ds^2 + g_\epsilon$. Given a measurable subset $S \subset \mathbb{R} \times Y$, we will write

$$\| \cdot \|_{L^p(S),\epsilon}$$

for the L^p-norm on S determined by $ds^2 + g_\epsilon$. We will not typically keep track of the underlying vector bundle (e.g., we use the same symbol $\| \cdot \|_{L^p(Y),\epsilon}$ to denote the norm on sections of TY as well as on sections of $\Lambda^2 T^*Y$). When S is clear from context, we will simply write

$$\|\cdot\|_\epsilon := \|\cdot\|_{L^2(S),\epsilon}, \qquad \text{and} \qquad (\cdot,\cdot)_\epsilon := (\cdot,\cdot)_{L^2(S),\epsilon}$$

for the L^2-norm and inner product. Denote by $*^S_\epsilon$ the Hodge star of the restriction $(ds^2 + g_\epsilon)|_S$. When $\epsilon = 1$ we will drop ϵ from the notation. For example, on $\Sigma_\bullet = \{(s,t)\} \times \Sigma_\bullet \subset \mathbb{R} \times Y$ we have

$$(\mu,\nu)_{L^2(\Sigma_\bullet),\epsilon} = \int_{\Sigma_\bullet} \langle \mu \wedge *^{\Sigma_\bullet}_\epsilon \nu\rangle = \epsilon^{2-2k} \int_{\Sigma_\bullet} \langle \mu \wedge *^{\Sigma_\bullet} \nu\rangle,$$

where μ, ν are k-forms on $\Sigma_\bullet$ with values in some vector bundle with inner product $\langle\cdot,\cdot\rangle$. We will often abuse notation and write $*^\Sigma$ for $*^{\Sigma_\bullet}$.

We will use $W^{k,p}(S,V)$ to denote the space of maps of Sobolev class $W^{k,p}$ from a manifold S to a Banach space (or bundle) V. We will specify the connection used to define these spaces only if it relevant. We also set $L^p(S,V) = W^{0,p}(S,V)$. As an example, if $f : \mathbb{R} \to W^{1,2}(Y)$ is a continuous compactly supported function, then

$$\|f\|^2_{L^2(\mathbb{R},W^{1,2}(Y))} = \int_{\mathbb{R}} \|f(s)\|^2_{W^{1,2}(Y)}\, ds.$$

We note that the usual Sobolev inequalities for $W^{k,p}(S,\mathbb{R})$ hold equally well for $W^{k,p}(S,V)$ for any Banach space V. When V has finite rank the usual compact Sobolev embedding statements hold as well.

2.1 Gauge theory

Principal $\mathrm{PU}(r)$-bundles $Q \to Y$ are classified, up to isomorphism, by a characteristic class t_2 taking values in $H^2(Y,\mathbb{Z}_r)$; see [30]. Fix a section γ of $f : Y \to S^1$, as well as a generator $d \in \mathbb{Z}_r$. Then we choose Q so that $t_2(Q) \in H^2(Y,\mathbb{Z}_r)$ is Poincaré dual to $d\,[\gamma] \in H_1(Y,\mathbb{Z}_r)$. This choice of bundle will allow us to avoid reducible connections throughout.

As with Y, for each $\epsilon > 0$, there is a canonical smooth structure on Q such that the projection $Q \to Y^\epsilon$ is smooth. The bundle Q induces bundles over the $Y_{j(j+1)}$ and Σ_j, and we set

$$Q_{j(j+1)} := Q|_{Y_{j(j+1)}}, \qquad P_j := Q|_{\Sigma_j}, \qquad Q_\bullet := \sqcup_j Q_{j(j+1)}, \qquad P_\bullet := \sqcup_j P_j.$$

As in (2), we have

$$Q = Q_\bullet \cup_{\partial Q_\bullet} (I \times P_\bullet).$$

We note that our choice of Q is such that

$$t_2(Q_{j(j+1)})[\Sigma_j] = t_2(P_j)[\Sigma_j] = d. \tag{3}$$

We equip the Lie algebra $\mathfrak{g} \cong \mathfrak{su}(r)$ of $G = \mathrm{PU}(r)$ with the Ad-invariant inner product given by

$$\langle \xi, \zeta\rangle := -\frac{1}{4\pi^2}\mathrm{Tr}(\xi\zeta).$$

This induces a fiberwise metric on the associated adjoint bundle $Q(\mathfrak{g}) \to Y$.

Let $\mathcal{G}(Q)$ denote the group of gauge transformations on Q, and $\mathcal{G}_0(Q) \subset \mathcal{G}(Q)$ the connected component of the identity. There is another preferred subgroup

$$\mathcal{G}_\Sigma \subset \mathcal{G}(Q)$$

consisting of the gauge transformations u on Q with the property that the restriction $u|_{\Sigma_0}$ can be homotoped over Σ_0 to the identity gauge transformation. Then $\mathcal{G}_\Sigma$ is a normal Lie subgroup properly containing the identity component $\mathcal{G}_0(Q)$.

In general, given a bundle $P \to X$, we will write $\mathcal{A}(P)$ for the space of connections on a bundle P, and $\mathcal{A}_{\text{flat}}(P)$ for the subspace of flat connections. Each connection $A \in \mathcal{A}(P)$ induces a covariant derivative

$$d_A : \Omega^k(X, P(\mathfrak{g})) \longrightarrow \Omega^{k+1}(X, P(\mathfrak{g}))$$

on the space of adjoint bundle-valued forms on X. We say that A is *irreducible* if d_A is injective on $\Omega^0(X, P(\mathfrak{g}))$. We will write $F_A \in \Omega^2(X, P(\mathfrak{g}))$ for the curvature.

Return now to the bundle Q. Since d is a generator of $\mathbb{Z}_r$, it can be shown that (3) implies the group identity component $\mathcal{G}_0(P_j)$ acts freely on the space $\mathcal{A}_{\text{flat}}(P_j)$ of flat connections on P_j; see [6, 26]. This implies that, at the 3-manifold level, the spaces $\mathcal{G}_0(Q_{j(j+1)})$ and $\mathcal{G}_\Sigma(Q)$ act freely on the spaces $\mathcal{A}_{\text{flat}}(Q_{j(j+1)})$ and $\mathcal{A}_{\text{flat}}(Q)$, respectfully. In particular, all flat connections that we will encounter are irreducible.

Consider the space $\mathcal{A}(\mathbb{R} \times Q)$ of connections on the pullback bundle $\mathbb{R} \times Q$ over $\mathbb{R} \times Y$. Using the product structure on $\mathbb{R} \times Y$, any $A \in \mathcal{A}(\mathbb{R} \times Q)$ can be written in components as

$$A = a + p\, ds,$$

where $a : \mathbb{R} \to \mathcal{A}(Q)$ is a path of connections on the 3-manifold Y, the object $p : \mathbb{R} \to \Omega^0(Y, Q(\mathfrak{g}))$ is a path of 0-forms, and s is the coordinate variable on $\mathbb{R}$. Similarly, the covariant derivative and curvature of A decompose as

$$d_A = d_a + ds \wedge \nabla_s, \qquad \nabla_s := \partial_s + [p, \cdot]\,,$$

$$F_A = F_a + ds \wedge b_s, \qquad b_s := \partial_s a - d_a p.$$

One can take this discussion one step further on $\mathbb{R} \times I \times \Sigma_\bullet$ by writing

$$A|_{\mathbb{R} \times I \times \Sigma_\bullet} = \alpha + \phi\, ds + \psi\, dt.$$

Here α is a map from the strip $\mathbb{R} \times I$ into the space of connections on $\Sigma_\bullet$, and ϕ, ψ are maps into the space of $P_\bullet(\mathfrak{g})$-valued 0-forms on $\Sigma_\bullet$. Then, on $\mathbb{R} \times I \times \Sigma_\bullet$, the curvature decomposes as

$$F_A = F_\alpha + ds \wedge \beta_s + dt \wedge \beta_t + \gamma\, ds \wedge dt,$$

where

$$\beta_s := \partial_s \alpha - d_\alpha \phi, \qquad \beta_t := \partial_t \alpha - d_\alpha \psi, \qquad \gamma := \partial_s \psi - \partial_t \phi - [\psi, \phi]\,.$$

Similarly, the covariant derivative satisfies

$$d_A = d_\alpha + ds \wedge \nabla_s + dt \wedge \nabla_t, \qquad \nabla_s := \partial_s + [\phi, \cdot]\,, \qquad \partial_t := \partial_t + [\psi, \cdot]\,.$$

(Note that $p|_{\mathbb{R}\times I\times\Sigma_\bullet} = \phi$ so the ∇_s here agrees with the one above.) In terms of the components of F_A, these operators satisfy the commutation relation

$$\nabla_s d_a - d_a \nabla_s = [b_s \wedge \cdot]\,.$$

Here the operators are acting on paths of forms on Y; the object $[\cdot \wedge \cdot]$ combines the wedge on forms with the Lie bracket on the adjoint bundle. Over $\mathbb{R} \times I \times \Sigma_\bullet$, these satisfy

$$\nabla_s d_\alpha - d_\alpha \nabla_s = [\beta_s \wedge \cdot]\,, \qquad \nabla_t d_\alpha - d_\alpha \nabla_t = [\beta_t \wedge \cdot]\,, \qquad \nabla_s \nabla_t - \nabla_t \nabla_s = [\gamma, \cdot]\,;$$

where the operators are acting on maps from the strip $\mathbb{R} \times I$ into $\Omega^k(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$.

Remark 2.2. *We adopt the convention that capital Latin letters are typically reserved for connections/forms on 4-manifolds, lower case Latin letters for 3-manifolds, and lower case Greek letters for 2-manifolds. For example, a 1-form V on the 4-manifolds $\mathbb{R} \times Y$ or $\mathbb{R} \times I \times \Sigma_\bullet$ may be written as*

$$V = v + r\, ds = (v, r), \qquad or \qquad V = \mu + \rho\, ds + \theta\, dt = (\mu, \rho, \theta),$$

depending on whether we wish to emphasize the 3- or 2-dimensional structure.

2.2 Symplectic geometry

Following [26], one can associate a quilted Floer cohomology group to $Q \to Y$. We briefly describe this construction here, placing emphasis on the aspects of the construction that are relevant to the proofs of our main results.

2.2.1 Symplectic data induced from a broken circle fibration

Let

$$M(P_j) := \mathcal{A}_{\text{flat}}(P_j)/\mathcal{G}_0(P_j)$$

denote the moduli space of flat connections on $P_j \to \Sigma_j$. The tangent space at $[\alpha] \in M(P_j)$ can be identified with the space

$$H^1_\alpha = H^1_\alpha(\Sigma_j) := \frac{\ker\left(d_\alpha | \Omega^1(\Sigma_j, P_j(\mathfrak{g}))\right)}{\operatorname{im}\left(d_\alpha | \Omega^0(\Sigma_j, P)j(\mathfrak{g}))\right)}$$

of α-harmonic 1-forms. By the Hodge theorem [24, Theorem 6.8], the metric on Σ_j induces a natural isomorphism

$$H^1_\alpha \cong \ker(d_\alpha \oplus d^*_\alpha) \subset \Omega^1(\Sigma_j, P_j(\mathfrak{g})).$$

We will typically treat this isomorphism as an identification. That being said, at times we will wish to emphasize that we are considering an α-harmonic 1-form η as an element of $\ker(d_\alpha \oplus d_\alpha^*)$, in which case we will refer to η as an α-harmonic 1-form *representative*.

The Hodge theorem also gives an L^2-orthogonal splitting

$$\Omega^1(\Sigma_j, P_j(\mathfrak{g})) = \ker(d_\alpha \oplus d_\alpha^*) \oplus \mathrm{im}\,(d_\alpha) \oplus \mathrm{im}\,(d_\alpha^*)\,.$$

It follows from these identifications that the space $M(P_j)$ inherits a natural symplectic structure from the symplectic form on Ω^1 given by integration. Similarly, the fixed metric on Σ_j induces a compatible complex structure J_j on $M(P_j)$ coming from the Hodge star on 1-forms.

Consider the space

$$M := M(P_0)^- \times M(P_1) \times M(P_2)^- \times \ldots \times M(P_{N-1}), \tag{4}$$

where the superscript indicates that we have replaced the symplectic form on $M(P_{2k})$ with its negative. Then M is compact, simply-connected, and monotone with monotonicity constant $1/2r$; see [26, 19]. The formula

$$J := \sum_{j=0}^{N-1} (-1)^{j+1} \mathrm{proj}_j^* J_j$$

defines a complex structure on M that is compatible with symplectic form, where proj_j is the obvious projection from M to $M(P_j)$.

Next set

$$L(Q_{j(j+1)}) := \mathcal{A}_{\mathrm{flat}}(Q_{j(j+1)})/\mathcal{G}_0(Q_{j(j+1)}).$$

As in the case with surfaces, the tangent space at $[a] \in L(Q_{j(j+1)})$ can be identified with the harmonic space

$$H_a^1 = H_a^1(Y_{j(j+1)}) := \frac{\ker\left(d_a|\Omega^1(Y_{j(j+1)}, Q_{j(j+1)}(\mathfrak{g}))\right)}{\mathrm{im}\left(d_a|\Omega^0(Y_{j(j+1)}, Q_{j(j+1)}(\mathfrak{g}))\right)}.$$

However, due to the boundary of $Y_{j(j+1)}$, the Hodge isomorphism takes the form

$$H_a^1 \cong \ker(d_a \oplus d_a^* \oplus \partial *) \subset \Omega^1(Y_{j(j+1)}, Q_{j(j+1)}(\mathfrak{g})),$$

where $\partial *$ is the map sending a 1-form v to $(*v)|_{\partial Y_{j(j+1)}}$. The Hodge decomposition is

$$\Omega^1(Y_{j(j+1)}, Q_{j(j+1)}(\mathfrak{g})) = \ker(d_a \oplus d_a^* \oplus \partial *) \oplus \mathrm{im}\,(d_a) \oplus \mathrm{im}\,(d_a^*|_{\ker \partial *})\,. \tag{5}$$

The proofs of the Hodge isomorphism and decomposition follow as in the closed case, with the operator $\partial *$ being used to kill off the boundary term that arises during integration by parts. As in the case of surfaces, the elements of the kernel $\ker(d_a \oplus d_a^* \oplus \partial *)$ are called *a-harmonic 1-form representatives*.

The space $L(Q_{j(j+1)})$ is compact and simply-connected. Moreover, restriction to each boundary component induces a Lagrangian embedding

$$L(Q_{j(j+1)}) \hookrightarrow M(P_j)^- \times M(P_{j+1}),$$

and we identify $L(Q_{j(j+1)})$ with its image under this map. The negation of the symplectic structure here ensures that $L(Q_{j(j+1)})$ is actually Lagrangian; this is ultimately the source of the negative exponents in (4).

The $L(Q_{j(j+1)})$ with j even determine a Lagrangian submanifold

$$L_{(0)} := L(Q_{01}) \times L(Q_{23}) \times \ldots \times L(Q_{(N-2)(N-1)}) \subset M.$$

in M. There is a second Lagrangian $L_{(1)} \subset M$ obtained from $Q_{j(j+1)}$ for odd j.

2.2.2 The perturbed symplectic action functional

Fix a function

$$H \in C^{\infty}(I \times M, \mathbb{R})$$

that vanishes to all orders at $\partial(I \times M)$. We will assume in addition that H is of *split-type*, meaning that

$$H(t; p_0, \ldots, p_{N-1}) = \sum_{j=0}^{N-1} (-1)^{j+1} H_j(t; p_j)$$

for some functions $H_j \in C^{\infty}(I \times M(P_j), \mathbb{R})$. The function H on M lifts to a $\mathcal{G}_0(P_\bullet)$-invariant function on the space of flat connections on $P_\bullet$, and we denote this lift by the same symbol.

Let $X^H : I \to \Gamma(TM)$ be the corresponding Hamiltonian vector field of H. We will often not explicitly write the dependence of X^H on the I-variable, and we will sometimes drop the superscript H as well. As with H, the vector field X^H lifts to a gauge equivariant vector field on the space of flat connections on $P_\bullet$.

Consider the space

$$\mathcal{P}(L_{(0)}, L_{(1)}) := \left\{ x : (I, 0, 1) \to (M, L_{(0)}, L_{(1)}) \right\}.$$

of paths with Lagrangian boundary condition. The tangent space

$$T_x \mathcal{P}(L_{(0)}, L_{(1)})$$

at a path x can be identified with the space of vector fields ξ along x with Lagrangian boundary conditions $\xi(j) \in T_{x(j)} L_{(j)}$, for $j = 0, 1$. Then $\mathcal{P}(L_{(0)}, L_{(1)})$ admits a natural 1-form λ_H defined by

$$\lambda_H : T_x \mathcal{P}(L_{(0)}, L_{(1)}) \longrightarrow \mathbb{R}, \qquad \xi \longmapsto \int_0^1 \omega_{x(t)} \left(\partial_t x - X^H(t; x(t)),\, \xi(t) \right) dt,$$

where ω is the symplectic form on M. This 1-form is closed, with cohomology class an integral multiple of the monotonicity constant $1/2r$ of M. Furthermore, since M is simply-connected and the $L_{(j)}$ are connected, the space $\mathcal{P}(L_{(0)}, L_{(1)})$ is path-connected. In particular, by fixing a base-point $x_0 \in \mathcal{P}(L_{(0)}, L_{(1)})$, we can integrate λ_H to a circle-valued function

$$\mathcal{SA}_H : \mathcal{P}(L_{(0)}, L_{(1)}) \longrightarrow \frac{\mathbb{R}}{(1/2r)\mathbb{Z}},$$

called the *perturbed symplectic action*. This sends a path $x \in \mathcal{P}(L_{(0)}, L_{(1)})$ to

$$\mathcal{SA}_H(x) := -\int_{I\times I} v^*\omega - \int_0^1 H(t;x(t))\,dt \quad \bmod (1/2r)\mathbb{Z}.$$

Here $v : I \times I \to M$ is any smooth map with $v(0,t) = x_0(t), v(1,t) = x(t)$ and $v(s,j) \in L_{(j)}$ for $j = 0,1$. Monotonicity implies that $\mathcal{SA}_H$ is independent of the choice of v, and $\mathcal{SA}_H$ depends on the choice of basepoint x_0 only up to an overall constant.

By definition, we have that $\lambda_H = d\mathcal{SA}_H$ is the differential of $\mathcal{SA}_H$, and so the formula for λ_H shows that the critical points of $\mathcal{SA}_H$ are the paths $x : (I,0,1) \to (M, L_{(0)}, L_{(1)})$ satisfying

$$\partial_t x = X^H(t;x(t));$$

these are just the Hamiltonian trajectories of H that have Lagrangian boundary conditions. We denote by

$$\mathcal{I}_H(L_{(0)}, L_{(1)})$$

the set of all of these critical points. There is a canonical identification

$$\mathcal{I}_H(L_{(0)}, L_{(1)}) \cong \Phi_1^H(L_{(0)}) \cap L_{(1)},$$

where Φ_1^H is the time-1 flow of X^H. Consequently, we refer to the elements of $\mathcal{I}_H(L_{(0)}, L_{(1)})$ as *H-Lagrangian intersection points*.

We say that an H-Lagrangian intersection point x is *non-degenerate* if the Hessian

$$\mathrm{Hess}_x \mathcal{SA}_H$$

of $\mathcal{SA}_H$ at x is non-degenerate as a bilinear form on the tangent space $T_x\mathcal{P}(L_{(0)}, L_{(1)})$. Then x is non-degenerate if and only if its associated element of $\Phi_1^H(L_{(0)}) \cap L_{(1)}$ is a transverse intersection point. We will always define the Hessian relative to the connection on the path space coming from the L^2-inner product $(\cdot,\cdot)_{L^2(M)}$. In particular, we can identify the Hessian with an L^2-self-adjoint linear operator $\mathcal{D}_x$ on $\Gamma(x^*TM)$ in the sense that

$$\mathrm{Hess}_x \mathcal{SA}_H(\eta_1, \eta_2) = (\eta_1, \mathcal{D}_x \eta_2)_{L^2(M)}$$

for all $\eta_1, \eta_2 \in T_x\mathcal{P}(L_{(0)}, L_{(1)})$. It follows from the definitions that we can write this operator as

$$\mathcal{D}_x \eta = J(\nabla_t^{LC} \eta - dX_x^H(\eta)),$$

where J is the almost complex structure on M, ∇_t^{LC} is the t-derivative on sections of $x^*TM \to I$ induced from the Levi-Civita connection on M, and dX_x^H is the linearization at x of the map $x \mapsto X^H(x)$. It follows that an H-Lagrangian intersection point x is non-degenerate if and only if $\mathcal{D}_x$ is injective when restricted to the space of sections η with Lagrangian boundary conditions.

2.2.3 Representatives of paths

In this section, we will represent path in M as connections on Y. To begin, fix $x \in \mathcal{P}(L_{(0)}, L_{(1)})$. Using the product structure (4), we can write x as a tuple

$$x(t) = (x_0(1-t), x_1(t), x_2(1-t), \ldots, x_{N-2}(1-t), x_{N-1}(t))$$

for paths $x_j : I \to M(P_j)$; the sign convention is to account for the change in symplectic structure in (4). Each x_j lifts to a map

$$\alpha_j : I \longrightarrow \mathcal{A}_{\text{flat}}(P_j)$$

into the space of flat connections, and α_j is unique up to the action of $\text{Maps}(I, \mathcal{G}_0(P_j))$. Let α denote the connection on $P_\bullet$ that restricts to α_j on P_j. Since α is irreducible, there is a unique path

$$\psi : I \longrightarrow \Omega^0(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$$

such that $\partial_t \alpha - d_\alpha \psi$ is an α-harmonic 1-form representative. Then define a connection on $I \times P_\bullet$ by

$$\alpha + \psi\, dt.$$

The assumption that x has Lagrangian boundary conditions means that this has a continuous extension a_x over $Y_\bullet$ to all of Y, with

$$F_{a_x}|_{Y_\bullet} = 0.$$

Moreover, the connection a_x is uniquely determined by x up to the action of $\mathcal{G}_\Sigma$. We call a_x a *representative* of the path x. Using the gauge freedom, it is not hard to see that for each $\epsilon > 0$, any x admits a representative that is ϵ-smooth. Moreover, if a is ϵ-smooth then for any other $\epsilon' > 0$ this connection is of Sobolev class $W^{1,\infty}$ with respect to the ϵ'-smooth structure.

Conversely, suppose a is a ϵ-smooth connection on Q satisfying

$$F_a|_{Y_\bullet} = 0, \qquad F_\alpha = 0, \qquad d_\alpha^*(\partial_t \alpha - d_\alpha \psi) = 0 \tag{6}$$

where $a|_{I \times \Sigma_\bullet} = \alpha + \psi\, dt$. Then a descends to a unique $x \in \mathcal{P}(L_{(0)}, L_{(1)})$, and a is a representative x.

One can check that x is an H-Lagrangian intersection point if and only it has a representative that satisfies

$$\partial_t \alpha(t) - d_{\alpha(t)} \psi(t) - X^H(t; \alpha(t)) = 0. \tag{7}$$

Any connection on Q satisfying (6) and (7) will be called *H-flat*, and we will denote the space of H-flat connections by

$$\mathcal{A}_{\text{flat}}(Q, H).$$

In Section 2.3.1, we will give a 3-manifold interpretation of the H-flat connections. For now, we emphasize that the above observations imply there is a natural bijection

$$\Psi : \mathcal{A}_{\text{flat}}(Q, H)/\mathcal{G}_\Sigma \xrightarrow{\cong} \mathcal{I}_H(L_{(0)}, L_{(1)}). \tag{8}$$

Fix a representative a of x, and write $a = \alpha + \psi\, dt$ on $I \times \Sigma_\bullet$. Having chosen a representative, there is a canonical way to realize elements of the tangent space $T_x\mathcal{P}(L_{(0)}, L_{(1)})$ as 1-forms. To describe this, first note that there is a natural subbundle of $T\mathcal{A}_{\text{flat}}(P_\bullet)$ whose fiber at a connection α' is the α'-harmonic space $H^1_{\alpha'}$. Let

$$H_\alpha \longrightarrow I$$

denote the pullback of this bundle under the path $\alpha : I \to \mathcal{A}_{\text{flat}}(P_\bullet)$ covering x. The fiber of H_α over $t \in I$ can be expressed in terms of the harmonic spaces as

$$(H_\alpha)_t := H^1_{\alpha_0(1-t)} \oplus H^1_{\alpha_1(t)} \oplus \ldots \oplus H^1_{\alpha_{N-2}(1-t)} \oplus H^1_{\alpha_{N-1}(t)},$$

where α_j is the restriction of α to Σ_j. Moreover, we can identify this with the pullback bundle

$$x^*TM \cong H_\alpha. \tag{9}$$

In summary, we may view any tangent vector $\xi \in T_x\mathcal{P}(L_{(0)}, L_{(1)})$ as a path from I with values in the space of harmonic 1-form representatives; in particular, it is a path of 1-forms on $\Sigma_\bullet$. (Using the Lagrangian boundary conditions, this path can be upgraded to a 1-form defined on all of Y; we carry this out in the next section.)

Fix a representative a of x. Consider the self-adjoint Hessian operator

$$\mathcal{D}_x : T_x\mathcal{P}(L_{(0)}, L_{(1)}) \longrightarrow T_x\mathcal{P}(L_{(0)}, L_{(1)}).$$

We claim that, relative to the identification (9), this can be written as

$$\mathcal{D}_x\xi = \text{proj}_\alpha \left[*^\Sigma \left(\nabla_t \xi - dX_\alpha \xi \right) \right] = J\text{proj}_\alpha \left(\nabla_t \xi - dX_\alpha \xi \right),$$

where $\nabla_t = \partial_t + [\psi, \cdot]$, and proj_α is the L^2-orthogonal projection to the harmonic bundle $H_\alpha \subset \Omega^1(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$. The only tricky thing to check is that $\text{proj}_\alpha \nabla_t \xi$ recovers the operator ∇_t^{LC} coming from the Levi-Civita connection. To see that this is the case, note that the image of $\text{proj}_\alpha \nabla_t \xi$ in x^*TM under (9) is independent of the choice of representative a. In particular, we may assume a is in a gauge in which $\psi = 0$ (this is, put a into temporal gauge on $I \times \Sigma_\bullet$). Then

$$\text{proj}_\alpha \nabla_t = \text{proj}_\alpha \partial_t.$$

The operator ∂_t can be viewed as the t-derivative along α induced from the trivial connection on the affine space of connections. By definition, the projection of this to the harmonic bundle is exactly the t-derivative along x on M coming from the Levi-Civita connection.

If we want to emphasize that we have chosen a representative a of x we will write

$$\mathcal{D}_{0,a} := \mathcal{D}_x.$$

The subscript of 0 is to stress an analogy with the instanton case considered in Section 2.3.

2.2.4 Representatives of strips

We will be interested in maps

$$(\mathbb{R} \times I, \mathbb{R} \times \{0\}, \mathbb{R} \times \{1\}) \longrightarrow (M, L_{(0)}, L_{(1)}). \tag{10}$$

We want to describe how to represent any such map v as a connection A on the 4-manifold $\mathbb{R} \times Y$. The process is entirely analogous to that of the previous section. Indeed, we set

$$A|_{\mathbb{R}\times I\times\Sigma_\bullet} := \alpha + \phi\, ds + \psi\, dt,$$

by first declaring $\alpha : \mathbb{R} \times I \to \mathcal{A}_{\text{flat}}(P_\bullet)$ to be any lift of v, and then finding ϕ, ψ such that

$$\partial_s \alpha - d_\alpha \phi, \qquad \partial_t \alpha - d_\alpha \psi$$

are α-harmonic 1-form representatives. As above, the strip α induces a vector bundle

$$H_\alpha \to \mathbb{R} \times I$$

with fibers given by the space of α-harmonic 1-form representatives. Turning now to $\mathbb{R} \times Y_\bullet$, we note that the existence of Lagrangian boundary conditions for v precisely means that there is some path $a : \mathbb{R} \to \mathcal{A}_{\text{flat}}(Q_\bullet)$ with

$$a(s)|_{\partial Y_\bullet} = \alpha(s, \cdot)|_{\partial I\times\Sigma_\bullet}.$$

Then a is unique up to gauge transformations that are the identity on the boundary. We can use this gauge freedom to ensure that the component of $a(s)$ normal to the boundary equals $\psi(s, \cdot)|_{\partial I\times\Sigma_\bullet}$. There is a unique path $p : \mathbb{R} \to \Omega^0(Y_\bullet, Q_\bullet(\mathfrak{g}))$ such that

$$\partial_s a - d_a p$$

is an a-harmonic 1-form representative. We then set

$$A|_{\mathbb{R}\times Y_\bullet} = a + p\, ds.$$

It follows that $A|_{\mathbb{R}\times Y_\bullet}$ and $A|_{\mathbb{R}\times I\times\Sigma_\bullet}$ patch together to define a continuous connection A on $\mathbb{R} \times Y$. We can use the gauge freedom to ensure this is ϵ-smooth as well. We call A a *representative* of the strip v, and we note that any two representatives of v are related by an element of $\text{Maps}(\mathbb{R}, \mathcal{G}_\Sigma)$.

Fix a strip representative A, and let $H_\alpha \to \mathbb{R} \times I$ be as above. Define a subbundle

$$H_{a,(0)} \longrightarrow \mathbb{R} \times \{0\}$$

of $H_\alpha|_{\mathbb{R}\times\{0\}}$ by declaring the fiber of $H_{a,(0)}$ over $(s, 0)$ to be the product of the harmonic spaces

$$(H_{a,(0)})_{(s,0)} = H^1_{a(s)}(Y_{01}) \oplus H^1_{a(s)}(Y_{23}) \oplus \ldots \oplus H^1_{a(s)}(Y_{(2N-2)(2N-1)}) \subset (H_\alpha)_{(s,0)}.$$

There is a second bundle

$$H_{a,(1)} \longrightarrow \mathbb{R} \times \{1\}$$

that is defined similarly, but with $Y_{j(j+1)}$ and j odd. For $j = 0, 1$, the subbundle $H_{a,(j)}$ is Lagrangian and can be identified with the pullback $v(\cdot, j)^* TL_{(j)}$. The space of strips of the form (10) can be viewed as the path space

$$\mathrm{Maps}(\mathbb{R}, \mathcal{P}(L_{(0)}, L_{(1)}))$$

for $\mathcal{P}(L_{(0)}, L_{(1)})$. Fixing such a strip v and a representative $A = a + p\,ds$, we can realize the tangent bundle $T_v\mathrm{Maps}(\mathbb{R}, \mathcal{P}(L_{(0)}, L_{(1)}))$ as the space of sections η of H_α with Lagrangian boundary conditions in the $H_{a,(j)}$.

We will now describe how to associate to such any such η a 1-form V_η on $\mathbb{R} \times Y$. This is the linearized version of the passage from v to A, except here the 1-form V_η is uniquely determined by η and the lift A (there is no more gauge freedom once we have fixed A). The construction will show that, when η is smooth, the 1-form V_η is continuous on $\mathbb{R} \times Y$ and smooth away from the seams $\mathbb{R} \times \partial Y_\bullet$; in particular, V_η is always locally of Sobolev class $W^{1,\infty}$. Moreover, it will be clear that the assignment

$$\eta \longmapsto V_\eta \tag{11}$$

is an injective linear map.

To define V_η, view η as a map

$$\eta : \mathbb{R} \times I \longrightarrow H_\alpha \subset \Omega^1(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$$

into the space of 1-forms on $\Sigma_\bullet$. Now define V_η on $\mathbb{R} \times I \times \Sigma_\bullet$ by

$$V_\eta|_{\mathbb{R}\times I\times\Sigma_\bullet} := \eta + \rho\,ds + \theta\,dt,$$

where ρ and θ are determined (uniquely) by the conditions

$$d_\alpha^*(\nabla_s\eta - d_\alpha\rho) = 0, \qquad d_\alpha^*(\nabla_t\eta - dX_\alpha\eta - d_\alpha\theta) = 0.$$

To define V_η on $\mathbb{R} \times Y_\bullet$ we use the Lagrangian boundary conditions. Working pointwise in $s \in \mathbb{R}$, it follows from the Hodge isomorphism for $Y_\bullet$ that there are unique paths v_1 and w_2 of 1-forms and 2-forms, respectively, on $Y_\bullet$ such that

$$\begin{array}{rclrclrclrcl} d_a v_1 &=& 0, & d_a^* v_1 &=& 0, & v_1|_{\partial Y_\bullet} &=& \eta, & (*^Y v_1)|_{\partial Y_\bullet} &=& 0 \\ d_a w_2 &=& 0, & d_a^* w_2 &=& 0, & w_2|_{\partial Y_\bullet} &=& *^\Sigma\theta, & (*^Y w_2)|_{\partial Y_\bullet} &=& 0. \end{array}$$

Then set

$$v := v_1 + *^Y w_2.$$

Note that this satisfies $d_a v = 0$, $d_a^* v = 0$. Next, define a 0-form r by the condition

$$d_a^*(\nabla_s v - d_a r)|_{Y_\bullet} = 0, \qquad r|_{\partial Y_\bullet} = \rho|_{\partial(I\times\Sigma_\bullet)};$$

a unique solution always exists by classical elliptic theory. Finally set

$$V_\eta|_{\mathbb{R}\times Y_\bullet} = v + r\,ds.$$

This patches together with $V_\eta|_{\mathbb{R}\times I\times\Sigma_\bullet}$ to define a continuous 1-form V_η defined on all of $\mathbb{R} \times Y$.

2.2.5 Quilted Floer cohomology

The quilted Floer cohomology of Q is the homology of the chain complex

$$(CF^\bullet_{\mathrm{symp}}, \partial_{\mathrm{symp}})$$

that arises when one applies the general framework of Morse theory to the perturbed symplectic action $\mathcal{SA}_H$; see Floer [13] and Oh [21, 22]. By definition then, the group $CF^\bullet_{\mathrm{symp}}$ is freely generated (over $\mathbb{Z}_2$, say) by the elements of $\mathcal{I}_H(L_{(0)}, L_{(1)})$. The boundary operator ∂_{symp} is defined by counting solutions $v : (\mathbb{R} \times I, \mathbb{R} \times \{j\})_{j=0,1} \to (M, L_{(j)})_{j=0,1}$ to the (J, H)-holomorphic curve equation

$$\partial_s v + J(\partial_t v - X^H(v)) = 0. \tag{12}$$

The relevant (J, H)-holomorphic curves v used to define ∂_{symp} are those that converge to H-Lagrangian intersection points $x^\pm$ at $\pm\infty$. In terms of representatives, this means that there are H-flat connections $a^\pm$ for which $a^\pm = \lim_{s \to \pm\infty} A|_{\{s\}\times Y}$, where A represents v.

Let $\mathcal{D}_v$ denote the linearization at v of the left-hand side of the (J, H)-holomorphic curve equation (12), with the linearization defined using the Levi-Civita connection. Then $\mathcal{D}_v$ is an operator with domain the space of $W^{1,2}$-sections of $v^*TM \to \mathbb{R} \times I$ with Lagrangian boundary conditions, and with codomain the space of L^2-sections of v^*TM. (Note, however, that $\mathcal{D}_v\eta$ makes sense even if η does not have Lagrangian boundary conditions.) In terms of a representative A for v, this operator can be written as

$$\mathcal{D}_{0,A}\eta := \mathcal{D}_v\eta = \mathrm{proj}_\alpha \left[\nabla_s \eta + *^\Sigma (\nabla_t \eta - dX_\alpha \eta) \right].$$

When all H-Lagrangian intersection points $x^\pm$ are non-degenerate the operator $\mathcal{D}_{0,A}$ is Fredholm, so the index

$$\mathrm{Ind}(\mathcal{D}_{0,A})$$

is well-defined. We note that $\mathcal{D}_v = \mathcal{D}_{0,A}$ is onto (when restricted to the space with Lagrangian boundary conditions) exactly when the formal adjoint

$$\mathcal{D}^*_{0,A} := \mathrm{proj}_\alpha \left[-\nabla_s + *^\Sigma (\nabla_t - dX_\alpha) \right]$$

is injective (when restricted to the space with Lagrangian boundary conditions).

Floer showed [13] that the chain complex $(CF^\bullet_{\mathrm{symp}}, \partial_{\mathrm{symp}})$ is well-defined (i.e., $\partial^2_{\mathrm{symp}} = 0$), provided H has been chosen so that (1) all H-flat connection are non-degenerate, and (2) $\mathcal{D}_v$ is surjective for all finite energy (J, H)-holomorphic curves v. It was shown in [14] that such Hamiltonians always exist, and the proof can be modified to show that H can be chosen to be of split-type. Assuming this is the case, we define the *quilted Floer cohomology* of Q by setting

$$HF^\bullet_{\mathrm{symp}}(Q) := \ker \partial_{\mathrm{symp}} / \mathrm{im}\, \partial_{\mathrm{symp}}.$$

The symplectic manifold M has minimal Chern number 2, see [2, 6]. It follows that, when taken mod 4, the index of $\mathcal{D}_v$ depends only on the H-Lagrangian intersection points $x^\pm$. We can therefore define a relative grading

$$\mu_{\mathrm{symp}} : CF^\bullet_{\mathrm{symp}} \times CF^\bullet_{\mathrm{symp}} \longrightarrow \mathbb{Z}_4$$

by declaring

$$\mu_{\text{symp}}(x^-, x^+) := \text{Ind}(\mathcal{D}_v) \mod 4,$$

for $x^\pm \in \mathcal{I}_H(L_{(0)}, L_{(1)})$; here $v : \mathbb{R} \times I \to M$ is any map with Lagrangian boundary conditions that limits to $x^\pm$ at $\pm\infty$. Strictly speaking, the Floer boundary operator only counts (J, H)-holomorphic curves v with $\text{Ind}(\mathcal{D}_v) = 1$. It follows that this relative grading descends to a well-defined relative grading on $HF^\bullet_{\text{symp}}(Q)$.

Remark 2.3. *(a) Floer showed in [10] that this definition of the index agrees with the geometrically defined relative grading coming from the Maslov-Viterbo index.*

(b) It follows from Floer's work [13] that $HF^\bullet_{\text{symp}}(Q)$ depends only on the symplectic manifold M and the Lagrangians $L_{(0)}, L_{(1)}$, up to Hamiltonian isotopy. This was refined by Wehrheim-Woodward [26, 28, 29] to show that this Floer group depends only on Y, the topological type of the bundle Q, and the homotopy class of the broken circle fibration f.

2.3 Instanton Floer cohomology

In this section, we will define the instanton Floer chain complex associated to Q. We do this in a way that is compatible with the Lagrangian Floer chain complex defined in the previous section. Most of the analytic details are due to Floer [11, 12]; Donaldson's book [4] is also a good reference.

2.3.1 The perturbed Chern-Simons functional

For each $0 \leq j \leq N-1$, let $H_j : I \times \mathcal{A}(P_j) \to \mathbb{R}$ be a $\mathcal{G}_0(P_j)$-invariant map that vanishes to infinite order on $\partial(I \times \mathcal{A}(P_j))$. Then H_j induces a Hamiltonian on $M(P_j)$, as we considered in the previous section. Define a map $H : \mathcal{A}(Q) \to \mathbb{R}$ by

$$H(a) := \sum_j \int_I H_j \left(t; a|_{\{t\} \times \Sigma_j} \right) \, dt. \tag{13}$$

The differential of H at $a \in \mathcal{A}(Q)$ is represented by a map $K : \mathcal{A}(Q) \to \Omega^2(Y, Q(\mathfrak{g}))$ in the sense that

$$(dH)_a v = \int_Y \langle K_a \wedge v \rangle$$

for all $v \in T_a\mathcal{A}(Q)$. The object $\langle \cdot \wedge \cdot \rangle$ combines the wedge on forms with the inner product on the adjoint bundle. It follows that

$$K_a|_{Y_\bullet} = 0 \qquad \text{and} \qquad K_a|_{I \times \Sigma_j} = dt \wedge X(\alpha),$$

where $X = X^{H_j}$ is the Hamiltonian vector field on $\mathcal{A}(P_j)$ associated to H_j. With this notation, the H-flat conditions (6) and (7) can be more concisely written as

$$F_a - K_a = 0.$$

We will denote the linearization of K at a (resp. of X at α) as dK_a (resp. dX_α). Hence $dK_a = dt \wedge dX_\alpha$. One can check that the gauge invariance of H implies

$$dK_a(d_a r) = [K_a, r], \qquad d_a (dK_a v) = [K_a \wedge v] \tag{14}$$

for all 0-forms r and 1-forms v on Y; similar statements hold for X_α on $\Sigma_\bullet$.

Consider the *perturbed Chern-Simons functional*

$$\mathcal{CS}_H := \mathcal{CS} - H : \mathcal{A}(Q) \longrightarrow \frac{\mathbb{R}}{(1/2r)\mathbb{Z}},$$

where $\mathcal{CS}$ is the Chern-Simons functional for Q. The constant $1/2r$ ensures that $\mathcal{CS}$ is invariant under the preferred gauge group $\mathcal{G}_\Sigma$; see [8].

Remark 2.4. *Recall in Section 2.2.3 that we represented elements of $\mathcal{P}(L_{(0)}, L_{(1)})$ as connections on Q. Let $\mathcal{A}_{\mathrm{rep}}(Q) \subset \mathcal{A}(Q)$ denote the space of all such representatives. Then the observations of Section 2.2.3 show*

$$\mathcal{A}_{\mathrm{rep}}(Q)/\mathcal{G}_\Sigma = \mathcal{P}(L_{(0)}, L_{(1)}).$$

Let $\iota : \mathcal{A}_{\mathrm{rep}}(Q) \hookrightarrow \mathcal{A}(Q)$ denote the inclusion, and $\pi : \mathcal{A}_{\mathrm{rep}}(Q) \to \mathcal{P}(L_{(0)}, L_{(1)})$ the projection. It follows from the definitions that $\mathcal{CS}_H$ is related to the perturbed symplectic action functional by $\iota^*\mathcal{CS}_H = \pi^*\mathcal{SA}_H$.

With respect to the metric g_ϵ on Y, the L^2-gradient of $\mathcal{CS}_H$ is the map

$$a \longmapsto *^Y_\epsilon (F_a - K_a).$$

It follows that the set of critical points of $\mathcal{CS}_H$ is precisely the space $\mathcal{A}_{\mathrm{flat}}(Q, H)$ of H-flat connections. From this perspective, we say an H-flat connection is *non-degenerate* if the Hessian of $\mathcal{CS}_H$ is injective, when taken modulo the gauge action. This is equivalent to saying that the self-adjoint operator

$$\mathcal{D}_{\epsilon,a} := \begin{pmatrix} *^Y_\epsilon (d_a - dK_a) & -d_a \\ -d_a^{*_\epsilon} & 0 \end{pmatrix} : \quad \Omega^1(Y, Q(\mathfrak{g})) \oplus \Omega^0(Y, Q(\mathfrak{g})) \longrightarrow \Omega^1(Y, Q(\mathfrak{g})) \oplus \Omega^0(Y, Q(\mathfrak{g}))$$

is injective, where $d_a^{*_\epsilon}$ is the adjoint taken with respect to the ϵ-dependent Hodge star. Note that $\mathcal{D}_{\epsilon,a}$ makes sense even if a is not H-flat.

At this point we have two notions of non-degeneracy for H-flat connections: the one just described, and the one on $\mathcal{I}_H(L_{(0)}, L_{(1)})$ described in the previous section. The next proposition says that these notions are equivalent.

Proposition 2.5. *Let Ψ be the map (8). An H-flat connection $a \in \mathcal{A}_{\mathrm{flat}}(Q, H)$ is non-degenerate as a critical point of $\mathcal{CS}_H$ if and only if the associated H-Lagrangian intersection point $\Psi([a]) \in \mathcal{I}_H(L_{(0)}, L_{(1)})$ is non-degenerate as a critical point of the symplectic action $\mathcal{SA}_H$.*

Proof. Suppose $\mathcal{D}_{0,a}$ is injective and $\mathcal{D}_{\epsilon,a}(v, r) = 0$ for some $(v, r) \in \Omega^1 \oplus \Omega^0$. Apply $*^Y_\epsilon (d_a - dK_a)$ to the top equation in $\mathcal{D}_{\epsilon,a}(v, r) = 0$. This gives

$$0 = (d_a - dK_a)^*(d_a - dK_a)v - *^Y_\epsilon ([F_a, r] - dK_a d_a r) = (d_a - dK_a)^*(d_a - dK_a)v,$$

where $(d_a - dK_a)^* := *^Y_\epsilon (d_a - dK_a) *^Y_\epsilon$ is the L^2-adjoint of $d_a - dK_a$, and in the second equality we used the H-flat condition together with (14). Taking the L^2-inner product with v then shows $(d_a - dK_a)v = 0$. Hence

$$d_a v|_{Y_\bullet} = 0, \qquad d_\alpha \mu = 0, \qquad \nabla_t \mu - d_\alpha \theta - dX_\alpha(\mu) = 0, \tag{15}$$

where we have written $v|_{I \times \Sigma_\bullet} = \mu + \theta\, dt$, and so the second two equations are equations on $\Sigma_\bullet$. The first two equations imply that μ descends to a section of the harmonic bundle $H_\alpha \to I$ with Lagrangian boundary conditions (i.e., an element in the domain of $\mathcal{D}_{0,a}$). The third equation implies that, when taken modulo the linearized gauge action, this section lies in the kernel of $\mathcal{D}_{0,a}$. We have assumed that $\mathcal{D}_{0,a}$ is injective, so it follows that $\mu = d_\alpha \chi$ and $v|_{Y_\bullet} = d_a x$ are both exact. Here x is a 0-form on $Y_\bullet$ and χ is a map from I into the space of 0-forms on $\Sigma_\bullet$. Plugging $\mu = d_\alpha \chi$ back into the third equation in (15) gives

$$0 = d_\alpha \nabla_t \chi - d_\alpha \theta - dX_\alpha(d_\alpha \chi) + [\beta_t, \chi] = d_\alpha(\nabla_t \chi - \theta),$$

where we used the H-flat condition and (14) again. Since α is irreducible, this implies $\nabla_t \chi = \theta$. This tells us that if we extend the definition of x to $I \times \Sigma_\bullet$ by $x|_{\{t\} \times \Sigma_\bullet} = \chi(t)$, then we have $v = d_a x$, which is an equation on the full 3-manifold Y. Finally, use the second equation in $\mathcal{D}_{\epsilon,a}(v, r) = 0$ to get

$$\|v\|^2_{L^2(Y),\epsilon} = (v, d_a x)_\epsilon = (d_a^{*_\epsilon} v, x) = 0.$$

Hence $v = 0$ from which $r = 0$ follows by irreducibility. This proves $\mathcal{D}_{\epsilon,a}$ is injective.

To prove the converse, suppose ξ lies in the kernel of $\mathcal{D}_{0,a}$. Since ξ is a harmonic section with Lagrangian boundary conditions it can be represented by a 1-form v_0 that satisfies (15). Then for any 0-form r, the quantity $v_0 - d_a r$ continues to represent ξ and satisfy (15). Since all H-flat connections are irreducible, we can choose r so that $v := v_0 - d_a r$ lies in the kernel of $d_a^{*_\epsilon}$. Then $(v, 0)$ lies in the kernel of $\mathcal{D}_{\epsilon,a}$. This implies $\xi = 0$ when $\mathcal{D}_{\epsilon,a}$ is injective.

□

2.3.2 Instanton Floer cohomology

The instanton Floer cohomology associated to $Q \to Y$ is the homology of a chain complex $(CF^\bullet_{\text{inst}}, \partial_{\text{inst}})$ that arises as the $\mathcal{G}_\Sigma$-equivariant Morse cohomology of $\mathcal{CS}_H$. It follows that the chain group $CF^\bullet$ is generated by the space $\mathcal{A}_{\text{flat}}(Q, H)/\mathcal{G}_\Sigma$. The boundary operator ∂_{inst} counts $\text{Maps}(\mathbb{R}, \mathcal{G}_\Sigma)$-equivalence classes of solutions $A = a + p\, ds$ to the $(ds^2 + g_\epsilon, H)$-instanton equation

$$\partial_s a - d_a p + *^Y_\epsilon (F_a - \widetilde{X}(a)) = 0.$$

The linearization of this equation at $A = a + p\, ds$ is the operator

$$\mathcal{D}_{\epsilon,A} := \nabla_s + \mathcal{D}_{\epsilon,a},$$

where $\mathcal{D}_{\epsilon,a}$ is as above. On $\mathbb{R}\times Y_\bullet$ this operator takes the form

$$\mathcal{D}_{\epsilon,A}V|_{\mathbb{R}\times Y_\bullet} = \begin{pmatrix} \nabla_s + \epsilon^{-1} *^Y d_a & -d_a \\ -\epsilon^{-2} d_a^* & \nabla_s \end{pmatrix} \begin{pmatrix} v \\ r \end{pmatrix},$$

where all Hodge stars are defined by the fixed metric, so all ϵ-dependence is explicit. Similarly on $\mathbb{R}\times I\times\Sigma_\bullet$, write $V|_{\{(s,t)\}\times\Sigma_\bullet} = \mu(s,t) + \rho(s,t)\,ds + \theta(s,t)\,dt$. Then we have

$$\mathcal{D}_{\epsilon,A}V|_{\mathbb{R}\times I\times\Sigma_\bullet} = \begin{pmatrix} \nabla_s + *^\Sigma(\nabla_t - dX_\alpha) & -d_\alpha & -*^\Sigma d_\alpha \\ \epsilon^{-2} *^\Sigma d_\alpha & \nabla_s & -\nabla_t \\ -\epsilon^{-2} d_\alpha^* & \nabla_t & \nabla_s \end{pmatrix} \begin{pmatrix} \mu \\ \rho \\ \theta \end{pmatrix};$$

all are defined by the fixed metric here as well. A natural domain for $\mathcal{D}_{\epsilon,A}$ is the $W^{1,2}$-completion of the set of compactly supported ϵ-smooth 1-forms on $\mathbb{R}\times Y$; the associated codomain is the L^2-completion. Suppose $a^\pm \in \mathcal{A}_{\text{flat}}(Q,H)$ are non-degenerate H-flat connections and A is any ϵ-smooth $W^{1,2}$-connection on $\mathbb{R}\times Q$ such that

$$\lim_{s\to\pm\infty} A|_{\{s\}\times Y} = a^\pm.$$

Then $\mathcal{D}_{\epsilon,A}$ is Fredholm. When all H-flat connections are non-degenerate, we may therefore use the index of $\mathcal{D}_{\epsilon,A}$ to define the relative grading on the generating set $\mathcal{A}_{\text{flat}}(Q,H)/\mathcal{G}_\Sigma$:

$$\mu_{\text{inst}}([a^-],[a^+]) := \text{Ind}\,\mathcal{D}_{\epsilon,A} \quad \mod 4.$$

Here the need for the reduction mod 4 comes from the choice in representatives $a^\pm \in [a^\pm]$. The desirable case is when $\mathcal{D}_{\epsilon,A}$ is onto. It is a standard fact that $\mathcal{D}_{\epsilon,A}$ is onto if and only if its formal adjoint $\mathcal{D}^*_{\epsilon,A} := -\nabla_s + \mathcal{D}_{\epsilon,a}$ is injective.

It can be shown that, for a suitably chosen H, the operator $\mathcal{D}_{\epsilon,A}$ is onto for all instantons A and the boundary operator squares to zero $\partial^2_{\text{inst}} = 0$. We then define the *instanton Floer cohomology of* Q by

$$HF^\bullet_{\text{inst}}(Q) := \ker\partial_{\text{inst}}/\operatorname{im}\partial_{\text{inst}};$$

see [11, 12], and also [8] for an extension to higher rank Lie groups. As with Lagrangian Floer cohomology, the boundary operator ∂_{inst} only counts instantons A with $\text{Ind}\,\mathcal{D}_{\epsilon,A} = 1$, and so the chain level relative grading descends to a relative grading on $HF^\bullet_{\text{inst}}(Q)$.

3 The index relation

For this section we fix a $\mathrm{PU}(r)$-bundle Q over a broken circle fibration $f: Y \to S^1$ as in Section 2.1. We also fix a perturbation $H : \mathcal{A}(Q) \to \mathbb{R}$ of the form (13).

Let $a^\pm \in \mathcal{A}_{\text{flat}}(Q,H)$ be H-flat connections, and assume H has been chosen so that these are non-degenerate. Then if $A \in \mathcal{A}(\mathbb{R}\times Q)$ is any ϵ-smooth connection limiting to $a^\pm$ at $\pm\infty$, the index $\text{Ind}\,(\mathcal{D}_{\epsilon,A})$ is well-defined and independent of the choice of A, ϵ. If, in addition, we assume that A represents a strip $\mathbb{R}\times I \to M$ with Lagrangian boundary conditions, then $\text{Ind}\,(\mathcal{D}_{0,A})$ is also well-defined. Our main theorem is that these indices are the same.

Theorem 3.1. *(Index Relation) With A, H as above, the Fredholm indices*

$$\mathrm{Ind}\left(\mathcal{D}_{0,A}\right) = \mathrm{Ind}\left(\mathcal{D}_{\epsilon,A}\right)$$

agree for all $\epsilon > 0$.

We prove this in Section 3.3. It follows immediately from the definitions that the gradings of both Floer theories agree.

Corollary 3.2. *Suppose H has been chosen so all H-flat connections are non-degenerate. Then the map Ψ from (8) respects the relative $\mathbb{Z}_4$-gradings in the sense that*

$$\mu_{\mathrm{inst}}\left(\left[a^-\right], \left[a^+\right]\right) = \mu_{\mathrm{symp}}\left(\Psi\left(\left[a^-\right]\right), \Psi\left(\left[a^+\right]\right)\right) \in \mathbb{Z}_4$$

for all $a^{\pm} \in \mathcal{A}_{\mathrm{flat}}(Q, H)$.

Most of the hard work is carried out in Sections 3.1 and 3.2 where we establish several elliptic estimates with ϵ-independent constants. We also prove that $\mathcal{D}_{\epsilon,A}$ is onto whenever $\mathcal{D}_{0,A}$ is onto and ϵ is sufficiently small. When this is the case, the index of each operator is just the dimension of its kernel, so to prove Theorem 3.1 it suffices to show the kernels have the same dimension. We prove this in Section 3.3 by constructing an isomorphism between the kernels.

3.1 Elliptic estimates on surfaces and cobordisms

This section establishes several standard estimates on surfaces and cobordisms, but with respect to the ϵ-dependent metric defined in the introduction to Section 2. Our primary interest is in the dependence of the constants on ϵ.

Given a flat connection A on a bundle $P \to X$ over an oriented Riemannian manifold X, we denote by

$$\mathrm{proj}_A : \Omega^1(X, P(\mathfrak{g})) \longrightarrow \ker(d_A \oplus d_A^* \oplus \partial *) \cong H_A^1$$

the L^2-orthogonal projection to the space of A-harmonic 1-form representatives.

Remark 3.3. *(a) This operator is independent of constant conformal scaling of the metric on X. In particular, taking $X = \Sigma_\bullet$ or $Y_\bullet$, the projection operator associated to the ϵ-dependent metric is identical to that obtained from the fixed metric.*

(b) If a is a small curvature connection on Y then it is possible that the space $H_a^1 = H_a^1(Y_\bullet)$ of a-harmonic forms on $Y_\bullet$ could be confused with the space $H_a^1(Y)$ of a-harmonic forms on Y. In this paper the latter space will not arise directly. We therefore take the liberty to write H_a^1 with the understanding that it always refers to $H_a^1(Y_\bullet)$. Similarly, proj_a will always refer to the projection to $H_a^1(Y_\bullet)$.

(c) Suppose $V = v + r\, ds$ is a 1-form on $\mathbb{R} \times Y$. Then we will sometimes write $\mathrm{proj}_a V$ for the path of 1-forms on $Y_\bullet$ given at $s \in \mathbb{R}$ by $\mathrm{proj}_a(v|_{\{s\} \times Y_\bullet})$. We will similarly write $\mathrm{proj}_\alpha V$ or $\mathrm{proj}_\alpha v$ for the strip of 1-forms on $\mathbb{R} \times I \times \Sigma_\bullet$ given at (s,t) by projecting the $\Sigma_\bullet$-component of $V|_{\{(s,t)\} \times \Sigma_\bullet}$ to the $\alpha(s,t)$-harmonic space.

The next lemma shows that, relative to the ϵ-dependent metric, the non-harmonic part of a form on $\Sigma_\bullet$ can be controlled (with small constants) by its covariant derivatives. This shows that the harmonic part controls the behavior of a form, when ϵ is small. We give a similar statement for $Y_\bullet$, except there is a boundary term that shows up as well.

Lemma 3.4. *There is some $C > 0$ such that if $\alpha \in \mathcal{A}_{\mathrm{flat}}(P_\bullet)$ is any flat connection, then*

$$\begin{array}{rcl} \|\rho\|_{L^2(\Sigma_\bullet),\epsilon} & \leq & \epsilon C \|d_\alpha \rho\|_{L^2(\Sigma_\bullet),\epsilon} \\ \|\mu\|_{L^2(\Sigma_\bullet),\epsilon} & \leq & C\Big(\|\mathrm{proj}_\alpha \mu\|_{L^2(\Sigma_\bullet),\epsilon} + \epsilon\|d_\alpha \mu\|_{L^2(\Sigma_\bullet),\epsilon} + \epsilon \|d_\alpha^{*_\epsilon}\mu\|_{L^2(\Sigma_\bullet),\epsilon}\Big) \end{array}$$

for all $\rho \in \Omega^0(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$, $\mu \in \Omega^1(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$, and $\epsilon > 0$. Similarly, if $a \in \mathcal{A}_{\mathrm{flat}}(Q_\bullet)$ is any flat connection, then

$$\begin{array}{rcl} \|r\|_{L^2(Y_\bullet),\epsilon} & \leq & \epsilon C\|d_a r\|_{L^2(Y_\bullet),\epsilon} \\ \|v\|_{L^2(Y_\bullet),\epsilon} & \leq & C\Big(\|\mathrm{proj}_a v\|_{L^2(Y_\bullet),\epsilon} + \epsilon\|d_a v\|_{L^2(Y_\bullet),\epsilon} + \epsilon\|d_a^{*_\epsilon} v\|_{L^2(Y_\bullet),\epsilon} \\ & & \qquad + \epsilon^{1/2}\|(*^Y v)|_{\partial Y_\bullet}\|_{W^{1/2,2}(\partial Y_\bullet)}\Big) \end{array} \qquad (16)$$

for all $r \in \Omega^0(Y_\bullet, Q_\bullet(\mathfrak{g}))$, $v \in \Omega^1(Y_\bullet, Q_\bullet(\mathfrak{g}))$, and $\epsilon > 0$.

We emphasize that, in the last line of (16), the Hodge star $*^Y$ is on Y and relative to the fixed metric; the $W^{1/2,2}$-norm is relative to the fixed metric as well (we have chosen to not explicitly define ϵ-dependent fractional Sobolev norms). All other norms are relative to the ϵ-dependent metric.

Proof of Lemma 3.4. The Fredholm theory for covariant derivatives d_A implies that, when X is compact, there is a constant C so that

$$\|V\|_{W^{1,p}(X)} \leq C\left(\|\mathrm{proj}_A V\|_{L^p(X)} + \|d_A V\|_{L^p(X)} + \|d_A^* V\|_{L^p(X)} + \|(*^X V)\|_{W^{1/p,p}(\partial X)}\right)$$

for all forms V and all flat connections A. (Note that when A is irreducible and V is a 0-form, the only non-zero term on the right-hand side is $d_A V$.) Then the lemma follows from this observation together with the conformal scaling properties of the L^2-norm on 2- and 3-manifolds. For example,

$$\|\rho\|_{L^2(\Sigma_\bullet),\epsilon} = \epsilon\|\rho\|_{L^2(\Sigma_\bullet)} \leq \epsilon C\|d_\alpha\rho\|_{L^2(\Sigma_\bullet)} = \epsilon C\|d_\alpha \rho\|_{L^2(\Sigma_\bullet),\epsilon}.$$

That the constants can be chosen to be independent of α, a follows from the compactness of the moduli spaces of flat connections on $\Sigma_\bullet$ and $Y_\bullet$. □

In preparation for a boundary-value problem, we want estimates that allow us to pass from boundary- and L^p-norms to $W^{1,2}$-norms. The relevant estimates are supplied by the next two results. We state them relative to the fixed metric, since that is the setting in which they will be used.

Lemma 3.5. *For any $\delta > 0$, there is a constant C_δ so that*

$$\begin{aligned}\|\nu\|_{L^2(\partial(I\times\Sigma_\bullet))} \quad \leq \quad & C_\delta\|\nu\|_{L^2(I\times\Sigma_\bullet)} \\ & +\delta\left(\|\nabla_t\nu\|_{L^2(I\times\Sigma_\bullet)}+\|d_\alpha\nu\|_{L^2(I\times\Sigma_\bullet)}+\|d_\alpha^*\nu\|_{L^2(I\times\Sigma_\bullet)}\right)\end{aligned}$$

for all smooth maps $\nu : I \to \Omega^1(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$.

Similarly, given $1 \leq p < 4$ and a compact set $B \subset \mathbb{R}\times I$, there is some constant $C_{p,B,\delta}$ so that

$$\begin{aligned}\|\mu\|_{L^p(B\times\Sigma_\bullet)} \quad \leq \quad & C_{p,B,\delta}\|\mu\|_{L^2(B\times\Sigma_\bullet)}+\delta\Big(\|\nabla_s\mu\|_{L^2(B\times\Sigma_\bullet)}+\|\nabla_t\mu\|_{L^2(B\times\Sigma_\bullet)} \\ & +\|d_\alpha\mu\|_{L^2(B\times\Sigma_\bullet)}+\|d_\alpha^*\mu\|_{L^2(B\times\Sigma_\bullet)}\Big)\end{aligned}$$

for all smooth maps $\mu : B \to \Omega^1(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$. For both estimates, the derivatives are defined relative to any connection on $\mathbb{R}\times Q$ and the constants can be chosen to be independent of the connection chosen.

Proof. Both estimates can be proved as in [17, Lemma 5.1.3]; we sketch a proof of the second estimate for convenience. By standard elliptic theory there is an a priori estimate of the form

$$\begin{aligned}\|\mu\|_{W^{1,2}(B\times\Sigma_\bullet)} \quad \leq \quad & C\Big(\|\mu\|_{L^2(B\times\Sigma_\bullet)}+\|\nabla_s\mu\|_{L^2(B\times\Sigma_\bullet)}+\|\nabla_t\mu\|_{L^2(B\times\Sigma_\bullet)} \\ & +\|d_\alpha\mu\|_{L^2(B\times\Sigma_\bullet)}+\|d_\alpha^*\mu\|_{L^2(B\times\Sigma_\bullet)}\Big)\end{aligned}$$

for some constant C that is independent of μ. The next key point is that the embedding $W^{1,2}(B\times\Sigma_\bullet) \hookrightarrow L^p(B\times\Sigma_\bullet)$ is compact for $1 \leq p < 4$. Suppose the second estimate in the lemma does not hold. Then there is some $\delta > 0$ and a sequence μ_n so that $\|\mu_n\|_{L^p(B\times\Sigma_\bullet)} = 1$ and

$$\begin{aligned}1 \quad \geq \quad & n\|\mu_n\|_{L^2(B\times\Sigma_\bullet)}+\delta\Big(\|\nabla_s\mu_n\|_{L^2(B\times\Sigma_\bullet)}+\|\nabla_t\mu_n\|_{L^2(B\times\Sigma_\bullet)} \\ & +\|d_\alpha\mu_n\|_{L^2(B\times\Sigma_\bullet)}+\|d_\alpha^*\mu_n\|_{L^2(B\times\Sigma_\bullet)}\Big)\end{aligned}$$

for all n. This implies $\|\mu_n\|_{L^2(B\times\Sigma_\bullet)} \to 0$, and that the derivatives of μ are bounded. It follows from the a priori estimate above that the sequence μ_n is bounded in $W^{1,2}$, and so this sequence has a subsequence that converges to zero in L^p. This contradicts $\|\mu_n\|_{L^p(B\times\Sigma_\bullet)} = 1$, and proves the second estimate for some constant.

That this constant can be chosen independent of the connection can be seen as follows: First prove the lemma in the case where μ is a real-valued 1-form, and with the de Rham operator as a connection. The case for more general vector bundles then follows from the real-valued case by Kato's inequality $|d|\mu|| \leq |\nabla\mu|$, which holds for any metric connection ∇. □

Corollary 3.6. *Let $p \in [1,\infty]$. Then for any $\delta > 0$ there is a constant C_δ with*

$$\begin{aligned}\|\mathrm{proj}_a v\|_{L^p(Y_\bullet)} \quad \leq \quad & C_\delta\|\mu\|_{L^2(I\times\Sigma_\bullet)} \\ & +\delta\left(\|\nabla_t\mu\|_{L^2(I\times\Sigma_\bullet)}+\|d_\alpha\mu\|_{L^2(I\times\Sigma_\bullet)}+\|d_\alpha^*\mu\|_{L^2(I\times\Sigma_\bullet)}\right)\end{aligned}$$

for all ϵ-smooth 1-forms v.

Proof. The map $h \mapsto \|h\|_{L^2(\partial Y_\bullet)}$ is a norm on the a-harmonic space $H^1_a(Y_\bullet)$ since a harmonic form vanishes if and only if it vanishes on the boundary. Since this harmonic space is finite-dimensional, the norm $\|h\|_{L^2(\partial Y_\bullet)}$ is equivalent to the $L^p(Y_\bullet)$-norm, and so for all $v \in \Omega^1(Y_\bullet, Q_\bullet(\mathfrak{g}))$ we have

$$\|\mathrm{proj}_a v\|_{L^p(Y_\bullet)} \leq C\|\mathrm{proj}_a v\|_{L^2(\partial Y_\bullet)}.$$

The result now follows from the first estimate in Lemma 3.5. □

The results of Lemma 3.4 and Corollary 3.6 can be packaged neatly to give the following bounds for ϵ-dependent norms on the full 3-manifold Y.

Corollary 3.7. *Assume $a \in \mathcal{A}(Q)$ represents a path $(I, 0, 1) \to (M, L_{(0)}, L_{(1)})$. Then there are constants C and $\epsilon_0 > 0$ so that*

$$\begin{aligned} \|r\|_{L^2(Y),\epsilon} &\leq \epsilon C\|d_a r\|_{L^2(Y),\epsilon} \\ \|v\|_{L^2(Y),\epsilon} &\leq C\Big(\|\mathrm{proj}_\alpha v\|_{L^2(I\times\Sigma_\bullet),\epsilon} + \epsilon\|d_a v\|_{L^2(Y),\epsilon} + \epsilon\|d_a^{*_\epsilon} v\|_{L^2(Y),\epsilon}\Big), \end{aligned}$$

for all $0 < \epsilon < \epsilon_0$, and for all ϵ-smooth $r \in \Omega^0(Y, Q(\mathfrak{g}))$ and $v \in \Omega^1(Y, Q(\mathfrak{g}))$.

The next two results provide more refined estimates. The first will be useful for bounding surface derivatives in terms of derivatives on the full 3-manifold Y.

Proposition 3.8. *(Surface to 3-manifold estimates) Assume $a \in \mathcal{A}(Q)$ represents a path in M with Lagrangian boundary conditions, and write $a|_{I\times\Sigma_\bullet} = \alpha + \psi\, dt$. Then there are constants C and $\epsilon_0 > 0$ so that*

$$\begin{aligned} &\|d_\alpha \mu\|_{L^2(I\times\Sigma_\bullet),\epsilon} + \|d_\alpha^{*_\epsilon} \mu\|_{L^2(I\times\Sigma_\bullet),\epsilon} + \|\nabla_t \mu\|_{L^2(I\times\Sigma_\bullet),\epsilon} \\ &\quad + \|d_\alpha \theta\|_{L^2(I\times\Sigma_\bullet),\epsilon} + \|\nabla_t \theta\|_{L^2(I\times\Sigma_\bullet),\epsilon} \\ &\leq C\left(\|v\|_{L^2(Y),\epsilon} + \|d_a v\|_{L^2(Y),\epsilon} + \|d_a^{*_\epsilon} v\|_{L^2(Y),\epsilon}\right) \end{aligned}$$

for all $0 < \epsilon < \epsilon_0$ and all ϵ-smooth 1-forms v on Y with $v|_{I\times\Sigma_\bullet} = \mu + \theta\, dt$

See [9] for a proof. Our last result of this section addresses the region complementary to that addressed by Proposition 3.8 above. It can be viewed as a sort of extension theorem stating that a function cannot behave too wildly on $Y_\bullet$ (relative to the fixed metric) if it has an extension to all of Y that has bounded ϵ-dependent $W^{1,2}$-norm.

Proposition 3.9. *Let $1 \leq p < 6$, and suppose $a \in \mathcal{A}(Q)$ represents a path in M with Lagrangian boundary conditions. Then for every $\delta > 0$, there are constants $C, \epsilon_0 > 0$ so that*

$$\|v\|_{L^p(Y_\bullet)} \leq C\|v\|_{L^2(Y),\epsilon} + \delta\left(\|d_a v\|_{L^2(Y),\epsilon} + \|d_a^{*_\epsilon} v\|_{L^2(Y),\epsilon}\right)$$

for all $0 < \epsilon < \epsilon_0$, and all ϵ-smooth $v \in \Omega^1(Y, Q(\mathfrak{g}))$.

We emphasize that the norm on the left is on $Y_\bullet$ and relative to the fixed metric, while the norms on the right are on all of Y and relative to the ϵ-dependent metric.

Proof of Proposition 3.9. Suppose the estimate does not hold. Then there is some $\delta_0 > 0$ and sequences ϵ_n, v_n with

$$\epsilon_n \to 0, \qquad \|v_n\|_{L^p(Y_\bullet)} = 1,$$

$$1 \geq n\|v_n\|_{L^2(Y),\epsilon_n} + \delta_0 \left(\|d_a v_n\|_{L^2(Y),\epsilon_n} + \|d_a^{*\epsilon_n} v_n\|_{L^2(Y),\epsilon_n} \right).$$

Then on $Y_\bullet$, and for ϵ_n small, we have

$$\begin{aligned} \|v_n\|_{W^{1,2}(Y_\bullet)} &\leq C_0 \left(\|v_n\|_{L^2(Y_\bullet)} + \|d_a v_n\|_{L^2(Y_\bullet)} + \|d_a^* v_n\|_{L^2(Y_\bullet)} \right) \\ &\leq C_1 \left(\|v_n\|_{L^p(Y_\bullet)} + \epsilon_n^{1/2} \|d_a v_n\|_{L^2(Y_\bullet),\epsilon_n} + \epsilon_n^{1/2} \|d_a^{*\epsilon_n} v_n\|_{L^2(Y_\bullet),\epsilon_n} \right) \\ &\leq C_1 (1 + \epsilon_n^{1/2} \delta_0^{-1}). \end{aligned}$$

Hence v_n has a subsequence that converges weakly in $W^{1,2}(Y_\bullet)$ and strongly in $L^p(Y_\bullet)$ for $1 \leq p < 6$ to some limit $v_\infty \in \Omega^1(Y_\bullet, Q_\bullet(\mathfrak{g}))$. In particular,

$$\|v_\infty\|_{L^p(Y_\bullet)} = 1.$$

To obtain a contradiction, we will show $v_\infty = 0$. To see this, first note that the above also shows

$$d_a v_\infty = 0, \qquad d^* v_\infty = 0.$$

It is not hard to show from the Hodge decomposition (5) that v_∞ is zero if and only if it restricts to zero on the boundary (this is just saying that the Dirichlet problem has a unique solution, which is zero in our case). We therefore aim to prove

$$v_\infty|_{\partial Y_\bullet} = 0.$$

The weak $W^{1,2}$-convergence implies strong L^2 convergence on the boundary, and so

$$\|v_\infty\|_{L^2(\partial Y_\bullet)} = \lim_n \|v_n\|_{L^2(\partial Y_\bullet)},$$

after possibly passing to a subsequence. Write $v_n = \mu_n + \theta_n \, dt$ on $I \times \Sigma_\bullet$. Let $C_{3.8}$ be the constant from Proposition 3.8. Fix any $\eta_0 > 0$, and apply Corollary 3.6 with

$$\delta := \frac{\eta_0 \delta_0}{2C_{3.8}}$$

to get some constant C_δ for which

$$\begin{aligned} \|v_n\|_{L^2(\partial Y_\bullet)} \leq\; & C_\delta \|\mu_n\|_{L^2(I\times\Sigma_\bullet)} \\ & + \delta \Big(\|\nabla_t \mu_n\|_{L^2(I\times\Sigma_\bullet)} + \|d_\alpha \mu_n\|_{L^2(I\times\Sigma_\bullet)} + \|d_\alpha^* \mu_n\|_{L^2(I\times\Sigma_\bullet)} \Big). \end{aligned}$$

Now convert to the ϵ_n-dependent norm, and use Proposition 3.8 to continue this as follows

$$\begin{aligned}
&= C_\delta \|\mu_n\|_{L^2(I\times\Sigma_\bullet),\epsilon_n} \\
&\quad +\delta\Big(\|\nabla_t\mu_n\|_{L^2(I\times\Sigma_\bullet),\epsilon_n}+\epsilon_n\|d_\alpha\mu_n\|_{L^2(I\times\Sigma_\bullet),\epsilon_n}+\epsilon_n\|d_\alpha^{*\epsilon_n}\mu_n\|_{L^2(I\times\Sigma_\bullet),\epsilon_n}\Big) \\
&\le \left(C_\delta+\frac{\eta_0\delta_0}{2}\right)\|v_n\|_{L^2(Y),\epsilon_n}+\frac{\eta_0\delta_0}{2}\Big(\|d_a v_n\|_{L^2(Y),\epsilon_n}+\|d_a^{*\epsilon_n}v_n\|_{L^2(Y),\epsilon_n}\Big), \\
&\le \left(C_\delta+\frac{\eta_0\delta_0}{2}\right)\frac{1}{n}+\frac{\eta_0}{2} \\
&\le \eta_0
\end{aligned}$$

where the last line follows provided n is large. This is true for all $\eta_0 > 0$, and so

$$\lim_n \|v_n\|_{L^2(\partial Y_\bullet)} = 0.$$

This implies $v_\infty = 0$, which is the desired contradiction. □

3.2 Uniform elliptic estimates for $\mathcal{D}_{\epsilon,A}$

We refer freely to the notation established in Section 2. In particular, we identify 1-forms $V = v + r\,ds$ on $\mathbb{R}\times Y$ with tuples (v, r) of (paths of) 1- and 0-forms on Y. Fix a holomorphic strip representative $A \in \mathcal{A}(\mathbb{R}\times Q)$, as well as a perturbation H as in Section 2.3. Unless otherwise stated, we make no regularity assumptions on H. Let $\mathcal{D}_{\epsilon,A}$ be the linearized instanton operator, and set

$$\widetilde{V} = (\widetilde{v}, \widetilde{r}) := \mathcal{D}_{\epsilon,A} V.$$

for a 1-form V. Breaking this up in components, the above equation becomes

$$\widetilde{v} = \nabla_s v + *_\epsilon d_{a,H} v - d_a r \tag{17}$$
$$\widetilde{r} = \nabla_s r - d_a^{*\epsilon} v, \tag{18}$$

Here we have set

$$d_{a,H} := d_a - dK_a$$

which is just the linearization of the map $a \mapsto F_a - K_a$; consequently we will only apply $d_{a,H}$ to 1-forms on Y. It is convenient to use the following ϵ-dependent $W^{1,2}$-norm on the tangent space to $\mathcal{A}(\mathbb{R}\times Q)$:

$$\|V\|^2_{W^{1,2}(\mathbb{R}\times Y),\epsilon} := \begin{array}{l}\|V\|^2_{L^2(\mathbb{R}\times Y),\epsilon}+\|\nabla_s r\|^2_{L^2(\mathbb{R}\times Y),\epsilon}+\|d_a r\|^2_{L^2(\mathbb{R}\times Y),\epsilon} \\ \quad +\|\nabla_s v\|^2_{L^2(\mathbb{R}\times Y),\epsilon}+\|d_{a,H} v\|^2_{L^2(\mathbb{R}\times Y),\epsilon}+\|d_a^{*\epsilon} v\|^2_{L^2(\mathbb{R}\times Y),\epsilon}.\end{array}$$

Remark 3.10. *(a) We want to keep track of the extent to which our constants depend on the connection A. To do this, let $U \subset \mathbb{R}$ be a measurable subset, and define c_A^U to be the maximum of the values*

$$\begin{array}{cc} \sup_{U\times I}\|\partial_s\alpha - d_\alpha\phi\|_{L^2(\Sigma_\bullet)}, & \sup_{U\times I}\|\partial_t\alpha - d_\alpha\psi - X_\alpha\|_{L^2(\Sigma_\bullet)}, \\ \|\partial_s\psi - \partial_t\phi - [\psi,\phi]\|_{L^4(U\times I\times\Sigma_\bullet)}, & \sup_U\|\partial_s a - d_a p\|_{L^2(Y_\bullet)}. \end{array} \tag{19}$$

We set $c_A := c_A^{\mathbb{R}}$.

(b) The significance of the quantities in (19) is that we have uniform estimates on these for ϵ*-ASD connections* A*; see [9]. In subsequent work, we intend to apply the results of Theorem 3.11 and Corollary 3.12 to the case where* A *is interpreted as an* ϵ*-ASD connection. In keeping track of the constants in (19), we ensure that the analysis here carries over the* ϵ*-ASD case with little extra work.*

The goal of this section is to prove several elliptic estimates for the linearized instanton operator. The first shows that the standard Fredholm estimate for $\mathcal{D}_{\epsilon,A}$, holds with an ϵ-independent constant.

Theorem 3.11. *Fix a perturbation* H*. There are constants* C *and* $\epsilon_0 > 0$ *with the following significance. Suppose* A *represents a strip* $\mathbb{R} \times I \to M$ *with Lagrangian boundary conditions, and let* c_A *be as in Remark 3.10. Then*

$$\|V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon} \quad \leq \quad C\Big(\|\mathcal{D}_{\epsilon,A}V\|_{L^2(\mathbb{R}\times Y),\epsilon} + c_A\|\mathrm{proj}_\alpha V\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}\Big), \tag{20}$$

for all $0 < \epsilon < \epsilon_0$*, and all* ϵ*-smooth 1-forms* V*. The same result holds with* $\mathcal{D}_{\epsilon,A}$ *replaced by its formal adjoint* $\mathcal{D}^*_{\epsilon,A}$*.*

This theorem is proved in Section 3.2.1. In Section 3.2.2, we prove the following corollary that improves the estimate in (20) by replacing the L^2-norm on the right with the L^2-norm over a compact set. This will be useful in our compactness arguments later.

Corollary 3.12. *Fix a perturbation* H *and a constant* c_0*. Then there are constants* $C, \delta_0, \epsilon_0 > 0$ *so the following holds. Suppose* A *represents a holomorphic strip with Lagrangian boundary conditions, and satisfies*

$$c_A^{\mathbb{R}} \leq c_0, \qquad \text{and} \qquad c_A^{(-\infty,-S_0]} + c_A^{[S_0,\infty)} < \delta_0$$

for some S_0*, where* c_A^U *is the constant from Remark 3.10. Then*

$$\|V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon} \quad \leq \quad C\Big(\|\mathcal{D}_{\epsilon,A}V\|_{L^2(\mathbb{R}\times Y),\epsilon} + c_A^{\mathbb{R}}\|\mathrm{proj}_\alpha V\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet),\epsilon}\Big),$$

for all $0 < \epsilon < \epsilon_0$*, and all* ϵ*-smooth 1-forms* V*. The same result holds with* $\mathcal{D}_{\epsilon,A}$ *replaced by its formal adjoint* $\mathcal{D}^*_{\epsilon,A}$*.*

Our final result shows that the operators $\mathcal{D}_{0,A}$ and $\mathcal{D}_{\epsilon,A}$ can be made simultaneously surjective when ϵ is sufficiently small. Its proof appears in Section 3.2.3.

Theorem 3.13. *Let* H, A *be as in the statement of Corollary 3.12. Assume further that the operator* $\mathcal{D}_{0,A}$ *is onto when restricted to sections with Lagrangian boundary conditions. Then* $\mathcal{D}_{\epsilon,A}$ *is onto. Moreover, there is a constant* C *such that*

$$\|V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon} \quad \leq \quad C\|\mathcal{D}^*_\epsilon V\|_{L^2(\mathbb{R}\times Y),\epsilon} \tag{21}$$

$$\|\mathcal{D}^*_{\epsilon,A}V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon} \quad \leq \quad C\|\mathcal{D}_{\epsilon,A}\mathcal{D}^*_{\epsilon,A}V\|_{L^2(\mathbb{R}\times Y),\epsilon}, \tag{22}$$

for all $0 < \epsilon < \epsilon_0$*, and all* ϵ*-smooth 1-forms* V*.*

3.2.1 Proof of Theorem 3.11

We prove the theorem for

$$\mathcal{D}_\epsilon := \mathcal{D}_{\epsilon,A};$$

the result for $\mathcal{D}^*_\epsilon$ is similar (e.g., changing the orientation on Y changes $\mathcal{D}_\epsilon$ to $-\mathcal{D}^*_\epsilon$). Throughout this proof, we use

$$\|\cdot\|_\epsilon, \quad \text{and} \quad (\cdot,\cdot)_\epsilon$$

to denote the L^2-norm and inner product on the 4-manifold $\mathbb{R} \times Y$ defined with respect to the ϵ-dependent metric. When $\epsilon = 1$, we will write $\|\cdot\| := \|\cdot\|_1$, etc. for the norms, etc. coming from the fixed metric. However, it is convenient to divert slightly from our usual convention and use $*_\epsilon$ to denote the Hodge star on the *3-manifold* Y. For example, if $x, y : \mathbb{R} \to \Omega^\bullet(Y, Q(\mathfrak{g}))$ are paths of forms on Y viewed as forms on $\mathbb{R} \times Y$, then $(x,y)_\epsilon = \int_{\mathbb{R}\times Y} ds \wedge \langle x(s) \wedge *_\epsilon y(s)\rangle$. All constants C, C' that appear below depend only on the perturbation H and the fixed metric on Y.

Let V be an ϵ-smooth 1-form, and set

$$\widetilde{V} := \mathcal{D}_\epsilon V.$$

To prove the theorem, we note that it suffices to show

$$\|\nabla_s v\|^2_\epsilon + \|d_a v\|^2_\epsilon + \|d_a^{*_\epsilon} v\|^2_\epsilon + \|\nabla_s r\|^2_\epsilon + \|d_a r\|^2_\epsilon \leq C\left(\|\widetilde{V}\|^2_\epsilon + c_A \|V\|^2_\epsilon\right). \tag{23}$$

This is because

$$\|V\|^2_\epsilon = \|V - \mathrm{proj}_\alpha V\|^2_\epsilon + \|\mathrm{proj}_\alpha V\|^2_\epsilon,$$

and Corollary 3.7 can be used to absorb the norm of $V - \mathrm{proj}_\alpha V$ into the left-hand side of (23) for ϵ sufficiently small. Our strategy for proving (23) is to integrate by parts to control the derivatives of v and r in terms of $\widetilde{V}$ plus some cross terms.

Bounds for the derivatives of v:

Apply d_a to (18) to get

$$\begin{aligned} d_a d_a^{*_\epsilon} v &= -d_a \widetilde{r} + d_a \nabla_s r \\ &= -d_a \widetilde{r} + \nabla_s d_a r - [b_s, r] \\ &= -d_a \widetilde{r} - \nabla_s \widetilde{v} + \nabla_s^2 v + *_\epsilon \nabla_s d_{a,H} v - [b_s, r], \end{aligned}$$

where we have used (17) in the third equality. Similarly, applying $*_\epsilon d_{a,H}$ to (17) gives

$$\begin{aligned} d_{a,H}^{*_\epsilon} d_{a,H} v &= d_{a,H}^{*_\epsilon} *_\epsilon \widetilde{v} - *_\epsilon d_{a,H} \nabla_s v + *_\epsilon [F_a, r] - *_\epsilon dK_a(d_a r) \\ &= d_{a,H}^{*_\epsilon} *_\epsilon \widetilde{v} - *_\epsilon \nabla_s d_{a,H} v + *_\epsilon [b_s \wedge v] + *_\epsilon d^2 K_a(b_s, v) + *_\epsilon [F_a - K_a, r], \end{aligned}$$

where $d^2 K_a$ is the Hessian of K_a, and we used (14) on the last term. Add these together (the $*_\epsilon \nabla_s d_{a,H} v$ terms cancel):

$$\begin{aligned} -\nabla_s^2 v + d_{a,H}^{*_\epsilon} d_{a,H} v + d_a d_a^{*_\epsilon} v &= -\nabla_s \widetilde{v} + d_{a,H}^{*_\epsilon} *_\epsilon \widetilde{v} - d_a \widetilde{r} \\ &\quad + [*_\epsilon(F_a - K_a) - b_s, r] + *_\epsilon [b_s \wedge v] \\ &\quad + *_\epsilon d^2 K_a(b_s, v). \end{aligned}$$

Now apply $(\cdot, v)_\epsilon$ and integrate by parts to get

$$\begin{aligned}
\|\nabla_s v\|_\epsilon^2 + \|d_{a,H} v\|_\epsilon^2 + \|d_a^{*_\epsilon} v\|_\epsilon^2 &= (\widetilde{v}, \nabla_s v)_\epsilon + (\widetilde{v}, *_\epsilon d_{a,H} v)_\epsilon - \left(\widetilde{r}, d_a^{*_\epsilon} v\right)_\epsilon \\
&\quad + ([*_\epsilon(F_a - K_a) - b_s, r], v)_\epsilon + (*_\epsilon [b_s \wedge v], v)_\epsilon \\
&\quad + (*_\epsilon d^2 K_a(b_s, v), v)_\epsilon \\
&\leq \|\widetilde{V}\|_\epsilon^2 + \frac{1}{2}\left(\|\nabla_s v\|_\epsilon^2 + \|d_{a,H} v\|_\epsilon^2 + \|d_a^{*_\epsilon} v\|_\epsilon^2\right) \\
&\quad + ([*_\epsilon(F_a - K_a) - b_s, r], v)_\epsilon + (*_\epsilon [b_s \wedge v], v)_\epsilon \\
&\quad + (*_\epsilon d^2 K_a(b_s, v), v)_\epsilon.
\end{aligned}$$

Moving the derivative terms to the other side gives

$$\begin{aligned}
\|\nabla_s v\|_\epsilon^2 + \|d_a v\|_\epsilon^2 + \|d_a^{*_\epsilon} v\|_\epsilon^2 &\leq 2\|\widetilde{V}\|_\epsilon^2 + 2\,([*_\epsilon(F_a - K_a) - b_s, r], v)_\epsilon \\
&\quad + 2\,(*_\epsilon [b_s \wedge v], v)_\epsilon + 2(*_\epsilon d^2 K_a(b_s, v), v)_\epsilon.
\end{aligned} \tag{24}$$

It remains to bound the last three terms. Before doing so, we carry out the analogous computation for the derivatives of r.

<u>Bounds for the derivatives of r:</u>

Apply $d_a^{*_\epsilon}$ to (17) to get

$$\begin{aligned}
d_a^{*_\epsilon} d_a r &= -d_a^{*_\epsilon} \widetilde{v} + d_a^{*_\epsilon} \nabla_s v - *_\epsilon [F_a - K_a \wedge v] \\
&= -d_a^{*_\epsilon} \widetilde{v} + \nabla_s d_a^{*_\epsilon} v + *_\epsilon [b_s \wedge *_\epsilon v] - *_\epsilon [F_a - K_a \wedge v] \\
&= -d_a^{*_\epsilon} \widetilde{v} - \nabla_s \widetilde{r} + \nabla_s^2 r + *_\epsilon [b_s \wedge *_\epsilon v] - *_\epsilon [F_a - K_a \wedge v]
\end{aligned}$$

where in the first line we used (14) on the perturbation term, and in the last line we used (18) on $d_a^{*_\epsilon} v$. Now take the L^2-inner product of both sides with r and integrate by parts:

$$\begin{aligned}
\|\nabla_s r\|_\epsilon^2 + \|d_a r\|_\epsilon^2 &= -(\widetilde{v}, d_a r)_\epsilon + (\widetilde{r}, \nabla_s r)_\epsilon + (*_\epsilon b_s - F_a + K_a, *_\epsilon [v, r])_\epsilon \\
&\leq \frac{1}{2}\left(\|\widetilde{v}\|_\epsilon^2 + \|\widetilde{r}\|_\epsilon^2\right) + \frac{1}{2}\left(\|d_a r\|_\epsilon^2 + \|\nabla_s r\|_\epsilon^2\right) \\
&\quad + (*_\epsilon b_s - F_a + K_a, *_\epsilon [v, r])_\epsilon.
\end{aligned}$$

The derivative terms on the right can be subtracted and absorbed on the left to give

$$\|\nabla_s r\|_\epsilon^2 + \|d_a r\|_\epsilon^2 \leq \|\widetilde{V}\|_\epsilon^2 + 2\,(*_\epsilon b_s - F_a + K_a, *_\epsilon [v, r])_\epsilon.$$

Combining this with (24), we have

$$\begin{aligned}
&\|\nabla_s v\|_\epsilon^2 + \|d_a v\|_\epsilon^2 + \|d_a^{*_\epsilon} v\|_\epsilon^2 + \|\nabla_s r\|_\epsilon^2 + \|d_a r\|_\epsilon^2 \\
&\leq 3\|\widetilde{V}\|_\epsilon^2 \\
&\quad + 2(*_\epsilon d^2 K_a(b_s, v), v)_\epsilon + 2\,(*_\epsilon [b_s \wedge v], v)_\epsilon + 3\,(*_\epsilon b_s - F_a + K_a, *_\epsilon [v, r])_\epsilon.
\end{aligned} \tag{25}$$

Our goal now is to estimate the last three terms on the right of (25). We will show that, for each $\delta > 0$, there are constants C_δ and $\epsilon_0 > 0$ so that the quantity

$$c_A C_\delta \|V\|_\epsilon^2 + \delta\left(\|\nabla_s v\|_\epsilon^2 + \|d_a v\|_\epsilon^2 + \|d_a^{*_\epsilon} v\|_\epsilon^2 + \|\nabla_s r\|_\epsilon^2 + \|d_a r\|_\epsilon^2\right)$$

bounds each of the last three terms in (25) for all $0 < \epsilon < \epsilon_0$. The estimate in (23) will then follow by taking δ sufficiently small.

The first term: $(*_\epsilon d^2 K_a(b_s, v), v)_\epsilon$

Recall that $K_a = dt \wedge X(\alpha)$ is supported on $\mathbb{R} \times I \times \Sigma_\bullet$, so the Hessian $d^2 K_a = dt \wedge d^2 X_\alpha$ is just the Hessian of the perturbation X on $\Sigma_\bullet$. Then writing $v = \mu + \theta\, dt$ and $b_s = \beta_s + \gamma\, dt$, we have

$$\begin{aligned}(*_\epsilon d^2 K_a(b_s, v), v)_\epsilon &= \int_{\mathbb{R}\times I\times\Sigma_\bullet} dt \wedge \langle d^2 X_\alpha(\beta_s, \mu) \wedge \mu\rangle \\ &\leq \|d^2 X_\alpha(\beta_s, \mu)\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} \|\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)}.\end{aligned}$$

Note that this quantity is actually independent of ϵ. The norm square of the perturbation is bounded by

$$\begin{aligned}\|d^2 X_\alpha(\beta_s, \mu)\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} &= \int_{\mathbb{R}\times I} \|d^2 X_\alpha(\beta_s, \mu)\|^2_{L^2(\Sigma_\bullet)} \\ &\leq \int_{\mathbb{R}\times I} C_{H,\alpha} \|\beta_s\|^2_{L^2(\Sigma_\bullet)} \|\mu\|^2_{L^2(\Sigma_\bullet)},\end{aligned}$$

where $C_{H,\alpha}$ is the operator norm of $d^2 X_\alpha$ relative to the L^2-topology. Note that since the moduli space of flat connection is compact, there is some C_H so that $C_{H,\alpha} \leq C_H$ for all flat α. In conclusion, using the definition of c_A from Remark 3.10, we have

$$(*_\epsilon d^2 K_a(b_s, v), v)_\epsilon \leq c_A C_H \|V\|_\epsilon^2,$$

as desired.

The second term: $(*_\epsilon [b_s \wedge v], v)_\epsilon$

We have

$$(*_\epsilon [b_s \wedge v], v)_\epsilon = \int_{\mathbb{R}\times I\times\Sigma_\bullet} ds \wedge \langle [b_s \wedge v] \wedge v\rangle + \int_{\mathbb{R}\times Y_\bullet} ds \wedge \langle [b_s \wedge v] \wedge v\rangle. \tag{26}$$

Note that this too is independent of ϵ. First estimate the integral over $\mathbb{R} \times I \times \Sigma_\bullet$. Write $v = \mu + \theta\, dt$ and $b_s = \beta_s + \gamma\, dt$, so

$$\begin{aligned}\int_{\mathbb{R}\times I\times\Sigma_\bullet} ds \wedge \langle [b_s \wedge v] \wedge v\rangle &= 2\int_{\mathbb{R}\times I\times\Sigma_\bullet} ds \wedge dt \wedge \langle [\beta_s \wedge \mu], \theta\rangle \\ &\quad + \int_{\mathbb{R}\times I\times\Sigma_\bullet} ds \wedge dt \wedge \langle [\gamma, \mu] \wedge \mu\rangle.\end{aligned} \tag{27}$$

Focusing on the first term on the right of (27). Here and below the notation

$$L^p(A, L^q(B))$$

stands for the space of L^p-maps from a space A into $L^q(B)$. Use the combined Hölder inequalities

$$\begin{aligned} \|f_1 f_2\|_{L^1(A\times B)} &\leq \|f_1\|_{L^\infty(A,L^2(B))}\|f_2\|_{L^1(A,L^2(B))} \\ \|g_1 g_2\|_{L^1(A,L^2(B))} &\leq \|g_1\|_{L^2(A,L^4(B))}\|g_2\|_{L^2(A,L^4(B))} \end{aligned}$$

for product spaces, as well as the estimate $|\,[\xi,\zeta]\,| \leq 2|\xi||\zeta|$, to write

$$\int_{\mathbb{R}\times I\times\Sigma_\bullet} ds\wedge dt\wedge\langle[\beta_s\wedge\mu]\,,\theta\rangle$$

$$\leq 2\|\beta_s\|_{L^\infty(\mathbb{R}\times I,L^2(\Sigma_\bullet))}\|\mu\|_{L^2(\mathbb{R}\times I,L^4(\Sigma_\bullet))}\|\theta\|_{L^2(\mathbb{R}\times I,L^4(\Sigma_\bullet))}.$$

Now fix $\delta>0$, and use the estimate $ab\leq\frac{1}{2\delta}a^2+\frac{\delta}{2}b^2$ to obtain

$$\int_{\mathbb{R}\times I\times\Sigma_\bullet} ds\wedge dt\wedge\langle[\beta_s\wedge\mu]\,,\theta\rangle$$

$$\leq \|\beta_s\|_{L^\infty(\mathbb{R}\times I,L^2(\Sigma_\bullet))}\left(\delta^{-1}\|\mu\|^2_{L^2(\mathbb{R}\times I,L^4(\Sigma_\bullet))}+\delta\|\theta\|^2_{L^2(\mathbb{R}\times I,L^4(\Sigma_\bullet))}\right).$$

Recall the definition of c_A from Remark 3.10, and use the Sobolev embedding $W^{1,2}\hookrightarrow L^4$ on $\Sigma_\bullet$ to continue this as

$$\begin{aligned} &\leq c_A C\delta^{-1}\left(\|\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)}+\|d_\alpha\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)}+\|d^*_\alpha\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)}\right) \\ &\qquad +\delta c_A C\|d_\alpha\theta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} \\ &= c_A C\delta^{-1}\left(\|\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}+\epsilon^2\|d_\alpha\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}+\epsilon^2\|d^{*_\epsilon}_\alpha\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}\right) \\ &\qquad +\delta c_A C\|d_\alpha\theta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon} \\ &\leq c_A C'\delta^{-1}\|v\|^2_\epsilon + c_A C'(\delta+\epsilon^2\delta^{-1})\left(\|d_a v\|^2_\epsilon+\|d^{*_\epsilon}_a v\|^2_\epsilon\right). \end{aligned}$$

In the second line we converted to the ϵ-dependent metric, and in the last line we used Proposition 3.8. By choosing $\epsilon_0<\delta$, this is the desired estimate for the first term on the right in (27).

For the second term on the right of (27), write

$$\int_{\mathbb{R}\times I\times\Sigma_\bullet} ds\wedge dt\wedge\langle[\gamma,\mu]\wedge\mu\rangle = \sum_{n\in\mathbb{Z}}\int_{[n,n+1]\times I\times\Sigma_\bullet} ds\wedge dt\wedge\langle[\gamma,\mu]\wedge\mu\rangle.$$

Setting $S_n := [n,n+1]\times I\times\Sigma_\bullet$, we have

$$\int_{S_n} ds\wedge dt\wedge\langle[\gamma,\mu]\wedge\mu\rangle \leq 2c_A\|\mu\|^2_{L^{8/3}(S_n)},$$

where we used $\|fg\|_{L^1}\leq\|f\|_{L^4}\|g\|_{L^{4/3}}$ on S_n together with the definition of c_A from Remark 3.10. Then by Lemma 3.5, for any $\delta>0$ there is some C so that this is bounded by

$$c_A C\|\mu\|^2_{L^2(S_n)}+c_A\delta\left(\|\nabla_s\mu\|^2_{L^2(S_n)}+\|\nabla_t\mu\|^2_{L^2(S_n)}+\|d_\alpha\mu\|^2_{L^2(S_n)}+\|d^*_\alpha\mu\|^2_{L^2(S_n)}\right).$$

Clearly C is independent of n due to translation invariance. Converting to the ϵ-dependent norms and summing over n gives

$$\int_{\mathbb{R}\times I\times\Sigma_\bullet} ds\wedge dt\wedge\langle[\gamma,\mu]\wedge\mu\rangle$$

$$\leq c_A C\|\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}+c_A\delta\Big(\|\nabla_s\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}$$
$$+\|\nabla_t\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}+\epsilon^2\|d_\alpha\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}+\epsilon^2\|d_\alpha^{*_\epsilon}\mu\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon}\Big).$$

Now use Proposition 3.8 again to bound these surface derivatives in terms of derivatives on Y. This yields the desired estimate for the integral in (26) over $\mathbb{R}\times I\times\Sigma_\bullet$.

It remains to bound the integral in (26) over the set $\mathbb{R}\times Y_\bullet$. We have

$$\begin{aligned}\int_{\mathbb{R}\times Y_\bullet} ds\wedge\langle[b_s\wedge v]\wedge v\rangle &\leq 2c_A\|v\|^2_{L^2(\mathbb{R},L^4(Y_\bullet))}\\ &\leq 2c_A C_\delta\|v\|^2_\epsilon+2c_A\delta\left(\|d_a v\|^2_\epsilon+\|d_a^{*_\epsilon}v\|^2_\epsilon\right),\end{aligned}$$

where we used Proposition 3.9 in the second step. This is the desired estimate here and completes our analysis of the second term.

<u>The third term</u>: $(*_\epsilon b_s-F_a+K_a,*_\epsilon[v,r])_\epsilon$

Write

$$\begin{aligned}(*_\epsilon b_s-F_a+K_a,*_\epsilon[v,r])_\epsilon &= \int_{\mathbb{R}\times Y_\bullet}\langle *_\epsilon b_s\wedge[v,r]\rangle\\ &\quad+\int_{\mathbb{R}\times I\times\Sigma_\bullet}\langle *_\epsilon b_s-F_a+K_a\wedge[v,r]\rangle.\end{aligned}\tag{28}$$

We prove the relevant estimate for the first term on the right of (28); the second term is similar except one also needs to use Proposition 3.8 to convert from surface derivatives to 3-manifold derivatives, just as we did for the second term above.

Note that on 1-forms on $Y_\bullet$ we have $*_\epsilon=\epsilon *$, so

$$\begin{aligned}|\,(*_\epsilon b_s,*_\epsilon[v,r])_{L^2(\mathbb{R}\times Y_\bullet),\epsilon}\,| &= \epsilon\left|\int_{\mathbb{R}\times Y_\bullet}\langle *b_s\wedge[v,r]\rangle\right|\\ &\leq 2\epsilon c_A\|\,[v,r]\,\|_{L^1(\mathbb{R},L^2(Y_\bullet))}\\ &\leq 2\epsilon c_A\left(\delta^{-1}\|v\|^2_{L^2(\mathbb{R},L^4(Y_\bullet))}+\delta\|r\|^2_{L^2(\mathbb{R},L^4(Y_\bullet))}\right);\end{aligned}$$

in the first inequality we used the definition of c_A, together with a combination of Hölder's inequality $\|fg\|_{L^1}\leq\|f\|_{L^\infty}\|g\|_{L^1}$ in the $\mathbb{R}$-variable with Hölder's inequality $\|fg\|_{L^1}\leq\|f\|_{L^2}\|g\|_{L^2}$ in the $Y_\bullet$-variables. The second inequality used similar Hölder estimates. Now use the Sobolev embedding $W^{1,2}\hookrightarrow L^4$ for $Y_\bullet$, and the irreducibility of a to bound this by

$$\epsilon c_A C\left(\delta^{-1}\left(\|v\|^2_{L^2(\mathbb{R}\times Y_\bullet)}+\|d_a v\|^2_{L^2(\mathbb{R}\times Y_\bullet)}+\|d_a^* v\|^2_{L^2(\mathbb{R}\times Y)}\right)+\delta\|d_a r\|^2_{L^2(\mathbb{R}\times Y_\bullet)}\right).$$

Converting back to the ϵ-dependent norms shows that this term is bounded by

$$c_A C\delta^{-1}\left(\|v\|^2_{L^2(\mathbb{R}\times Y_\bullet),\epsilon}+\epsilon\|d_a v\|^2_{L^2(\mathbb{R}\times Y_\bullet),\epsilon}+\epsilon\|d_a^{*_\epsilon}v\|^2_{L^2(\mathbb{R}\times Y),\epsilon}\right)$$
$$+\delta c_A C\|d_a r\|^2_{L^2(\mathbb{R}\times Y_\bullet),\epsilon},$$

as desired. □

3.2.2 Proof of Corollary 3.12

Theorem 3.11 gives

$$\begin{aligned}\|V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon} &\leq C_{3.11}\left(\|\mathcal{D}^*_\epsilon V\|_\epsilon + c_A^{\mathbb{R}}\|\mathrm{proj}_\alpha V\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)}\right)\\ &\leq C_{3.11}\Big(\|\mathcal{D}^*_\epsilon V\|_\epsilon + c_A^{\mathbb{R}}\|\mathrm{proj}_\alpha V\|_{L^2([-S_0-1,S_0+1]\times I\times\Sigma_\bullet)}\\ &\qquad + c_A^{\mathbb{R}}\|V\|_{L^2([S_0+1,\infty)\times Y),\epsilon} + c_A^{\mathbb{R}}\|V\|_{L^2((-\infty,-S_0-1]\times Y),\epsilon}\Big).\end{aligned}$$

where $C_{3.11}$ is the constant from Theorem 3.11. Focus on the second-to-last term. With the use of a bump function, it is not hard to show that Theorem 3.11 implies

$$\|V\|_{L^2([S_0+1,\infty)\times Y),\epsilon}\leq C\left(\|\mathcal{D}^*_\epsilon V\|_{L^2([S_0,\infty)),\epsilon}+c_A^{[S_0,\infty)}\|V\|_{L^2([S_0,\infty)),\epsilon}\right)$$

for some constant C. A similar estimate holds for $\|V\|_{L^2((-\infty,-S_0-1]\times Y),\epsilon}$. Recall we are assuming

$$c_A^{[S_0,\infty)}+c_A^{(-\infty,-S_0]}<\delta_0,$$

so by taking δ_0 sufficiently small, we have

$$\begin{aligned}\|V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon} &\leq C'\Big(\|\mathcal{D}^*_\epsilon V\|_\epsilon + c_A^{\mathbb{R}}\|\mathrm{proj}_\alpha V\|_{L^2([-S_0-1,S_0+1]\times I\times\Sigma_\bullet)}\Big)\\ &\quad+\tfrac{1}{2}\|V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon}.\end{aligned}$$

This gives

$$\|V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon}\leq 2C'\Big(\|\mathcal{D}^*_\epsilon V\|_\epsilon + c_A^{\mathbb{R}}\|\mathrm{proj}_\alpha V\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet)}\Big),$$

as desired. □

3.2.3 Proof of Theorem 3.13

Set $\mathcal{D}_0 := \mathcal{D}_{0,A}$ and $\mathcal{D}_\epsilon := \mathcal{D}_{\epsilon,A}$. We will write $\|\cdot\|_\epsilon$ and $(\cdot,\cdot)_\epsilon$ for the L^2-norm and inner product on $\mathbb{R}\times Y$ defined using the ϵ-dependent metric.

First we show how (21) can be used to prove (22). By Theorem 3.11 applied to $\mathcal{D}^*_\epsilon V$, we have

$$\|\mathcal{D}^*_\epsilon V\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon}\leq C_1\left(\|\mathcal{D}_\epsilon\mathcal{D}^*_\epsilon V\|_\epsilon+\|\mathrm{proj}_\alpha\mathcal{D}^*_\epsilon V\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)}\right).$$

To verify (22), it therefore suffices to bound

$$\|\mathrm{proj}_\alpha \mathcal{D}_\epsilon^* V\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} \leq \|\mathcal{D}_\epsilon^* V\|_\epsilon$$

in terms of $\|\mathcal{D}_\epsilon\mathcal{D}_\epsilon^* V\|_\epsilon$. To obtain such a bound note that for any $\delta > 0$ we have

$$\|\mathcal{D}_\epsilon^* V\|_\epsilon^2 = (\mathcal{D}_\epsilon\mathcal{D}_\epsilon^* V, V)_\epsilon \leq \frac{1}{2\delta}\|\mathcal{D}_\epsilon\mathcal{D}_\epsilon^* V\|_\epsilon^2 + \frac{\delta}{2}\|V\|_\epsilon^2.$$

Assuming (21) holds, we can continue this to get

$$\|\mathcal{D}_\epsilon^* V\|_\epsilon^2 \leq \frac{1}{2\delta}\|\mathcal{D}_\epsilon\mathcal{D}_\epsilon^* V\|_\epsilon^2 + \frac{C_1^2\delta}{2}\|\mathcal{D}_\epsilon^* V\|_\epsilon^2.$$

Then (22) follows by taking δ small.

To prove the theorem, it therefore suffices to show (21) (note that this implies $\mathcal{D}_\epsilon$ is surjective). In light of Corollary 3.12, to establish (21) it suffices to prove the following estimate

$$\|\mathrm{proj}_\alpha V\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet)} \leq C\|\mathcal{D}_\epsilon^* V\|_\epsilon. \tag{29}$$

If (29) does not hold, then we can find sequences ϵ_n, V_n with

$$\epsilon_n \to 0, \qquad \|\mathrm{proj}_\alpha V_n\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet)} = 1, \qquad \|\mathcal{D}_{\epsilon_n}^* V_n\|_{\epsilon_n} \to 0.$$

Write $V_n = \mu_n + \rho_n\, ds + \theta_n\, dt$ on $\mathbb{R}\times I\times\Sigma_\bullet$. By Corollary 3.12 and Proposition 3.8, there is a bound of the form

$$\begin{aligned}
&\|\nabla_s\mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} + \|\nabla_t\mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} \\
&\qquad + \epsilon_n^{-1}\|d_\alpha\mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} + \epsilon_n^{-1}\|d_\alpha^{*\epsilon_n}\mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} \\
&= \|\nabla_s\mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon_n} + \|\nabla_t\mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon_n} \\
&\qquad + \|d_\alpha\mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon_n} + \|d_\alpha^{*\epsilon_n}\mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon_n} \\
&\leq C_2\left(\|\mathcal{D}_{\epsilon_n}^* V_n\|_{\epsilon_n} + \|\mathrm{proj}_\alpha V_n\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet)}\right) \\
&\leq 2C_2
\end{aligned}$$

provided n is large enough so $\|\mathcal{D}_{\epsilon_n}^* V_n\|_{\epsilon_n} \leq 1$. After passing to a subsequence, it follows that for each compact $B \subset \mathbb{R}$, the sequence μ_n converge weakly in $W^{1,2}(B\times I, L^2(\Sigma_\bullet))$ to some limiting η_∞ with

$$d_\alpha\eta_\infty = 0, \qquad d_\alpha^*\eta_\infty = 0.$$

In fact, writing

$$\eta_n := \mathrm{proj}_\alpha\mu_n$$

it is not hard to see that the difference $\mu_n - \eta_n$ is going to zero in $W^{1,2}(\mathbb{R} \times I, L^2(\Sigma_\bullet))$, and so the η_n are converging to η_∞ weakly in $W^{1,2}$ on compact subsets of $\mathbb{R} \times I$ (e.g., see the proof of Claim 1, below). Moreover, the η_n take values in the finite-dimensional harmonic space, so this weak $W^{1,2}$-convergence implies strong L^2-convergence on compact subsets. In particular, we have

$$\|\eta_\infty\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet)} = \lim_n \|\eta_n\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet)} = 1,$$

which shows η_∞ is non-zero.

Claim 1: η_∞ lies in the kernel of $\mathcal{D}_0^*$.

Claim 2: η_∞ has Lagrangian boundary conditions.

Assuming these claims for now, it follows from the fact that η_∞ is nonzero that the operator $\mathcal{D}_0^*$ is not injective when restricted the space of sections with Lagrangian boundary conditions. However, this implies that its adjoint $\mathcal{D}_0$ is not onto when restricted to the space of sections with Lagrangian boundary conditions. We have assumed otherwise in the statement of Theorem 3.13, and so this contradiction establishes the estimate (29), and hence proves the theorem.

We therefore set out to verify the claims; we begin with Claim 1. Use the Hodge theorem on $\Sigma_\bullet$ to write

$$\mu_n = \eta_n + d_\alpha \zeta_n + * d_\alpha \xi_n$$

for some 0-forms ζ_n, ξ_n on $\Sigma_\bullet$. There is an identity

$$\mathrm{proj}_\alpha \mathcal{D}^*_{\epsilon_n} V_n - \mathcal{D}_0^* \eta_n = \mathrm{proj}_\alpha \omega_n$$

where we have set

$$\omega_n := [\beta_s - *\beta_t, \zeta_n] + [*\beta_s + \beta_t, \xi_n] + *dX_\alpha(*d_\alpha \xi_n),$$

with β_s, β_t the ds, dt-components of the curvature of A. The form ω_n can be bounded as follows:

$$\begin{aligned}
\|\omega_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} &\leq C_3 \left(\|d_\alpha \zeta_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} + \|d_\alpha \xi_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} \right) \\
&\leq 2C_3 \|\mu_n - \eta_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon_n} \\
&\leq \epsilon_n C_4 \left(\|d_\alpha \mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon_n} + \|d_\alpha^{*\epsilon_n} \mu_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon_n} \right) \\
&\leq \epsilon_n C_5 \|V_n\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon_n},
\end{aligned}$$

where we used Lemma 3.4 in the penultimate line. Then we have

$$\begin{aligned}
\|\mathcal{D}_0^* \eta_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} &\leq \|\mathcal{D}^*_{\epsilon_n} V_n\|_{\epsilon_n} + \epsilon_n C_5 \|V_n\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon_n} \\
&\leq C_6 \left(\|\mathcal{D}^*_{\epsilon_n} V_n\|_{\epsilon_n} + \epsilon_n \|\mathrm{proj}_\alpha V_n\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet)} \right).
\end{aligned}$$

This is going to zero in n, and so the statement of Claim 1 follows from Fatou's lemma

$$\|\mathcal{D}_0^* \eta_\infty\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} \leq \liminf_n \|\mathcal{D}_0^* \eta_n\|_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} = 0.$$

Moving on to Claim 2, write $V_n = v_n + r_n\, ds$. Recall that $\mathrm{proj}_a v_n$ denotes the L^2-orthogonal projection to the space of a-harmonic 1-form representatives on $Y_\bullet$. In particular, for each $s \in \mathbb{R}$, the form $\mathrm{proj}_a v_n(s)|_{\partial Y_\bullet}$ lies in the Lagrangian subbundle, essentially by definition. To prove Claim 2, we will show that for any $s \in \mathbb{R}$ the form $\mathrm{proj}_a v_n(s)|_{\partial Y_\bullet}$ can be made arbitrarily close to $(\eta_\infty(s,0), \eta_\infty(s,1))$ by taking n sufficiently large. Then since the Lagrangian subbundle is closed, we must have that $(\eta_\infty(s,0), \eta_\infty(s,1))$ lies in the Lagrangian subbundle as well. Strictly speaking, it may not be clear what this means since η_∞ is only $W^{1,2}$ and so its pointwise evaluation may not make sense. Nevertheless, this sketch can be made precise as follows.

First note that restriction to the boundary of I extends to a compact operator

$$W^{1,2}(B \times I, H_\alpha) \to L^2(B \times \partial I, H_\alpha|_{B \times \partial I})$$

for any compact $B \subset \mathbb{R}$. By the uniform $W^{1,2}$ bound on the η_n, this implies that, after possibly passing to a further subsequence, we have

$$\|\eta_n - \eta_\infty\|_{L^2(B \times \partial I \times \Sigma_\bullet)} \longrightarrow 0, \tag{30}$$

for each compact $B \subset \mathbb{R}$. Note that since V_n is continuous, the forms μ_n and v_n agree on the common boundary

$$\mu_n|_{\partial(\mathbb{R} \times I \times \Sigma_\bullet)} = v_n|_{\partial(\mathbb{R} \times Y_\bullet)}.$$

However, their respective projections

$$\eta_n = \mathrm{proj}_\alpha \mu_n, \qquad \text{and} \qquad \mathrm{proj}_a v_n$$

generally will not agree on their common boundary, reflecting the fact that V_n generally will not have Lagrangian boundary conditions. Nevertheless, the following subclaim says they almost agree.

Subclaim: For every compact $B \subset \mathbb{R}$ we have

$$\left\| \eta_n|_{B \times \partial I \times \Sigma_\bullet} - \{\mathrm{proj}_a v_n\}|_{B \times \partial Y_\bullet} \right\|_{L^2(B \times \partial I \times \Sigma_\bullet)} \longrightarrow 0. \tag{31}$$

Assuming this for now, the triangle inequality applied to (30) and (31) shows that the $\{\mathrm{proj}_a v_n\}|_{B \times \partial Y_\bullet}$ are converging to $\eta_\infty|_{B \times \partial I}$ in L^2. The Lagrangian subbundle is closed (it has finite rank), and the uniform $W^{1,2}$-estimates above imply that the values of $\mathrm{proj}_a v_n|_{\{s\} \times \partial Y_\bullet}$ are confined to a bounded set in this subbundle. In particular, this forces $\eta_\infty|_{\{s\} \times \partial I}$ to lie in this Lagrangian subbundle for almost every $s \in B$. Since B was arbitrary, this shows that η_∞ takes values in the Lagrangian subbundle almost everywhere on the boundary of $\mathbb{R} \times I$. By Lemmas B.4.6 and B.4.9 in [18], elliptic regularity holds for holomorphic curves with almost everywhere Lagrangian boundary conditions and so η_∞ has Lagrangian boundary conditions in the strong sense, as desired.

By the previous paragraph, to prove Claim 2, it suffices to verify that (31) holds for every compact $B \subset \mathbb{R}$. Fix B and let $s \in B$. Use the Hodge decomposition (5) on $Y_\bullet$ to write

$$v_n|_{\{s\} \times Y_\bullet} = \mathrm{proj}_a v_n + d_a x_n + d_a^* y_n,$$

where x_n is a 0-form on $Y_\bullet$ and y_n is a 2-form on $Y_\bullet$ with the property that

$$(*^Y y_n)|_{\{s\}\times\partial Y_\bullet} = 0.$$

Then

$$\begin{aligned}\eta_n|_{\{s\}\times\partial I\times\Sigma_\bullet} - \{\mathrm{proj}_a v_n\}|_{\{s\}\times\partial Y_\bullet} &= \mathrm{proj}_\alpha\left(v_n - \{\mathrm{proj}_a v_n\}|_{\{s\}\times\partial Y_\bullet}\right)\\ &= \mathrm{proj}_\alpha\left(d_a^* y_n|_{\{s\}\times\partial Y_\bullet}\right).\end{aligned}$$

Let η be any element of the α-harmonic space over $\Sigma_\bullet$, and set

$$\|\eta\|_{\mathrm{ext}} := \inf_h \|h\|_{W^{1,2}(Y_\bullet)}$$

where the infimum is over all smooth 1-forms h on $Y_\bullet$ that extend η in the sense that

$$\mathrm{proj}_\alpha(h|_{\partial Y_\bullet}) = \eta.$$

This defines a norm on the α-harmonic space. Since this harmonic space is finite-dimensional, all norms are equivalent, and so there is a bound of the form

$$\|\eta\|_{L^2(\partial I\times\Sigma_\bullet)} \leq C_7\|\eta\|_{\mathrm{ext}}.$$

Apply this with

$$\eta = \mathrm{proj}_\alpha\left(d_a^* y_n|_{\{s\}\times\partial Y_\bullet}\right), \qquad h = d_a^* y_n|_{\{s\}\times\partial Y_\bullet}$$

to get

$$\begin{aligned}\left\|\eta_n|_{\{s\}\times\partial I\times\Sigma_\bullet} - \{\mathrm{proj}_a v_n\}|_{\{s\}\times\partial Y_\bullet}\right\|_{L^2(\{s\}\times\partial I\times\Sigma_\bullet)} &\leq C_7\|d_a^* y_n\|_{W^{1,2}(\{s\}\times\partial Y_\bullet)}\\ &\leq C_8\|d_a d_a^* y_n\|_{L^2(\{s\}\times Y_\bullet)}\\ &= C_8\|d_a v_n\|_{L^2(\{s\}\times Y_\bullet)},\end{aligned}$$

where we used (16) in the second line (with $v = d_a^* y_n$ and $\epsilon = 1$; note that $d_a^* y_n$ is d_a^*-closed, has no harmonic part, and $*d_a^* y_n = d_a * y_n$ restricts to zero on the boundary). This estimate holds for a fixed s, and it is not hard to see that all constants can be taken to be independent of s (they depend only on the connection A). Integrate the above estimate over the compact set $B \subset \mathbb{R}$ and convert to the ϵ_n-dependent metric

$$\begin{aligned}\|\eta_n|_\partial - \{\mathrm{proj}_a v_n\}|_\partial\|_{L^2(B\times\partial I\times\Sigma_\bullet)} &\leq C_8\|d_a v_n\|_{L^2(B\times Y_\bullet)}\\ &= \epsilon_n^{1/2} C_8\|d_a v_n\|_{L^2(B\times Y_\bullet),\epsilon_n}\\ &\leq \epsilon_n^{1/2} C_9\Big(\|\mathcal{D}_{\epsilon_n}^* V_n\|_{\epsilon_n}\\ &\qquad + \|\mathrm{proj}_\alpha V_n\|_{L^2([-S_0,S_0]\times I\times\Sigma_\bullet)}\Big).\end{aligned}$$

We used Corollary 3.12 in the last line. This is going to zero in n, which is exactly (31). $\square$

3.3 Proof of the index relation

Here we prove Theorem 3.1. Let H be a perturbation as in (13), and assume that $a^{\pm}$ are non-degenerate H-flat connections. Suppose A is any connection on $\mathbb{R} \times Q$ that represents a holomorphic strip. For simplicity, we will assume A is constantly equal to the $a^{\pm}$ outside of a compact set in the sense that

$$A|_{[S_0,\infty)\times Y} = a^+, \qquad A|_{(-\infty,-S_0]\times Y)} = a^-$$

for some S_0 sufficiently large. Hence

$$c_A^{[S_0,\infty)} = c_A^{(-\infty,-S_0]} = 0$$

where c_A^U is the constant from Remark 3.10. By applying a suitable gauge transformation, we may assume A is in $W^{1,2}$ relative to the ϵ-dependent smooth structure for all $\epsilon > 0$. Then the operators

$$\mathcal{D}_\epsilon := \mathcal{D}_{\epsilon,A}, \qquad \mathcal{D}_0 := \mathcal{D}_{0,A}$$

are Fredholm and have well-defined Fredholm indices. Assume that H, A have been chosen so that the operator $\mathcal{D}_0$ is surjective when restricted to the sections with Lagrangian boundary conditions. This can always be arranged by making a small perturbation of H. It follows from Theorem 3.13 that the operator $\mathcal{D}_\epsilon$ is onto as well, provided ϵ is sufficiently small. Hence

$$\text{Ind}\,\mathcal{D}_\epsilon = \dim\,\ker \mathcal{D}_\epsilon, \qquad \text{Ind}\,\mathcal{D}_0 = \dim\,\ker \mathcal{D}_0|,$$

where $\mathcal{D}_0|$ denotes the restriction of $\mathcal{D}_0$ to the sections with Lagrangian boundary conditions. To prove the theorem, we will construct an isomorphism between these kernels.

Write

$$\|\cdot\|_\epsilon := \|\cdot\|_{L^2(\mathbb{R}\times Y),\epsilon},$$

and recall the norm $\|\cdot\|_{W^{1,2},\epsilon} := \|\cdot\|_{W^{1,2}(\mathbb{R}\times Y),\epsilon}$ defined in Section 3.2. We denote by

$$L_\epsilon, \qquad \text{and} \qquad W_\epsilon$$

the completion of the space of compactly supported ϵ-smooth elements of $\Omega^1(\mathbb{R} \times Y, \mathbb{R} \times Q(\mathfrak{g}))$ with respect to $\|\cdot\|_\epsilon$, and $\|\cdot\|_{W^{1,2},\epsilon}$, respectively. Then $\mathcal{D}_\epsilon$ extends to a map of the form $\mathcal{D}_\epsilon : W_\epsilon \longrightarrow L_\epsilon$. By (22), it follows that $\mathcal{D}_\epsilon$ restricts to a Banach space isomorphism of the form

$$\mathcal{D}_\epsilon| : \text{im}\,\mathcal{D}_\epsilon^* \xrightarrow{\cong} L_\epsilon,$$

and we denote by

$$Q_\epsilon := (\mathcal{D}_\epsilon|)^{-1}$$

the inverse of this restriction; Q_ϵ is therefore a right inverse for $\mathcal{D}_\epsilon$. Then (22) implies that there is some constant $C_Q > 0$ with

$$\|Q_\epsilon V\|_{W^{1,2},\epsilon} \leq C_Q \|V\|_\epsilon \tag{32}$$

for all $0 < \epsilon < \epsilon_0$ and all $V \in L_\epsilon$. It follows that the linear map

$$W_\epsilon \longrightarrow W_\epsilon, \qquad V \longmapsto (I - Q_\epsilon \mathcal{D}_\epsilon)V, \tag{33}$$

is bounded and has image $\ker \mathcal{D}_\epsilon$. Identify $\ker \mathcal{D}_0|$ with its image in W_ϵ under the embedding (11). Then restricting the map (33) to $\ker \mathcal{D}_0|$ gives a map

$$\mathcal{F}_\epsilon : \ker \mathcal{D}_0| \longrightarrow \ker \mathcal{D}_\epsilon.$$

Our goal is to show $\mathcal{F}_\epsilon$ is an isomorphism.

As a preliminary step, in Lemma 3.14 below we make precise the statement that elements of $\ker \mathcal{D}_0|$ are close to elements of $\ker \mathcal{D}_\epsilon$ when ϵ is small. To state the lemma, we note that the $W^{1,2}$-norm on sections η of $H_\alpha \to \mathbb{R} \times I$ pushes forward under (11) to the norm

$$\begin{aligned} \|V_\eta\|^2_{W^{1,2},0} \quad := \quad & \|\mathrm{proj}_\alpha V_\eta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} \\ & + \|\mathrm{proj}_\alpha \nabla_s V_\eta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} + \|\mathrm{proj}_\alpha \nabla_t V_\eta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)}. \end{aligned}$$

on the space of strip representatives V_η. Note also that this definition of $\|V_\eta\|_{W^{1,2},0}$ makes sense with V_η replaced by any 1-form V on $\mathbb{R} \times Y$; however $\|\cdot\|_{W^{1,2},0}$ is no longer a norm on this larger space of all 1-forms (e.g., there is no control over $V|_{\mathbb{R}\times Y_\bullet}$). With this understood, there is an obvious bound of the form

$$\|V\|_{W^{1,2},0} \leq \|V\|_{W^{1,2},\epsilon}$$

that holds for all ϵ-smooth 1-forms V.

Lemma 3.14. *There is a constant $C > 0$ such that*

$$\|\mathcal{D}_\epsilon V_\eta\|_\epsilon \leq \epsilon^{1/2} C \|V_\eta\|_{W^{1,2},0}$$

for all $0 < \epsilon \leq 1$ and all $V_\eta \in \ker \mathcal{D}_0|$ with Lagrangian boundary conditions.

Proof. Write $V_\eta = v + r\,ds$ and $v|_{\mathbb{R}\times I\times\Sigma_\bullet} = \eta + \rho\,ds + \theta\,dt$. The Lagrangian boundary conditions and $\mathcal{D}_0\eta = 0$ imply

$$\mathcal{D}_\epsilon V_\eta|_{\mathbb{R}\times Y_\bullet} = \begin{pmatrix} \nabla_s v - d_a r \\ -\nabla_s r \end{pmatrix}, \qquad \mathcal{D}_\epsilon V_\eta|_{\mathbb{R}\times I\times\Sigma_\bullet} = \begin{pmatrix} 0 \\ \nabla_s \rho - \nabla_t \theta \\ \nabla_t \rho + \nabla_s \theta \end{pmatrix};$$

this uses the coordinate description of $\mathcal{D}_\epsilon$ and the detailed construction of the embedding (11). Taking the ϵ-dependent L^2-norm and then converting to the standard

L^2-norm gives

$$\begin{aligned}\|\mathcal{D}_\epsilon V_\eta\|_\epsilon^2 &= \|\nabla_s v - d_a r\|^2_{L^2(\mathbb{R}\times Y_\bullet),\epsilon} + \|\nabla_s r\|^2_{L^2(\mathbb{R}\times Y_\bullet),\epsilon} \\ &\quad + \|\nabla_s \rho - \nabla_t \theta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon} + \|\nabla_t \rho + \nabla_s \theta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet),\epsilon} \\ &\leq \epsilon\Big\{\|\nabla_s v - d_a r\|^2_{L^2(\mathbb{R}\times Y_\bullet)} + \|\nabla_s r\|^2_{L^2(\mathbb{R}\times Y_\bullet)} \\ &\qquad + \|\nabla_s \rho - \nabla_t \theta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)} + \|\nabla_t \rho + \nabla_s \theta\|^2_{L^2(\mathbb{R}\times I\times\Sigma_\bullet)}\Big\} \\ &\leq \epsilon\left(\|d_A V_\eta\|^2_{L^2(\mathbb{R}\times Y)} + \|d_A^* V_\eta\|^2_{L^2(\mathbb{R}\times Y)}\right),\end{aligned}$$

where the first inequality holds because $\epsilon^2 \leq \epsilon \leq 1$ and the last follows by writing d_A and d_A^* in components. Since the space $\ker \mathcal{D}_0|$ is finite-dimensional, any two norms are equivalent. In particular, there is some constant $C > 0$ such that

$$\|d_A V_\eta\|^2_{L^2(\mathbb{R}\times Y)} + \|d_A^* V_\eta\|^2_{L^2(\mathbb{R}\times Y)} \leq C\|V_\eta\|^2_{W^{1,2},0} \qquad \forall\, V_\eta \in \ker \mathcal{D}_0|.$$

(Strictly speaking, the object on the left may only be a semi-norm, but this does not affect the argument.) Combining this with the above estimate of $\|\mathcal{D}_\epsilon V_\eta\|_\epsilon^2$ finishes the proof. □

Now we prove that $\mathcal{F}_\epsilon$ is injective. Suppose $\mathcal{F}_\epsilon(V_\eta) = 0$ for some $V_\eta \in \ker \mathcal{D}_0|$. Then it follows that $V_\eta = Q_\epsilon \mathcal{D}_\epsilon V_\eta$, and so

$$\|V_\eta\|_{W^{1,2},0} \leq \|V_\eta\|_{W^{1,2},\epsilon} = \|Q_\epsilon \mathcal{D}_\epsilon V_\eta\|_{W^{1,2},\epsilon} \leq C_Q\|\mathcal{D}_\epsilon V_\eta\|_\epsilon \leq \epsilon^{1/2} C_Q C \|V_\eta\|_{W^{1,2},0};$$

the last two inequalities follow by (32) and Lemma 3.14, respectively. Since $\|\cdot\|_{W^{1,2},0}$ is a (non-degenerate) norm on $\ker \mathcal{D}_0|$, this implies $V_0 = 0$ when ϵ is small.

For surjectivity, suppose $\mathcal{F}_\epsilon$ is not onto regardless of how small we take ϵ. Then for each n, we can find a positive number ϵ_n, and a section $V_n \in W_{\epsilon_n}$ such that

$$\epsilon_n \to 0, \qquad \mathcal{D}_{\epsilon_n} V_n = 0, \qquad \|V_n\|_{W^{1,2},\epsilon_n} = 1.$$

We may assume V_n is $W^{1,2}_{\epsilon_n}$-orthogonal to the image of $\mathcal{F}_{\epsilon_n}$. Write

$$V_n = \mu_n + \rho_n\, ds + \theta_n\, dt, \qquad \text{on } \mathbb{R}\times I\times\Sigma.$$

Then by the same type of compactness argument that appeared in Section 3.2.3, it follows that there is a subsequence of the μ_n (still denoted μ_n) that converges to some $\eta_\infty \in \ker \mathcal{D}_0|$, where the convergence is in L^2 on compact subsets of $\mathbb{R}\times I$. Then since V_n is orthogonal to the image of $\mathcal{F}_{\epsilon_n}$, we have

$$1 = \|V_n\|_{W^{1,2},\epsilon_n} \quad \leq \quad \|V_n - \mathcal{F}_{\epsilon_n} V_{\eta_\infty}\|_{W^{1,2},\epsilon_n}.$$

Now apply Corollary 3.12 to $V = V_n - \mathcal{F}_{\epsilon_n} V_{\eta_\infty} \in \ker \mathcal{D}_{\epsilon_n}$ to continue this and get

$$
\begin{aligned}
1 &\leq \|V_n - \mathcal{F}_{\epsilon_n} V_{\eta_\infty}\|_{W^{1,2},\epsilon_n} \\
&\leq C\|\mathrm{proj}_\alpha (V_n - \mathcal{F}_{\epsilon_n} V_{\eta_\infty})\|_{L^2([-S_0,S_0]\times I\times\Sigma)} \\
&\leq C\Big(\|\mu_n - \eta_\infty\|_{L^2([-S_0,S_0]\times I\times\Sigma)} + \|Q_{\epsilon_n}\mathcal{D}_{\epsilon_n} V_{\eta_\infty}\|_{L^2([-S_0,S_0]\times I\times\Sigma)}\Big) \\
&\leq C\Big(\|\mu_n - \eta_\infty\|_{L^2([-S_0,S_0]\times I\times\Sigma)} + C_Q\|\mathcal{D}_{\epsilon_n} V_{\eta_\infty}\|_{L^2([-S_0,S_0]\times I\times\Sigma)}\Big).
\end{aligned}
$$

By Lemma 3.14 and the L^2-convergence of the μ_n, the right-hand side is going to zero in n, which is a contradiction.

References

[1] M. Atiyah. New invariants of three and four dimensional manifolds. *The Math. Heritage of Hermann Weyl, Proc. Sympos. Pure Math.* 48, 285-299, 1988.

[2] M. Atiyah, R. Bott. The Yang-Mills equations over Riemann surfaces. *Phil. Trans. Roy. Soc. London Ser. A*, 308:523-615, 1982.

[3] P.J Braam, S.K. Donaldson. Floer's work on instanton homology, knots and surgery. *The Floer Memorial Volume*, Ed. Hofer et al. Birkhäuser 195-256 1995.

[4] S. Donaldson. *Floer homology groups in Yang-Mills theory*. Cambridge University Press, Cambridge, 2002. With the assistance of M. Furuta and D. Kotschick.

[5] S. Dostoglou, D. Salamon. Cauchy-Riemann operators, self-duality and the spectral flow. *First European Congress of Mathematics, Vol.* I (Paris, 1992), 511-545, Progr. Math., 119, Birkhäuser, Basel, 1994.

[6] S. Dostoglou, D. Salamon. Instanton homology and symplectic fixed points. *Symplectic geometry*, 57-93, *London Math. Soc. Lecture Note Ser.*, 192, Cambridge Univ. Press, Cambridge, 1993.

[7] S. Dostoglou, D. Salamon. Self-dual instantons and holomorphic curves. *Ann. of Math.* (2) 139 (1994), no. 3, 581-640.

[8] D. Duncan. Higher-rank instanton cohomology and the quilted Atiyah-Floer conjecture. arXiv:1311.5609.

[9] D. Duncan. Compactness results for neck-stretching limits of instantons. arXiv:1212.1547v3.

[10] A. Floer. A relative Morse index for the symplectic action. *Comm. on Pure and Applied Math., Vol.* XLI 393-407, 2988.

[11] A. Floer. An instanton-invariant for 3-manifolds. *Comm. Math. Phys.*, 118(2):215-240, 1988.

[12] A. Floer. Instanton homology and Dehn surgery. *Floer Memorial Volume, Progress in Mathematics, Vol.* 133, Birkhäuser Verlag, 1995. ISBN 3-7643-5044-X, Basel; ISBN 0-8176-5044-8, Boston.

[13] A. Floer. Morse theory for Lagrangian intersections. *J. Differential Geom.* 28 (1988), no. 3, 513-547.

[14] A. Floer, H. Hofer, D. Salamon. Transversality in elliptic Morse theory for the symplectic action. Duke Mathematical Journal (C) October 1995 80, No. 1.

[15] D. Gay, R. Kirby. Indefinite Morse 2-functions: broken fibrations and generalizations, 2011. arXiv:1102.0750.

[16] D. Gay, K. Wehrheim, C. Woodward. Connected Cerf theory. In preparation.

[17] P. Kronheimer, T. Mrowka. *Monopoles and Three-Manifolds*. Cambridge University Press, New York, 2007.

[18] D. McDuff, D. Salamon. *J-holomorphic curves and symplectic topology*. American Mathematical Society Colloquium Publications, Vol. 52. American Mathematical Society, Providence, RI, 2004.

[19] E. Meinrenken, C. Woodward. *Hamiltonian loop group actions and Verlinde factorization*.

[20] J. Milnor. *Lectures on the h-Cobordism Theorem. Notes by L. Siebenmann and J. Sondow.* Princeton University Press, Princeton, N.J., 1965.

[21] Y.-G. Oh. Floer cohomology of Lagrangian intersections and pseudo-holomorphic disks I. *Comm. Pure Appl. Math.*, 46(7):949-993, 1993.

[22] Y.-G. Oh. On the structure of pseudo-holomorphic discs with totally real boundary conditions. *The Journal of Geometric Analysis*, June 1997, Volume 7, Issue 2, pp 305-327.

[23] D. Salamon, K. Wehrheim. Instanton Floer homology with Lagrangian boundary conditions. arXiv:math/0607318.

[24] F. Warner. *Foundations of differentiable manifolds and Lie groups*. Scott-Foreman, Glenview, Illinois, 1971.

[25] K. Wehrheim. *Uhlenbeck compactness*. EMS Series of Lectures in Mathematics, 2004.

[26] K. Wehrheim, C. Woodward. Floer field theory for coprime rank and degree. Preprint.

[27] K. Wehrheim, C. Woodward. Pseudoholomorphic quilts. *J. Symp. Geom.* To appear.

[28] K. Wehrheim, C. Woodward. Quilted Floer cohomology. *Geometry and Topology*, 14:833-902, 2010. arXiv:0905.1370.

[29] K. Wehrheim, C. Woodward. Erratum: Quilted Floer cohomology. *Geometry and Topology* 16 (2012) 127-154.

[30] L. M. Woodward. The classification of principal PU_n-bundles over a 4-complex. *J. London Math. Soc.* (2) 25 (1982), no. 3, 513-524.