

Higher-rank instanton cohomology and 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. The quilted Atiyah-Floer conjecture states that these cohomology groups are isomorphic. We initiate a program for proving this conjecture.

Contents

1	Introduction	1
2	Background on the quilted Atiyah-Floer conjecture	2
2.1	ϵ -dependent smooth structures	6
2.2	Gauge theory	8
2.3	Topological aspects of principal $PU(r)$ -bundles	14
2.4	Compatible bundles for broken circle fibrations	17
2.5	Quilted Floer theory	18
3	Higher rank instanton Floer theory	23
3.1	Instanton Floer cohomology for $G = PU(r)$	24
3.2	Gradings	33
3.3	Instanton Floer cohomology for broken circle fibrations	37
4	The quilted Atiyah-Floer conjecture	39
4.1	A program for proving the conjecture	40
4.2	Compatible perturbations	44

1 Introduction

This paper is the first in a series [11, 13] that prove various aspects of the quilted Atiyah-Floer conjecture. The present paper is dedicated entirely to the underlying Floer theory and differential geometry of the project.

Section 2 begins with a brief history and overview of the conjecture. We also introduce our notation and conventions. The section ends with a description of the quilted Floer group associated to a 3-manifold. This group was first introduced by Wehrheim and Woodward in [29], and is defined via a 2+1 field theoretic scheme that relates symplectic and gauge theoretic 3-manifold invariants. We point out that in [29] the authors define their quilted invariants for 3-manifolds via gauge theory on suitable $PU(r)$ -bundles (one traditionally considers the case $PU(2) = SO(3)$).

In Section 3 we discuss Floer’s instanton homology [14, 15]. We note that, as in the symplectic theory, the standard references (e.g., Donaldson’s book [5], as well as Floer’s original papers) focus almost entirely on the Lie groups $SU(2)$ and $SO(3)$. However, it is possible to define instanton homology for other Lie groups as well. To maintain the analogy with the Woodward-Wehrheim theory, we focus attention on $PU(r)$ for $r \geq 2$. A full treatment of higher-rank instanton Floer theory does not seem to appear in the literature, so we have spent some time developing the theory in this case, proving several folklore results that are well-known in the case $r = 2$. We note also that Kronheimer has defined higher-rank Donaldson invariants [23]. These can be viewed as the 4-dimensional version of the higher-rank instanton Floer theory discussed here.

In Section 4 we make a precise statement of the quilted Atiyah-Floer conjecture, as well as a chain level version of the conjecture. We then describe our overall approach, and prove several chain level statements. This includes a detailed discussion of how to obtain transversality *simultaneously* in both Floer theories.

Acknowledgments: The author is grateful to his thesis advisor Chris Woodward for his insight and valuable suggestions. In addition, the author benefited greatly from discussions with Katrin Wehrheim, via Chris Woodward, of her unpublished work on the Atiyah-Floer conjecture. This unpublished work analyzes various bubbles that appear in the limit of instantons with Lagrangian boundary conditions with degenerating metrics, and it outlines the remaining problems in such an approach. The present paper takes a different approach that avoids instantons with Lagrangian boundary conditions.

The author would also like to thank Penny Smith, Dietmar Salamon, Fabian Ziltener, Peter Ozsváth, and Tom Parker for enlightening discussions. Finally, thanks are certainly due to the author’s wife Claudette D’Souza for her help with early versions of the paper, and for her support along the way. This work was partially supported by NSF grants DMS 0904358 and DMS 1207194.

2 Background on the quilted Atiyah-Floer conjecture

Let Y be a closed, connected, oriented 3-manifold, and fix an $SO(3)$ -bundle $Q \rightarrow Y$ such that $w_2(Q)$ pairs non-trivially with an element of $H_2(Y, \mathbb{Z})$ (so Y

is necessarily *not* a homology 3-sphere). In [15, 14], Floer showed that this data defines a chain complex $(CF_{\text{inst}}(Q), \partial_{\text{inst}})$ where $CF_{\text{inst}}(Q)$ is the abelian group freely generated by the gauge equivalence classes of flat connections on Q . (Strictly speaking, one needs to suitably perturb the defining equations in order to obtain $\partial_{\text{inst}}^2 = 0$.) Then the associated homology group $HF_{\text{inst}}(Q)$, called *instanton Floer cohomology*, is well-defined and depends only on the topological type of $Q \rightarrow Y$.

In [7, 8, 9, 10], Dostoglou and Salamon considered the following special case. Let Σ be a closed, connected, oriented surface, $P \rightarrow \Sigma$ a non-trivial $SO(3)$ -bundle, and $\Phi : P \rightarrow P$ a bundle map covering an orientation-preserving diffeomorphism $\varphi : \Sigma \rightarrow \Sigma$. Then Dostoglou-Salamon applied Floer's construction to the mapping tori $Y = (I \times \Sigma)/\varphi$ and $Q = (I \times P)/\Phi$. This is particularly interesting due to the following symplectic interpretation: Let $M(\Sigma)$ denote the space of flat connections on P modulo the action of the identity component of the gauge group. This is naturally a smooth finite-dimensional symplectic manifold, and pullback by Φ induces a symplectomorphism $\Phi^* : M(\Sigma) \rightarrow M(\Sigma)$. Then the graph $\text{Gr}(\Phi^*) \subset M(\Sigma)^- \times M(\Sigma)$ is a Lagrangian submanifold, where the superscript in $M(\Sigma)^-$ means we have replaced the symplectic form by its negative. Similarly, the diagonal $\Delta \subset M(\Sigma)^- \times M(\Sigma)$ is also Lagrangian, and one can use this data to define a *Lagrangian Floer group* $HF_{\text{symp}}(\text{Gr}(\Phi^*), \Delta)$. This can be thought of as the homology of a chain complex $(CF_{\text{symp}}(\text{Gr}(\Phi^*), \Delta), \partial_{\text{symp}})$ generated by the intersection points $\text{Gr}(\Phi^*) \cap \Delta$. On the other hand, the elements of $\text{Gr}(\Phi^*) \cap \Delta$ can be naturally identified with the gauge equivalence classes of flat connections on Q . That is, $HF_{\text{symp}}(\text{Gr}(\Phi^*), \Delta)$ and $HF_{\text{inst}}(Q)$ are both well-defined homology groups that measure the flat connections on Q . Dostoglou and Salamon showed that these Floer groups are naturally isomorphic

$$HF_{\text{inst}}(Q) \cong HF_{\text{symp}}(\text{Gr}(\Phi^*), \Delta).$$

They prove this by demonstrating a chain-level isomorphism in the spirit of Atiyah's idea to 'stretch the neck' [1].

The problem we consider in the present paper is a generalization of the Dostoglou-Salamon set-up to 3-manifolds Y with $b_1(Y) > 0$. Our aim now is to describe the relevant symplectic Floer group in this more general setting. More explicitly, we assume Y is a closed, connected, oriented 3-manifold equipped with a Morse function $f : Y \rightarrow S^1$ that is not homotopically trivial. We assume further that f has connected fibers. Such functions were considered by Lekili [24]; following him we call these *broken circle fibrations*. It follows from work of Gay and Kirby [19] that Y admits a broken circle fibration if and only if Y has positive first Betti number. Moreover, each critical point of f has index 1 or 2 (the index is defined with respect to the canonical orientation on the circle). We will see that the invariants we are interested in only depend on the homotopy class of f . In particular, by homotoping f , we can assume the critical points of f have distinct critical values, and we let N denote the number of critical points. If $N = 0$, then by performing a suitable homotopy we can replace f by

a Morse function with two critical points; this is the Morse-theoretic version of destabilization. We may therefore assume $N > 0$.

Following [29] and [20], we will use f to decompose Y into a union of elementary cobordisms (a so-called *Cerf decomposition*): Identify $S^1 \cong \mathbb{R}/C\mathbb{Z}$ for some $C > 0$. Find regular values $r_i \in S^1$ such that, for each $0 \leq i \leq N-1$, there is exactly one critical value $c_{i(i+1)}$ with $r_i + \delta < c_{i(i+1)} < r_{i+1} - \delta$, for some fixed $\delta > 0$. Here and below we work with i modulo N . We may assume the circumference C is large enough to take $\delta = 1/2$. Define

$$\Sigma_i := f^{-1}(r_i - 1/2), \quad Y_{i(i+1)} := f^{-1}([r_i + 1/2, r_{i+1} - 1/2]),$$

which are compact, connected, oriented surfaces and 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 $r_i - 1/2$ and $r_i + 1/2$, so the normalized gradient vector field $V := \nabla f / |\nabla f|$ is well-defined on $f^{-1}([r_i - 1/2, r_i + 1/2])$. The time-1 gradient flow of V provides an identification

$$f^{-1}([r_i - 1/2, r_i + 1/2]) \cong I \times \Sigma_i, \quad (1)$$

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

$$Y_{01} \cup_{\Sigma_1} (I \times \Sigma_1) \cup_{\Sigma_1} Y_{12} \cup_{\Sigma_2} \dots \cup_{\Sigma_{N-1}} Y_{(N-1)0} \cup_{\Sigma_0} (I \times \Sigma_0) \cup_{\Sigma_0} \quad (2)$$

Note that this 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, reflecting the fact that f maps to the circle. By construction, each $Y_{i(i+1)}$ is an *elementary cobordism*; that is, it admits an I -valued Morse function that (i) preserves the cobordism structure between $Y_{i(i+1)}$ and I , and (ii) has at most one critical point (necessarily in the interior of $Y_{i(i+1)}$). Since each critical point has index either 1 or 2, it follows from standard Morse theory considerations [25, Theorem 3.14] that the genus of Σ_i differs from that of Σ_{i+1} by one. In particular, the number of critical points N is *even*.

It is possible to choose an $\mathrm{SO}(3)$ -bundle $Q \rightarrow Y$ so its restriction to each Σ_i is non-trivial. Then $HF_{\mathrm{inst}}(Q)$ is well-defined. On the symplectic side, define $M(\Sigma_i)$ for each Σ_i as above. This choice of Q ensures that each $M(\Sigma_i)$ is smooth. Restricting to each of the two boundary components of the $Y_{i(i+1)}$ defines two smooth Lagrangian submanifolds $L^{(0)}, L^{(1)} \subset M(\Sigma_0) \times \dots \times M(\Sigma_{N-1})$ for which $HF_{\mathrm{symp}}(L^{(0)}, L^{(1)})$ is well-defined; see Section 2.5.2. Then the *quilted Atiyah-Floer conjecture* states that there is a natural isomorphism $HF_{\mathrm{inst}}(Q) \cong HF_{\mathrm{symp}}(L^{(0)}, L^{(1)})$.

It follows from the definitions that there is a natural isomorphism of abelian groups

$$CF_{\mathrm{inst}}(Q) \cong CF_{\mathrm{symp}}(L^{(0)}, L^{(1)}), \quad (3)$$

where these are (suitably defined) chain groups inducing the Floer homology groups $HF_{\mathrm{inst}}(Q)$, $HF_{\mathrm{symp}}(L^{(0)}, L^{(1)})$; see Theorem 4.2 below. Moreover, these chain groups each admit natural relative gradings. In [11] it is shown that the isomorphism (3) preserves these gradings. The quilted Atiyah-Floer conjecture would follow if one could show that this isomorphism intertwines the boundary operators. One way to show this is to prove that there is a one-to-one correspondence between instantons and holomorphic curves, at least for suitably chosen data (metric, almost complex structure and perturbations). The main result of [13] implies that, when this data is chosen well, every instanton trajectory counted by the boundary operator ∂_{inst} is close to some holomorphic curve trajectory counted by the boundary operator ∂_{symp} . (Proving the converse statement, and hence the conjecture, is still a work in progress.) In the present paper we describe how the aforementioned data can be chosen so that both chain level groups are comparable.

Remark 2.1. (a) Fukaya [17, 18] describes a program similar to the one presented here. Though quite similar, there is one striking difference between his approach and ours: Fukaya deals with a fixed smooth (but not definite) metric, whereas we deal with a sequence of metrics that degenerate to a singular metric, and these metrics are not smooth.

(b) Recently M. Lipyanskiy has developed compactness results for quilts with patches consisting both of instantons and holomorphic curves. The motivation is that these quilts may define a chain map that interpolates between the trajectories defining each of the two Floer theories.

(c) The quilted Atiyah-Floer conjecture is a variant of the older Atiyah-Floer conjecture; see [1]. In this latter statement, one considers a homology 3-sphere Y , and works with $\mathrm{SU}(2)$, rather than $\mathrm{SO}(3)$. Then the Atiyah-Floer conjecture posits an isomorphism between an instanton Floer homology and a symplectic Floer homology associated to this data. Unfortunately, at the time of writing, the relevant symplectic Floer homology is not well-defined due to issues arising from reducible connections; however, see [26]. The homology theories relevant for the quilted version are well-posed essentially because the bundles are chosen so that one can avoid the use of reducible connections.

(d) In [1], Atiyah outlined a metric degeneration approach to the Atiyah-Floer conjecture. The degenerating metric g_ϵ we define in the next section is, up to conformal equivalence, a direct analogue of Atiyah's metric for the quilted Atiyah-Floer conjecture.

2.1 ϵ -dependent smooth structures

As a matter of notational convenience, we set

$$\Sigma_\bullet := \bigsqcup_i \Sigma_i, \quad Y_\bullet := Y \setminus (I \times \Sigma_\bullet).$$

We refer to the connected components of the boundary $\partial Y_\bullet = \{0, 1\} \times \Sigma_\bullet$ as the *seams*, and we use the letter t to denote the coordinate variable on the interval I . In particular, $dt \in \Omega^1(I \times \Sigma)$ is identified with

$$df / |df| \in \Omega^1 \left(\cup_i f^{-1}([r_i - 1/2, r_i + 1/2]) \right)$$

under the identification (1).

Over $I \times \Sigma_\bullet$ the metric g has the form

$$dt^2 + g_\Sigma$$

where g_Σ is a path of metrics on Σ_\bullet . To simplify the exposition, we assume that g has been chosen so that g_Σ is a *constant* path, which can always be achieved. For $\epsilon > 0$ define a new metric g_ϵ by

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

We will be interested in taking the limit as ϵ approaches 0. See Figure 2.

Let \mathcal{S}_1 denote the smooth structure on Y ; that is, the smooth structure in which g and f are smooth. We call this the *standard smooth structure*. It is important to note that when $\epsilon \neq 1$, the metric g_ϵ is *not* smooth with the standard smooth structure. For example, take $V = \nabla f / |\nabla f|$, where the norm and gradient are taken with respect to $g = g_1$. Then away from the critical points of f , V is smooth on (Y, \mathcal{S}_1) , but

$$g_\epsilon(V, V) = \begin{cases} 1 & \text{on } I \times \Sigma_\bullet \\ \epsilon^2 & \text{on } Y_\bullet \setminus \{\text{critical points}\} \end{cases} \quad (4)$$

is not even continuous, so g_ϵ cannot be continuous on (Y, \mathcal{S}_1) .

However, there is a *different* smooth structure \mathcal{S}_ϵ in which g_ϵ is smooth. Moreover, $(Y, \mathcal{S}_\epsilon)$ is *diffeomorphic* to (Y, \mathcal{S}_1) . This can be seen as follows: View Y as the *topological* manifold in (2). Following Milnor [25, Theorem 1.4], any choice of collar neighborhoods of the seams determines a smooth structure on Y , and any two choices are given by isotopic data. These isotopies determine a




Figure 2: Above are four illustrations of the manifold Y ; in each, Y is viewed from the top. Moving from left to right, the metric on Y is being deformed in such a way that the volume of the Σ_i and the $Y_{i(i+1)}$ are going to zero. However, the volume in the I -direction (the ‘neck’) is remaining fixed. In the picture on the far right, the $Y_{i(i+1)}$ have collapsed entirely to the critical points of f , which are represented by dots.




Figure 3: The vector field V is not continuous with respect to the topology on TY defined using the standard smooth structure. However, it is continuous with respect to the topology on TY defined using the ϵ -dependent smooth structure.

diffeomorphism between the smooth structures. In this language, the smooth structure \mathcal{S}_1 arises by choosing collar neighborhoods of the seams determined by the time- δ gradient flow of f , and then using the identity to glue these neighborhoods on the overlap. Here $\delta > 0$ is small, but fixed. On the other hand, the smooth structure \mathcal{S}_ϵ arises by taking the time- δ gradient flow on the $Y_{i(i+1)}$ side of the seam $\{1\} \times \Sigma_i$, but the time- $\delta\epsilon$ gradient flow on the $[0, 1] \times \Sigma_i$ side of the seam, and then gluing using the gluing map $(t, \sigma) \mapsto (\epsilon t, \sigma)$.

By the discussion of the previous paragraph, there is a diffeomorphism

$$F_\epsilon : (Y, \mathcal{S}_1) \longrightarrow (Y, \mathcal{S}_\epsilon). \tag{5}$$

In fact, there is a canonical choice of F_ϵ given by taking the obvious straight line homotopy between the gluing maps described above. When the specific smooth structure on Y is relevant, we will write Y^ϵ for $(Y, \mathcal{S}_\epsilon)$. When a function, section of a bundle, connection, etc. is smooth on Y^ϵ , we will say that it is ϵ -smooth. We will write f^ϵ for the pullback of $f : Y \rightarrow S^1$ under F_ϵ^{-1} .

Observe that g_ϵ only fails to be smooth on Y^1 at the seams $\{0, 1\} \times \Sigma_i$. Fur-

thermore, even at the seams, the metric g_ϵ is smooth in directions *parallel* to the seams. So the discontinuity illustrated in Figure 3 is the only thing that goes wrong.

Remark 2.2. Fix $1 \leq p \leq \infty$ and let $W^{1,p}(Y^1)$ denote the Sobolev space associated to the standard smooth structure on $Y = Y^1$. By passing to local coordinates, it is straightforward to show that every ϵ -smooth function h on Y^ϵ is of Sobolev class $W^{1,p}(Y^1)$. That is, the identity map on Y determines a pullback map of the form $\text{Id}^* : C^\infty(Y^\epsilon) \rightarrow W^{1,p}(Y^1)$. This can be seen as follows: The underlying topologies on Y^1 and Y^ϵ are identical, so the ϵ -smooth function h is continuous on Y^1 . Moreover, on the complement of the seams, h is 1-smooth with bounded derivative. This implies Id^*h is of Sobolev class $W^{1,p}$.

However, in general, an ϵ -smooth tensor will only be in $L^p(Y^1)$ with respect to the standard smooth structure. Indeed, any tensor that is non-zero in directions transverse to the seam will necessarily have a jump discontinuity as in Figure 3. So taking a derivative transverse to the seam will introduce a delta function. This applies to connections as well.

2.2 Gauge theory

Let X be an oriented Riemannian n -manifold, possibly with boundary. Given a fiber bundle $E \rightarrow X$, we denote the space of smooth sections by $\Gamma(E)$. Now suppose E is a vector bundle equipped with a metric $\langle \cdot, \cdot \rangle$. Then we write

$$\Omega^\bullet(X, E) := \bigoplus_k \Omega^k(X, E), \quad \Omega^k(X, E) := \Gamma(\Lambda^k T^*X \otimes E),$$

for the spaces of smooth E -valued forms on X . Given $\mu \in \Omega^j(X, E), \nu \in \Omega^k(X, E)$ we can combine the inner product with the wedge to define a new section $\langle \mu \wedge \nu \rangle \in \Omega^{j+k}(X)$. When X is compact, this can be integrated to obtain a non-degenerate bilinear pairing on forms of dual degree:

$$\Omega^k(X, E) \otimes \Omega^{n-k}(X, E) \longrightarrow \mathbb{R}, \quad \mu \otimes \nu \longmapsto \int_X \langle \mu \wedge \nu \rangle. \quad (6)$$

The metric on X induces a Hodge star $*$: $\Omega^k(M, E) \rightarrow \Omega^{n-k}(M, E)$, and this satisfies $** = (-1)^{k(n-k)}$ on k -forms. Sticking $*$ in the second slot of (6) defines an L^2 -inner product on the vector space $\Omega^k(X, E)$. More generally, for $p \geq 1$ we can define the L^p -norm on $\Omega^k(M, E)$ by

$$\|\mu\|_{L^p}^p := \int_X |\mu|^p \, d\text{vol},$$

where $|\mu| := (*\langle \mu \wedge *\mu \rangle)^{1/2}$. Similarly, fixing a metric connection on E , we can define the Sobolev norms $\|\cdot\|_{W^{k,p}}$ on $\Omega^\bullet(X, E)$ by the usual formula. We denote by $W^{k,p}(E)$ the closure, with respect to $\|\cdot\|_{W^{k,p}}$, of the space of compactly supported smooth E -valued forms on X . When the bundle E is clear from context,

we will write $W^{k,p}$ for $W^{k,p}(E)$. In situations where it is particularly important to emphasize the underlying base manifold, we will write $W^{k,p}(X)$ instead of $W^{k,p}(E)$.

Let G be a compact Lie group with Lie algebra \mathfrak{g} , and let $\pi : P \rightarrow X$ be a principal G -bundle. Given any matrix representation $\rho : G \rightarrow \text{GL}(V)$, we can form the *associated bundle*

$$P(V) := P \times_G V = (P \times V)/G.$$

This is naturally equipped with the structure of vector bundle $P(V) \rightarrow X$. Of particular interest is the case when $V = \mathfrak{g}$ is the Lie algebra of G , and ρ is the adjoint representation. We will assume that \mathfrak{g} is equipped with a choice of Ad-invariant inner product $\langle \cdot, \cdot \rangle$. This is always possible since G is compact, and this is unique (up to multiplication by a scalar) when \mathfrak{g} is simple. The Ad-invariance implies that the inner product descends to a bundle metric $\langle \cdot, \cdot \rangle$ on the vector bundle $P(\mathfrak{g})$.

The Lie bracket $[\cdot, \cdot] : \mathfrak{g} \otimes \mathfrak{g} \rightarrow \mathfrak{g}$ is Ad-invariant, and so combines with the wedge to endow $\Omega^\bullet(X, P(\mathfrak{g}))$ and $\Omega^\bullet(P, \mathfrak{g})$ with the structure of a graded algebra; we denote this multiplication by $[\mu \wedge \nu]$ in both cases. Then pullback by $\pi : P \rightarrow X$ induces a graded algebra homomorphism

$$\pi^* : \Omega^\bullet(X, P(\mathfrak{g})) \longrightarrow \Omega^\bullet(P, \mathfrak{g}).$$

Moreover, π^* is injective with image given by the *basic* forms

$$\Omega^\bullet(P, \mathfrak{g})_{\text{basic}} := \left\{ \mu \in \Omega^\bullet(P, \mathfrak{g}) \mid \begin{array}{l} (u_P)^* \mu = \text{Ad}(u^{-1})\mu, \quad \forall u \in G \\ \iota_{\xi_P} \mu = 0, \quad \forall \xi \in \mathfrak{g} \end{array} \right\}.$$

Here u_P (resp. ξ_P) is the image of $u \in G$ (resp. $\xi \in \mathfrak{g}$) under the map $G \rightarrow \text{Diff}(P)$ (resp. $\mathfrak{g} \rightarrow \text{Vect}(P)$) afforded by the group action. We will use π^* to identify the spaces $\Omega^\bullet(X, P(\mathfrak{g}))$ and $\Omega^\bullet(P, \mathfrak{g})_{\text{basic}}$.

Denote by

$$\mathcal{A}(P) = \left\{ A \in \Omega^1(P, \mathfrak{g}) \mid \begin{array}{l} (u_P)^* A = \text{Ad}(u^{-1})A, \quad \forall u \in G \\ \iota_{\xi_P} A = \xi, \quad \forall \xi \in \mathfrak{g} \end{array} \right\}$$

the space of *connections* on P . Then $\mathcal{A}(P)$ is an affine space modeled on $\Omega^1(X, P(\mathfrak{g}))$, and the affine action is given by $A + \pi^* \mu$ for $A \in \mathcal{A}(P)$ and $\mu \in \Omega^1(X, P(\mathfrak{g}))$. In particular, $\mathcal{A}(P)$ is a smooth (infinite dimensional) manifold with tangent space $\Omega^1(X, P(\mathfrak{g}))$. Each connection $A \in \mathcal{A}(P)$ determines a covariant derivative

$$d_A : \Omega^\bullet(X, P(\mathfrak{g})) \longrightarrow \Omega^{\bullet+1}(X, P(\mathfrak{g})), \quad \mu \longmapsto (\pi^*)^{-1} (d(\pi^* \mu) + [A \wedge \pi^* \mu]),$$

where d is the trivial connection on the trivial bundle $P \times \mathfrak{g}$. Define the *curvature* of A by

$$F_A = (\pi^*)^{-1} \left(dA + \frac{1}{2} [A \wedge A] \right) \in \Omega^2(X, P(\mathfrak{g})).$$

It follows that $d_A(d_A(\mu)) = [F_A \wedge \mu]$ for all $\mu \in \Omega^\bullet(X, P(\mathfrak{g}))$. It is convenient to view $\Omega^\bullet(X, P(\mathfrak{g}))$ as acting on itself by left multiplication, in which case we can write $d_A \circ d_A = F_A$. In identifying $\Omega^\bullet(X, P(\mathfrak{g}))$ with $\Omega^\bullet(P, \mathfrak{g})_{\text{basic}}$ we will typically drop π^* from the notation. For example, the covariant derivative and curvature satisfy the following

$$d_{A+\mu} = d_A + \mu, \quad F_{A+\mu} = F_A + d_A\mu + \frac{1}{2}[\mu \wedge \mu], \quad d_A F_A = 0.$$

The last is the *Bianchi identity*.

Given a connection $A \in \mathcal{A}(P)$ with covariant derivative d_A , we define the *formal adjoint* on k -forms by $d_A^* := -(-1)^{(n-k)(k-1)} * d_A * \mu$. Stokes' theorem shows that this satisfies $(d_A \mu, \nu) = (\mu, d_A^* \nu)$, where (\cdot, \cdot) is the L^2 -pairing, and μ, ν have compact support in the interior of X . It follows from the scaling properties of forms that $d_A^{*c} = c^{-2} d_A^*$, where d_A^{*c} is the adjoint defined with respect to the metric $c^2 g$.

We will be interested in the *flat* connections $A \in \mathcal{A}(P)$. By definition, these satisfy $F_A = 0$, and we will denote the set of flat connections on P by $\mathcal{A}_{\text{flat}}(P)$. If A is flat then $\text{im } d_A \subseteq \ker d_A$ and we can form the *harmonic spaces*

$$H_A^k := H_A^k(X, P(\mathfrak{g})) := \frac{\ker(d_A|_{\Omega^k(X, P(\mathfrak{g}))})}{\text{im}(d_A|_{\Omega^{k-1}(X, P(\mathfrak{g}))})}, \quad H_A^\bullet := \bigoplus_k H_A^k.$$

In general, we say that a connection A is *irreducible* if d_A is injective on $\Omega^0(X, P(\mathfrak{g}))$ (when A is flat this is equivalent to $H_A^0 = 0$). It follows from (9) below that if A is irreducible, then the gauge action at A is locally free.

Suppose now that X is closed. Then the Hodge isomorphism [27, Theorem 6.8] says

$$H_A^\bullet \cong \ker(d_A \oplus d_A^*), \quad \Omega^\bullet(X, P(\mathfrak{g})) \cong H_A^\bullet \oplus \text{im}(d_A) \oplus \text{im}(d_A^*), \quad (7)$$

for any flat connection A on X . Here the decomposition is orthogonal with respect to the L^2 -inner product defined above. We will treat these isomorphisms as identifications. From the first isomorphism in (7) we see that H_A^\bullet is finite dimensional since $d_A \oplus d_A^*$ is elliptic (locally its leading order term is $d \oplus d^*$). Furthermore, it is clear that $* : H_A^\bullet \rightarrow H_A^{n-\bullet}$ restricts to an isomorphism on the harmonic spaces, and so the pairing (6) continues to be non-degenerate when restricted to the harmonic spaces.

Example 2.3. Suppose $X = \Sigma$ is a closed, oriented surface equipped with a metric. Then on 1-forms $*$ squares to -1 , and so defines a complex structure on $\Omega^1(\Sigma, P(\mathfrak{g}))$. Furthermore, the pairing in (6) determines a symplectic structure ω on $\Omega^1(\Sigma, P(\mathfrak{g}))$ and the triple $(\Omega^1(\Sigma, P(\mathfrak{g})), *, \omega)$ is Kähler. Similarly, $(H_A^1, *, \omega)$ is Kähler and finite-dimensional.

The *Yang-Mills functional* is defined to be the map

$$\mathcal{YM} : \mathcal{A}(P) \longrightarrow \mathbb{R}, \quad A \longmapsto \frac{1}{2} \|F_A\|_{L^2}^2.$$

We will refer to the value at A as the *energy* of the connection A . Suppose X is a 4-manifold. On 2-forms the Hodge star squares to the identity, and we denote by $\Omega^\pm(X, P(\mathfrak{g}))$ the ± 1 eigenspace of $*$. Then we have an L^2 -orthogonal decomposition $\Omega^2(X, P(\mathfrak{g})) = \Omega^+(X, P(\mathfrak{g})) \oplus \Omega^-(X, P(\mathfrak{g}))$, with the orthogonal projection to $\Omega^\pm(X, P(\mathfrak{g}))$ given by $\mu \mapsto \frac{1}{2}(1 \pm *)\mu$. The elements of $\Omega^-(X, P(\mathfrak{g}))$ are called *anti-self dual 2-forms*. A connection A is said to be *anti-self dual (ASD)* or an *instanton* if its curvature $F_A \in \Omega^-(X, P(\mathfrak{g}))$ is an anti-self dual 2-form. The instantons are the minimizers of the Yang-Mills functional.

A *gauge transformation* is an equivariant bundle map $u : P \rightarrow P$ covering the identity. The set of gauge transformations on P forms an infinite-dimensional Lie group, called the *gauge group*, and is denoted $\mathcal{G}(P)$. This can be equivalently described as (i) the group of G -equivariant maps $P \rightarrow G$, or (ii) the group of sections of the fiber bundle $P \times_G G \rightarrow X$; in both cases G is acting on itself by conjugation. From this latter description it is clear that the space $\Omega^0(X, P(\mathfrak{g})) = \Gamma(P \times_G \mathfrak{g})$ can be identified with the Lie algebra of $\mathcal{G}(P)$. Moreover, given $\xi \in \Omega^0(X, P(\mathfrak{g}))$, we define $\exp(\xi) \in \mathcal{G}(P)$ to be the section $x \mapsto \exp(-\xi(x))$, where in this latter formula \exp is the exponential for G (the negative sign is introduced to absorb a negative sign later). We define $\mathcal{G}_0(P) \subseteq \mathcal{G}(P)$ to be the connected component of the identity gauge transformation.

Remark 2.4. Fix a faithful matrix representation $G \rightarrow \mathrm{GL}(\mathbb{C}^n) \subset \mathrm{End}(\mathbb{C}^n)$. This induces a Lie algebra representation $\mathfrak{g} \rightarrow \mathrm{End}(\mathbb{C}^n)$. Then $\Omega^0(X, P(\mathfrak{g}))$ and $\mathcal{G}(P)$ both embed into $\Gamma(P \times_G \mathrm{End}(\mathbb{C}^n))$. With respect to this embedding, it makes sense to multiply elements of $\mathcal{G}(P)$ with elements of its Lie algebra.

The gauge group acts on the left on the space $\Omega^\bullet(P, \mathfrak{g})$ by $(u, A) \mapsto (u^{-1})^*A$, for $u \in \mathcal{G}(P)$, $A \in \Omega^\bullet(P, \mathfrak{g})$. Here the star indicates the pullback map induced by the bundle map $u^{-1} : P \rightarrow P$. The action of $\mathcal{G}(P)$ on $\Omega^\bullet(P, \mathfrak{g})$ restricts to actions on $\Omega^\bullet(P, \mathfrak{g})_{\mathrm{basic}}$ and $\mathcal{A}(P)$. Viewing a gauge transformation as a map $u : P \rightarrow G$ we can write this latter action as

$$(u^{-1})^*A = uAu^{-1} + ud(u^{-1}), \quad (8)$$

where we have fixed a faithful matrix representation as in Remark 2.4, and so the concatenation appearing on the right is just matrix multiplication. The curvature of $A \in \mathcal{A}(P)$ transforms under $u \in \mathcal{G}(P)$ by $F_{(u^{-1})^*A} = \mathrm{Ad}(u)F_A$. This shows that $\mathcal{G}(P)$ restricts to an action on $\mathcal{A}_{\mathrm{flat}}(P)$ and, in 4-dimensions, the instantons. The infinitesimal action of $\mathcal{G}(P)$ at $A \in \mathcal{A}(P)$ takes the form

$$\Omega^0(X, P(\mathfrak{g})) \longrightarrow \Omega^1(X, P(\mathfrak{g})), \quad \xi \longmapsto d_A \xi. \quad (9)$$

Suppose X is compact. Fix a reference connection $A_{\mathrm{ref}} \in \mathcal{A}(P)$. This combines with the Levi-Civita connection to define a norm on $W^{k,p}(\wedge^1 T^*X \otimes P(\mathfrak{g}))$, as above. Then we set

$$\mathcal{A}^{k,p}(P) := A_{\mathrm{ref}} + W^{k,p}(T^*X \otimes P(\mathfrak{g})).$$

The space $\mathcal{A}^{k,p}(P)$ is independent of the choice of reference connection A_{ref} , and the norm only depends on A_{ref} up to norm equivalence. Moreover, $\mathcal{A}^{k,p}(P)$ is an affine space modeled on the vector space $W^{k,p}(T^*X \otimes P(\mathfrak{g}))$, and so Sobolev's embedding theorems carry over directly to $\mathcal{A}^{k,p}(P)$.

Given an embedding $G \subseteq U(r)$, define $\mathcal{G}^{k,p}(P)$ to be the subset of sections in $W^{k,p}(\text{End}(\mathbb{C}^r))$ whose images lie in $G \subset U(r) \subset \text{End}(\mathbb{C}^r)$. Under such an embedding, G necessarily has measure zero in $\text{End}(\mathbb{C}^r)$; nonetheless, this is a meaningful definition whenever we are in the continuous range for Sobolev embedding (e.g. $kp \geq n$ and $k \geq 2$, or $kp > n$ and $k = 1$). The space $\mathcal{G}^{k,p}(P)$ forms a group when we are in the continuous range, and the group operations are smooth, making $\mathcal{G}^{k,p}(P)$ a Banach Lie group. Moreover, in the continuous Sobolev range the group $\mathcal{G}^{k,p}(P)$ acts smoothly on $\mathcal{A}^{k-1,p}(P)$. See [28, Appendix B] for more details. When X is non-compact, it is convenient to consider the space $\mathcal{G}_{\text{loc}}^{k,p}(P)$ of sections that are *locally* $W^{k,p}$, defined in the obvious way.

Next we address the case where $X = \mathbb{R} \times Y$, and Y is compact. We assume that $P = \mathbb{R} \times Q$ is the pullback of a principal G -bundle Q over Y . Fix connections $a^\pm \in \mathcal{A}(Q)$ on Y . Let A_{ref} be any connection on $\mathbb{R} \times Q$ that is constantly equal to a^- on $(-\infty, -S_0] \times Y$, and a^+ on $[S_0, \infty) \times Y$ for some $S_0 > 0$. Then we set

$$\mathcal{A}^{k,p}(a^-, a^+) := A_{\text{ref}} + W^{k,p}(T^*X \otimes P(\mathfrak{g})). \quad (10)$$

This is a Banach affine space, and all elements limit to a^\pm at $\pm\infty$. Moreover, as a topological space, it is independent of the choice of reference connection that agrees with the a^\pm off of a compact set. However, the particular choice of Banach norm depends on the choice of A_{ref} . The appropriate gauge group is

$$\mathcal{G}_c^{k,p}(\mathbb{R} \times Q), \quad (11)$$

which we define to be the $W^{k,p}$ -closure of the elements of $\mathcal{G}_{\text{loc}}^{k,p}(\mathbb{R} \times Q)$ that restrict to the identity element off of compact sets. It follows that $\mathcal{G}_c^{k,p}(\mathbb{R} \times Q)$ is a Banach Lie group with Lie algebra the $W^{k,p}$ -completion of the compactly supported 0-forms with values in $Q(\mathfrak{g})$. Moreover, $\mathcal{G}_c^{k,p}(\mathbb{R} \times Q)$ acts smoothly on the space $\mathcal{A}^{k-1,p}(a^-, a^+)$, at least when k, p are in the appropriate Sobolev range.

Let s denote the coordinate variable on \mathbb{R} , and $\frac{\partial}{\partial s} \in \text{Vect}(\mathbb{R} \times P)$ the obvious vector field. Then a connection $A \in \Omega^1(\mathbb{R} \times P, \mathfrak{g})$ is in *temporal gauge* if its contraction with $\partial/\partial s$ vanishes $\iota_{\partial/\partial s} A = 0$. It can be shown that for each connection A , there is an identity-component gauge transformation $u \in \mathcal{G}_0(\mathbb{R} \times P)$ with $u^* A$ in temporal gauge. Note that any connection $A \in \mathcal{A}(\mathbb{R} \times P)$ can be written in the form $A = a(s) + p(s) ds$ for unique $a : \mathbb{R} \rightarrow \mathcal{A}(P)$ and $p : \mathbb{R} \rightarrow \Omega^0(X, P(\mathfrak{g}))$, where $ds \in \Omega^1(\mathbb{R} \times Y)$ is the obvious 1-form. Then A is in temporal gauge if and only if $p = 0$.

We will need a *perturbed* version of Uhlenbeck's compactness theorem for flat connections and instantons. Specifically, we fix a $\mathcal{G}(P)$ -equivariant map

$V : \mathcal{A}(P) \rightarrow \Omega^2(X, P(\mathfrak{g}))$ which is smooth with respect to the $W^{k,p}$ -topology on each space, where $k, p \geq 1$. We assume the following:

- (a) V is uniformly L^∞ -bounded, $\sup_{A \in \mathcal{A}(P)} \|V(A)\|_{L^\infty(X)} < \infty$.
- (b) If $\{A_n\}$ is a sequence of connections such that $\|A_n - A_0\|_{L^p(X)}$ is bounded, then $V(A_n)$ has an L^p -convergent subsequence.

In particular, Kronheimer proves that these conditions are satisfied whenever V is the differential of a holonomy perturbation [23, Lemma 10].

Theorem 2.5. [23, Proposition 11] *Suppose V satisfies conditions (a) and (b) above, and $\{A_n\}$ is a sequence of connections on a fixed bundle over X such that $\|F_{A_n} - V(A_n)\|_{L^p(X)}$ is bounded. Then there is a finite set of points $B \subset X$, a subsequence (still denoted $\{A_n\}$), and a sequence of gauge transformations $\{u_n\}$ such that $u_n^* A_n$ converges weakly in $W^{1,p}$ on compact subsets of $X \setminus B$ to a connection A_∞ defined on $X \setminus B$. If $F_{A_n} - V(A_n) = 0$ for all n , then this limiting connection satisfies $F_{A_\infty} - V(A_\infty) = 0$, the sequence $u_n^* A_n$ converges strongly in $W^{1,p}$ on compact subsets of $X \setminus B$, and A_∞ is defined on a (possibly different) bundle over X .*

If $\dim X = 4$, then the same result holds with $F_A - V(A)$ replaced by $(F_A - V(A))^+$, where the superscript $+$ denotes the anti-self dual component.

Consider the case where $(F_{A_n} - V(A_n))^+ = 0$ for all n . When V is zero, the (unperturbed version of) Uhlenbeck's compactness theorem implies that a subsequence of $u_n^* A_n$ converges in C^∞ to A_∞ . Kronheimer points out that when V is a (non-zero) holonomy perturbation, then one can no longer expect to get anything stronger than the $W^{1,p}$ convergence of Theorem 2.5. This is due to the non-locality inherent in holonomy perturbations.

As we have already begun to see, when the base manifold X is a product, it is useful to have a notation that distinguishes between a connection or gauge transformation and its components. We therefore adopt the following notation:

- On 4-manifolds we will use A and U for connections and gauge transformations.
- We tend to denote 3-manifolds by Y , and use a and u for connections and gauge transformations on Y .
- We tend to denote 2-manifolds by Σ , and use α and μ for connections and gauge transformations on Σ .

On the other hand, when the base manifold has no relevant product structure, then we continue to use A and u for connections and gauge transformations.

2.3 Topological aspects of principal $\mathrm{PU}(r)$ -bundles

Throughout the remainder of this paper we work primarily with the Lie group $\mathrm{PU}(r) := \mathrm{SU}(r)/\mathbb{Z}_r = \mathrm{U}(r)/\mathrm{U}(1)$, for $r \geq 2$. It follows that $\pi_1(\mathrm{PU}(r)) = \mathbb{Z}_r$, since $\mathrm{SU}(r)$ is simply-connected. We fix an Ad-invariant inner product $\langle \cdot, \cdot \rangle$ on the Lie algebra $\mathfrak{pu}(r)$ to $\mathrm{PU}(r)$. This Lie algebra is simple, and so $\langle \mu, \nu \rangle = -\kappa_r \mathrm{tr}(\mu \cdot \nu)$, for some $\kappa_r > 0$, where the trace is the one induced from the identification $\mathfrak{pu}(r) \cong \mathfrak{su}(r) \subset \mathrm{End}(\mathbb{C}^r)$. We leave κ_r arbitrary, but fixed.

In [32], L.M. Woodward exploited the adjoint representation to classify the principal $\mathrm{PU}(r)$ -bundles over spaces of dimension ≤ 4 . This classification scheme assigns cohomology classes

$$t_2(P) \in H^2(X, \mathbb{Z}_r), \quad q_4(P) \in H^4(X, \mathbb{Z})$$

to each principal $\mathrm{PU}(r)$ -bundle $P \rightarrow X$. For example, q_4 is defined to be the second Chern class of the complexified adjoint bundle $P(\mathfrak{g})_{\mathbb{C}} := P(\mathfrak{g}) \otimes \mathbb{C}$, where $\mathfrak{g} = \mathfrak{pu}(r)$. The class t_2 is defined as the mod r reduction of a suitable first Chern class. We will be mostly interested in the case where X is a smooth manifold, but these classes are defined for CW complexes as well.

Example 2.6. *When $r = 2$ we have $\mathrm{PU}(2) = \mathrm{SO}(3)$, and the classes t_2 and q_4 are the second Stiefel-Whitney class and the (negative of the) first Pontryagin class, respectively.*

We summarize the properties of these classes that we will need.

- If $\dim(X) \leq 4$ and X is a manifold, then two bundles P and P' over X are isomorphic if and only if $t_2(P) = t_2(P')$ and $q_4(P) = q_4(P')$;
- The class $t_2(P)$ is zero if and only if the structure group of P can be lifted to $\mathrm{SU}(r)$.
- The class $t_2(P)$ is the mod r reduction of an integral class if and only if the structure group of P can be lifted to $\mathrm{U}(r)$.

These characteristic classes can be used to study the components of the gauge group $\mathcal{G}(P)$ for a principal $\mathrm{PU}(r)$ -bundle $P \rightarrow X$. We take as a starting point the following observation by Donaldson. We refer the reader to [12] for more details of the assertions in the remainder of this section.

Proposition 2.7. [5, Section 2.5.2] *Suppose G is a compact Lie group, X is a smooth manifold and $P \rightarrow X$ is a principal G -bundle. Then there is a bijection between $\pi_0(\mathcal{G}(P))$ and the set of isomorphism classes of principal G -bundles over $S^1 \times X$ that restrict to P on a fiber. This bijection is induced from the map that sends a gauge transformation u to the bundle $P_u := [0, 1] \times P / (0, u(p)) \sim (1, p)$, given by the mapping torus of u .*

Now we combine this with L.M. Woodward's classification. Fix a principal $\mathrm{PU}(r)$ -bundle $P \rightarrow X$. To focus the discussion, we assume X is a manifold of dimension at most 3. Let $u \in \mathcal{G}(P)$, define $P_u \rightarrow S^1 \times X$ as in the proposition above, and consider the classes $t_2(P_u) \in H^2(S^1 \times X, \mathbb{Z}_r)$, $q_4(P_u) \in H^4(S^1 \times X, \mathbb{Z})$. By the Künneth formula, we have an isomorphism

$$H^k(S^1 \times X, R) \cong H^k(X, R) \oplus H^{k-1}(X, R),$$

where $R = \mathbb{Z}$ or \mathbb{Z}_r . The image of $t_2(P_u) \in H^2(X, \mathbb{Z}_r) \oplus H^1(X, \mathbb{Z}_r)$ in the first factor is exactly $t_2(P)$, so the dependence of $t_2(P_u)$ on the gauge transformation u is contained entirely in the projection of $t_2(P_u)$ to the second factor $H^1(X, \mathbb{Z}_r)$. We denote this projection by

$$\eta(u) := \eta_X(u) \in H^1(X, \mathbb{Z}_r),$$

and call this the *parity* of u .

In a similar fashion, we will define the degree of a gauge transformation using q_4 . Before defining this explicitly, we note that a straight-forward characteristic class argument shows that the image of $q_4(P_u)$ in $H^3(X, \mathbb{Z})$ is always even. Then we define the *degree* of u to be the class

$$\deg(u) := \frac{1}{2} \mathrm{proj}_{H^3} q_4(P_u) \in H^3(X, \mathbb{Z}).$$

where $\mathrm{proj}_{H^3} : H^4(S^1 \times X, \mathbb{Z}) \rightarrow H^3(X, \mathbb{Z})$ is given by contraction with the generator of $H^1(S^1, \mathbb{Z})$. See [8] and [15] for alternative realizations of the parity and degree.

Proposition 2.7, and the properties of t_2 and q_4 immediately imply that the parity and degree detect the components of the gauge group when $\dim(X) \leq 3$. The next proposition describes some information captured by the parity. It is standard when $r = 2$ and the extension to larger r is no different.

Proposition 2.8. *Fix a gauge transformation $u : P \rightarrow \mathrm{PU}(r)$. Then the following are equivalent:*

- $\eta(u) = 0$;
- $u : P \rightarrow \mathrm{PU}(r)$ lifts to a $\mathrm{PU}(r)$ -equivariant map $\tilde{u} : P \rightarrow \mathrm{SU}(r)$;
- When restricted to the 1-skeleton of X , u is homotopic to the identity map.

If X is a compact, connected, oriented 3-manifold and $\eta(u) = 0$, then $\deg(u)$ is divisible by r .

Next we give a formula for computing the degree. Let $Q \rightarrow Y$ be a $\mathrm{PU}(r)$ -bundle over a closed, connected, oriented 3-manifold Y , and let $u \in \mathcal{G}(Q)$. Fix a connection $a_0 \in \mathcal{A}(Q)$, and let $a : I \rightarrow \mathcal{A}(Q)$ be any path from a_0 to u^*a_0 . This defines a connection A on $I \times Q \rightarrow I \times Y$ by declaring $A|_{\{s\} \times Y} = a(s)$ for $s \in I$. Moreover, A descends to a connection on the mapping torus Q_u .

The usual Chern-Weil formula for the second Chern class of $SU(r)$ -bundles gives $q_4(Q_u) = -r(4\pi^2\kappa_r)^{-1} \int_{I \times Y} \langle F_A \wedge F_A \rangle$. The curvature decomposes into components as $F_A = F_{a(s)} + ds \wedge \partial_s a(s)$, and so

$$\deg(u) = \frac{r}{4\pi^2\kappa_r} \int_I \left(\int_Y \langle F_{a(s)} \wedge \partial_s a(s) \rangle \right). \quad (12)$$

(Our orientation convention is that $ds \wedge d\text{vol}_Y$ is a positive volume form on $I \times Y$.) This formula also shows that the degree operator is a homomorphism of groups.

We end the discussion of the degree and parity by mentioning two applications. These will be useful in our treatment of instanton Floer theory later.

Application 1: Free actions on flat connections. We suppose X is a connected and oriented manifold of dimension 2 or 3, equipped with a principal $PU(r)$ -bundle $P \rightarrow X$. In case X is 2-dimensional, we assume that $t_2(P) [X] \in \mathbb{Z}_r$ is a generator. If X is 3-dimensional then we assume there is an embedding $\Sigma \hookrightarrow X$ of a connected, oriented surface such that $t_2(P) [\Sigma] \in \mathbb{Z}_r$ is a generator.

Lemma 2.9. [8, Lemma 2.5] *Suppose $P \rightarrow X$ is as above. Then all of the flat connections on P are irreducible. In fact, any flat connection a on P has trivial stabilizer in $\ker \eta$:*

$$\{u \in \mathcal{G}(P) \mid \eta(u) = 0, \quad u^*a = a\} = \{e\}.$$

Application 2: Existence of degree d gauge transformations. Fix a closed, connected, oriented 3-manifold X and a principal $PU(r)$ -bundle $P \rightarrow X$. We assume that $t_2(P)$ is the reduction of an integral class. This implies $P = \bar{P} \times_{U(r)} PU(r)$ is induced from a principal $U(r)$ -bundle $\bar{P} \rightarrow X$, and $t_2(P) = c_1(\bar{P}) \bmod r$. Let $d \in \mathbb{Z}$, and assume that $c_1(\bar{P}) [\Sigma] = d$ for some embedded connected, oriented surface $\Sigma \hookrightarrow X$. (In particular, P satisfies the conditions of Application 1 when d, r are relatively prime.)

Proposition 2.10. *Let $d \in \mathbb{Z}$, $P \rightarrow X$ and $\Sigma \subset X$ be as above.*

(a) *There exists a gauge transformation $u \in \mathcal{G}(P)$ of degree d . Moreover, the parity $\eta(u) : H_1(X) \rightarrow \mathbb{Z}_r$ is given by the mod r reduction of the Poincaré dual $PD_X [\Sigma] : H_1(X) \rightarrow \mathbb{Z}$.*

(b) *Suppose d and r are relatively prime. Then there exists a gauge transformation $u \in \mathcal{G}(P)$ of degree 1. Moreover, writing $md + nr = 1$, the parity $\eta(u) : H_1(X) \rightarrow \mathbb{Z}_r$ is given by the mod r reduction of m times the Poincaré dual $PD_X [\Sigma] : H_1(X) \rightarrow \mathbb{Z}$.*

In the case $r = 2$, Dostoglou and Salamon [8, Lemmas 2.3 and A.2] prove this by explicitly constructing the desired gauge transformation. In [12], we present a proof for general $r \geq 2$ using the properties of the characteristic classes t_2, q_4 .

2.4 Compatible bundles for broken circle fibrations

Fix an integer $r \geq 2$ and $d \in \mathbb{Z}_r$. Let $f : Y \rightarrow S^1$ be a broken circle fibration. Suppose we are given a map $\gamma : S^1 \rightarrow Y$ that is a section of f (these always exist). Then by the classification of principal $\text{PU}(r)$ -bundles on 3-manifolds, there is a unique, up to bundle isomorphism, principal $\text{PU}(r)$ -bundle $Q = Q(\gamma) \rightarrow Y$ such that $t_2(Q) \in H^2(Y, \mathbb{Z}_r)$ is Poincaré dual to the class $d[\gamma] \in H_1(Y, \mathbb{Z}_r)$. We will say that a bundle Q is *d-compatible* with $f : Y \rightarrow S^1$ if it can be obtained in this way. Since $[\gamma]$ is the reduction of an integral class, it follows that $t_2(Q)$ is as well. Moreover, $t_2(Q)[\Sigma_i] = d$ for all i .

In general, there may be multiple isomorphism classes of bundles that are *d-compatible* with f . For example, suppose γ is a section of f , and $\alpha : S^1 \rightarrow \Sigma_0$ is a loop in the fiber Σ_0 . Then the concatenation $\gamma * \alpha$ is homologous in $H_1(Y)$ to a section γ' of f . The associated bundles $Q(\gamma)$ and $Q(\gamma')$ will be isomorphic if and only if γ and γ' induce the same element of $H_1(Y, \mathbb{Z}_r)$.

Recall the decomposition (2) of a broken circle fibration $Y \rightarrow S^1$ into elementary cobordisms $I \times \Sigma_j$ and $Y_{j(j+1)}$. Given a *d-compatible* bundle $Q \rightarrow Y$, we set

$$Q_\bullet := \bigcup_i Q_{i(i+1)}, \quad Q_{i(i+1)} := Q|_{Y_{i(i+1)}}, \quad P_\bullet := \bigcup_i P_i, \quad P_i := Q|_{\Sigma_i}.$$

Then Q_\bullet and P_\bullet are naturally bundles over Y_\bullet and Σ_\bullet . Let $F_\epsilon : Y^1 \rightarrow Y^\epsilon$ be the diffeomorphism (5), and define $Q^\epsilon := (F_\epsilon^{-1})^*Q$. This is a smooth $\text{PU}(r)$ -bundle over Y^ϵ enjoying the same topological properties as $Q \rightarrow Y$. Since F_ϵ is the identity on Y_\bullet and $\{0\} \times \Sigma_\bullet$, it follows that the restrictions $Q_\bullet = Q^\epsilon|_{Y_\bullet}$ and $P_\bullet = Q^\epsilon|_{\{0\} \times \Sigma_\bullet}$ do not depend on ϵ .

We will use a subscript η_Y (resp. η_{Σ_i}) to denote the parity operator for $Q \rightarrow Y$ (resp. $P_i \rightarrow \Sigma_i$). Let $u \in \mathcal{G}(Q)$, and notice that $\eta_{\Sigma_i}(u|_{\Sigma_i}) = 0$ for some i if and only if $\eta_{\Sigma_i}(u|_{\Sigma_i})$ for all i . We will be particularly interested in the following subgroup

$$\mathcal{G}_{\Sigma_\bullet} := \{u \in \mathcal{G}(Q) \mid \eta_{\Sigma_0}(u|_{\Sigma_0}) = 0\}.$$

There is a sequence of inclusions

$$\mathcal{G}_0(Q) \subset \ker \eta_Y \subset \mathcal{G}_{\Sigma_\bullet} \subset \mathcal{G}(Q).$$

Each of these inclusions is strict. For example, the first is strict since any gauge transformation in $\ker \eta_Y$ lifts to an $\text{SU}(r)$ -gauge transformation, and there are certainly $\text{SU}(r)$ -gauge transformations with non-trivial degree. That the second inclusion is strict follows from Proposition 2.10 (a) and the fact that the restriction of $\eta_Y(u)$ to $H^1(\Sigma_0) \subset H^1(Y)$ is exactly $\eta_{\Sigma_0}(u|_{\Sigma_0})$.

The following is the key result regarding $\mathcal{G}_{\Sigma_\bullet}$.

Proposition 2.11. *Suppose $d \in \mathbb{Z}_r$ is a generator and $Q \rightarrow Y$ is *d-compatible* with a broken circle fibration $f : Y \rightarrow S^1$. Then the degree descends to an isomorphism $\mathcal{G}_{\Sigma_\bullet}/\mathcal{G}_0(Q) \cong \mathbb{Z}$.*

See [12, Proposition 4.8] for a proof. It follows that the positive generator of $\mathcal{G}_{\Sigma_\bullet}/\mathcal{G}_0(Q) \cong \mathbb{Z}$ is the homotopy class of the degree 1 gauge transformation from Proposition 2.10 (b). Moreover, by Proposition 2.8, the image of $\ker \eta_Y$ in $\mathcal{G}_{\Sigma_\bullet}/\mathcal{G}_0(Q)$ is identified with the multiples of r :

$$\ker \eta_Y / \mathcal{G}_0(Q) \cong r\mathbb{Z} \subset \mathbb{Z}.$$

2.5 Quilted Floer theory

In this section we describe how to assign a quilted Floer cohomology group to a d -compatible bundle Q over a broken circle fibration $f : Y \rightarrow S^1$. We begin by introducing the moduli spaces of flat connections on surfaces and cobordisms. These bring symplectic geometry into the picture.

2.5.1 Moduli spaces of flat connections

Fix a principal $\mathrm{PU}(r)$ -bundle $P \rightarrow X$, and we assume $\dim(X) \leq 3$. For $2 \leq q < \infty$ and $k \geq 1$, define

$$M(P) := \mathcal{A}_{\mathrm{flat}}^{k,q}(P) / \mathcal{G}_0^{k+1,q}(P),$$

where $\mathcal{G}_0^{k+1,q}(P) \subseteq \mathcal{G}^{k+1,q}(P)$ is the identity component of the gauge group, and let $\Pi : \mathcal{A}_{\mathrm{flat}}^{1,q}(P) \rightarrow M(P)$ be the quotient map. Then Π and $M(P)$ are independent of the choice of $2 \leq q < \infty$ and $k \geq 1$ in the sense that the natural inclusion $\mathcal{A}_{\mathrm{flat}}^{k,q}(P) \hookrightarrow \mathcal{A}_{\mathrm{flat}}^{1,2}(P)$ induces a homeomorphism

$$\mathcal{A}_{\mathrm{flat}}^{k,q}(P) / \mathcal{G}_0^{k+1,q}(P) \cong \mathcal{A}_{\mathrm{flat}}^{1,2}(P) / \mathcal{G}_0^{2,2}(P).$$

Moreover, this is a diffeomorphism whenever the gauge group acts freely.

Now suppose $X = \Sigma$ is a closed, connected, oriented surface, and $P \rightarrow \Sigma$ is a principal $\mathrm{PU}(r)$ -bundle, with $t_2(P) [\Sigma] \in \mathbb{Z}_r$ a generator. It follows from Lemma 2.9 that $\mathcal{G}_0^{2,q}(P)$ acts freely on the space of flat connections, and so it is immediate that the moduli space $M(P)$ is a smooth manifold. The next theorem goes back to Atiyah-Bott [2], and captures the relevant properties of $M(P)$.

Theorem 2.12. *Let Σ be a closed, connected, oriented surface, and $P \rightarrow \Sigma$ a principal $\mathrm{PU}(r)$ -bundle with $t_2(P) [\Sigma] \in \mathbb{Z}_r$ a generator.*

1. [29, Prop. 3.2.4, Thm. 3.3.2] *If Σ has genus $g(\Sigma) \geq 1$, then the moduli space $M(P)$ is a nonempty compact, symplectic manifold of dimension $(2g(\Sigma) - 2)(r^2 - 1)$, with even minimal Chern number, and with monotonicity constant $2\pi^2\kappa_r/r$. The tangent space at $[\alpha] \in M(P)$ is isomorphic to H_α^1 and the symplectic form $\omega_{M(P)}$ is given by restricting the pairing (6). If $g(\Sigma) = 0$, then $M(P) = \emptyset$. Moreover, $M(P)$ is always connected and simply-connected.*

2. [29, Lemma 3.3.5] *If $P' \rightarrow \Sigma$ is a second principal $\mathrm{PU}(r)$ -bundle then any $\mathrm{PU}(r)$ -equivariant bundle isomorphism $\psi : P \rightarrow P'$ covering the identity induces a*

symplectomorphism $\psi^* : M(P') \rightarrow M(P)$ by pullback. Furthermore, if $\phi : P \rightarrow P'$ is a second bundle map then $\psi^* = \phi^*$, so the moduli spaces $M(P)$ and $M(P')$ are canonically symplectomorphic.

Here, κ_r is the constant coming from our choice of metric on \mathfrak{g} as in Section 2.3. In [29] the authors work with $G = \mathrm{U}(r)$ and the space of *central curvature connections with fixed determinant*, rather than the space of flat $\mathrm{PU}(r)$ connections, as we consider here. Their theorems carry over verbatim to our situation: On a surface, the space of central curvature $\mathrm{U}(r)$ -connections with fixed determinant is naturally diffeomorphic to the space of flat $\mathrm{PU}(r)$ -connections, and this diffeomorphism intertwines the relevant gauge group actions (see [29, Lemma 3.2.5]).

Remark 2.13. (a) Since Σ is assumed to be oriented, any choice of metric g on Σ induces a Hodge star, $*$. Then $*$ descends to a compatible complex structure on the tangent bundle $TM(P)$, which we denote by $*$ or J_g , depending on the context.

(b) The symplectic form on $M(P)$ is given by integration over the surface Σ , so reversing the orientation of Σ changes $(M(P), \omega_{M(P)})$ to $M(P)^- := (M(P), -\omega_{M(P)})$.

The manifold $M(P)$ has two alternative descriptions that are often quite useful. We refer the reader to [2, 8, 29] for more details. For the first description, fix a basepoint $x_0 \in \Sigma$. Then the holonomy determines a map $\mathcal{A}_{\mathrm{flat}}(P)/\mathcal{G}(P) \rightarrow \mathrm{Hom}(\pi_1(\Sigma, x_0), \mathrm{PU}(r))/\mathrm{PU}(r)$ that is a homeomorphism onto a component of the right-hand side, and $\mathrm{PU}(r)$ is acting on the Hom-set by conjugation. We will use the notation $\mathrm{Hom}_{\mathrm{PU}(r)}^d$ to refer to this component. On the other hand, quotienting $\mathcal{A}_{\mathrm{flat}}(P)$ only by the identity component $\mathcal{G}_0(P)$ of the gauge group recovers $M(P)$. In particular, we obtain a map $M(P) \rightarrow \mathrm{Hom}_{\mathrm{PU}(r)}^d$ that is a $\mathbb{Z}_r = \mathcal{G}(P)/\mathcal{G}_0(P)$ -covering space away from the singularities of $\mathrm{Hom}_{\mathrm{PU}(r)}^d$.

For the second description, fix a basepoint $p_0 \in P$ over $x_0 \in \Sigma$. The fundamental group of P is generated by tuples $(A_1, B_1, \dots, A_g, B_g, z)$ subject to the relations $\prod_{j=1}^g [A_j, B_j] = z$ and $[A_j, z] = [B_j, z] = z^r = \mathrm{Id}$. Consider the space

$$\mathrm{Hom}_{\mathrm{SU}(r)}^d := \left\{ \rho \in \mathrm{Hom}(\pi_1(P, p_0), \mathrm{SU}(r)) \mid \rho(z) = e^{2\pi i d/r} \mathrm{Id} \right\}$$

The group $\mathrm{SU}(r)$ acts freely on $\mathrm{Hom}_{\mathrm{SU}(r)}^d$ by conjugation and quotienting by this action recovers $M(P) \cong \mathrm{Hom}_{\mathrm{SU}(r)}^d / \mathrm{SU}(r)$. This diffeomorphism arises by sending a flat connection to its $\mathrm{SU}(r)$ -valued holonomy in $\mathrm{Hom}(\pi_1(P), \mathrm{SU}(r))$.

Next, suppose Y is a compact, connected, oriented cobordism between non-empty, connected, oriented surfaces Σ^\pm . Our orientation convention is that $\partial Y = \overline{\Sigma^-} \sqcup \Sigma^+$, where the bar denotes the manifold with the opposite orientation. Let $Q \rightarrow Y$ be a principal $\mathrm{PU}(r)$ -bundle and assume that $t_2(Q) [\Sigma^+] \in \mathbb{Z}_r$ is a generator. It follows that $t_2(Q) [\Sigma^+] = t_2(Q) [\Sigma^-]$. In particular, Lemma 2.9 implies that the moduli spaces $M(Q)$ and $M(Q|_{\Sigma^\pm})$ of flat connections are all *smooth* manifolds.

Restriction to each boundary component of Y induces a $\mathcal{G}(Q)$ -equivariant map $\rho : \mathcal{A}(Q) \rightarrow \mathcal{A}(Q|_{\Sigma^-}) \times \mathcal{A}(Q|_{\Sigma^+})$ that preserves the flat connections. It follows that ρ descends to a map, still denoted by ρ , at the level of moduli spaces. Then we set

$$L(Q) := \rho(M(Q)) \subset M(Q|_{\Sigma^-})^- \times M(Q|_{\Sigma^+}).$$

Theorem 2.14. *Let $Q \rightarrow Y$ be as above. Assume, in addition, that Y is an elementary cobordism.*

1. [29, Theorem 3.4.1] *The map $\rho : M(Q) \rightarrow L(Q) \subset M(P^-)^- \times M(P^+)$ is a Lagrangian embedding. Furthermore, $L(Q)$ is compact, oriented, and simply-connected (hence monotone).*

2. [29, Lemma 3.4.4] *The Lagrangian $L(Q)$ is independent of the choice of Q under the canonical symplectomorphisms of Theorem 2.12. If $Y = I \times \Sigma^+$ is a product cobordism from $\Sigma^- = \Sigma^+$ to itself, then $L(Q) \subset M(P^-)^- \times M(P^-)$ is the diagonal.*

2.5.2 Quilted Floer cohomology for broken circle fibrations

Fix a broken circle fibration $f : Y \rightarrow S^1$ and a d -compatible principal $\mathrm{PU}(r)$ -bundle $Q \rightarrow Y$, where $d \in \mathbb{Z}_r$ is a generator. In this section we review the definition from [29] of the *quilted Floer group* associated to this data. Recall the decomposition (2) of Y into elementary cobordisms $Y_{j(j+1)}$ and $I \times \Sigma_j$, and the associated decomposition of the bundle Q from Section 2.4. By construction, each $P_j \rightarrow \Sigma_j$ and $Q_{j(j+1)} \rightarrow Y_{j(j+1)}$ satisfies the conditions of Theorems 2.12 and 2.14, respectively. In the terminology of [30], this implies that the tuple

$$\underline{L}(Q) := \left(L(Q_{01}), \dots, L(Q_{(N-1)0}) \right)$$

determines a cyclic Lagrangian correspondence in the product symplectic manifold

$$\underline{M} := \left(M(P_j) \right)_{j=0}^{N-1}.$$

It also implies that the quilted Floer cohomology group of $\underline{L}(Q)$ is well-defined. We denote this cohomology group by $HF_{\mathrm{symp}}^\bullet(Y, f, Q)_{r,d}$ and call this the *quilted Floer group* of (Y, f, Q, r, d) . Since each $M(P_j)$ has even minimal Chern number, this quilted Floer group admits a relative \mathbb{Z}_4 -grading. Moreover, $HF_{\mathrm{symp}}^\bullet(Y, f, Q)_{r,d}$ only depends on the bundle isomorphism type of Q . The following theorem, due to Gay, Wehrheim and Woodward, addresses the dependence of this group on f . Note that if f and f' are two homotopic broken circle fibrations on Y , then any bundle Q that is d -compatible with f is also d -compatible with f' .

Theorem 2.15. [20, 29] *If f and f' are broken circle fibrations that are homotopic as maps $Y \rightarrow S^1$, then there is a canonical isomorphism $HF_{\mathrm{symp}}^\bullet(Y, f, Q)_{r,d} \cong HF_{\mathrm{symp}}^\bullet(Y, f', Q)_{r,d}$, and this isomorphism preserves the relative gradings.*

The remainder of this section is spent describing $HF_{\text{symp}}^\bullet(Y, f, Q)_{r,d}$ in more detail. We first remark that the group $HF_{\text{symp}}^\bullet(Y, f, Q)_{r,d}$ is the cohomology of a chain complex $(CF_{\text{symp}}^\bullet, \partial_{\text{symp}})$. We will give a precise definition of CF_{symp}^\bullet and ∂_{symp} , beginning with the former.

For each j , fix a time-dependent Hamiltonian $H_j = H_{t,j} : M(P_j) \rightarrow \mathbb{R}$, which we assume vanishes for $t \notin (0, 1)$. Let \underline{H} denote the tuple $(H_j)_j$, and call this a *split-type* Hamiltonian for \underline{M} . The time-dependent function H_j lifts to a $\mathcal{G}_0(P_j)$ -invariant time-dependent function on $\mathcal{A}_{\text{flat}}(P_j)$. Similarly, the Hamiltonian vector field of H_j on $M(P_j)$ lifts to a Hamiltonian vector field $X_{t,j} : \mathcal{A}_{\text{flat}}(P_j) \rightarrow \Omega^1(\Sigma_j, P_j(\mathfrak{g}))$ satisfying

$$X_{t,j}(u^* \alpha_j) = \text{Ad}(u^{-1})X_{t,j}(\alpha_j), \quad \text{and} \quad d_{\alpha_j} X_{t,j}(\alpha_j) = 0,$$

for all $u \in \mathcal{G}_0(P_j)$ and $\alpha_j \in \mathcal{A}_{\text{flat}}(P_j)$. Note that $X_{t,j}$ vanishes when $t \notin (0, 1)$. Then the chain group CF_{symp}^\bullet is freely generated over \mathbb{Z}_2 by the set of *\underline{H} -perturbed generalized Lagrangian intersection points* \underline{e} . These are tuples

$$\underline{e} = \left([a_{01}], [a_{12}], \dots, [a_{(N-1)0}] \right),$$

where $a_{j(j+1)} \in \mathcal{A}_{\text{flat}}(Q_{j(j+1)})$, $[a_{j(j+1)}]$ denotes the $\mathcal{G}_0(Q_{j(j+1)})$ -equivalence class, and each tuple is required to satisfy the following: For each j , there is a map $\alpha_j : I \rightarrow \mathcal{A}_{\text{flat}}(P_j)$ such that

$$\partial_t \alpha_j(t) = X_{t,j}(\alpha_j(t)), \quad [a_{(j-1)j} | \Sigma_j] = [\alpha_j(0)], \quad \text{and} \quad [\alpha_j(1)] = [a_{j(j+1)} | \Sigma_j], \quad (13)$$

where, here, the brackets denote $\mathcal{G}_0(P_j)$ -equivalence class. We use $\mathcal{I}_{\underline{H}}(\underline{L}(Q))$ to denote the set of \underline{H} -perturbed generalized Lagrangian intersection points. This arises naturally as the critical point set of a suitably defined *\underline{H} -perturbed symplectic action functional*, see [30, Section 5.2]. We say $\underline{e} \in \mathcal{I}_{\underline{H}}(\underline{L}(Q))$ is *non-degenerate* if it is a non-degenerate critical point in the Morse-theoretic sense. This is equivalent to requiring that the linearization at \underline{e} of the defining equations (13) becomes an injective operator, modulo the linearized gauge action on the P_j . Unless otherwise specified, we will assume \underline{H} has been chosen so that all elements of $\mathcal{I}_{\underline{H}}(\underline{L}(Q))$ are non-degenerate. This can always be done [30, Proposition 5.2.1] and, when this is the case, it follows that $\mathcal{I}_{\underline{H}}(\underline{L}(Q))$ is a finite set.

Now we move on to discuss the boundary operator ∂_{symp} . By linearity, it suffices to define ∂_{symp} in terms of its matrix coefficient $\langle \partial_{\text{symp}} \underline{e}^-, \underline{e}^+ \rangle$, for $\underline{e}^\pm \in \mathcal{I}_{\underline{H}}(\underline{L}(Q))$. This can be defined concisely by saying that it is the mod-2 count of the isolated $(\underline{J}, \underline{H})$ -holomorphic quilted cylinders that limit to \underline{e}^\pm at $\pm\infty$. Now we unravel this. Here, $\underline{J} = (J_j)_j$ is a *split-type* almost complex structure on \underline{M} , meaning that each J_j is a compatible almost complex structure on $M(P_j)$. In this paper we will always assume that each J_j arises as follows: Fix a metric on Y , and let $J_j = *$ be the compatible almost complex structure on $M(P_j)$ induced by the Hodge star on Σ_j as in Remark 2.13 (a).

Next, let $\underline{H} = (H_j)_j$ be a split-type Hamiltonian as above. Then we define a $(\underline{J}, \underline{H})$ -holomorphic quilted cylinder to be a tuple $\underline{v} = ([\alpha_0], \dots, [\alpha_{N-1}])$ where $[\cdot]$ denotes the $\mathcal{G}_0(P_j)$ -equivalence class for the relevant j , and α_j is a map $\mathbb{R} \times I \rightarrow \mathcal{A}_{\text{flat}}(P_i)$ satisfying the following conditions:

- (($\underline{J}, \underline{H}$)-holomorphic) $\text{proj}_{H_{\alpha_i}^1} (\partial_s \alpha_i + * (\partial_t \alpha_i - X_{t,i}(\alpha_i))) = 0$
- (Lagrangian seam) $([\alpha_i(s, 1)], [\alpha_{i+1}(s, 0)]) \in L(Q_{i(i+1)})$

where $\text{proj}_{H_{\alpha}^1}$ denotes the L^2 -orthogonal projection to the harmonic space H_{α}^1 . We assume that \underline{v} limits to \underline{e}^{\pm} at $\pm\infty$, in the sense that $\lim_{s \rightarrow \pm\infty} \alpha_j(s, t) = \alpha^{\pm}(t)$ for all j , where the α^{\pm} are coming from \underline{e}^{\pm} as in (13). Here we require that the convergence is uniform in t , and when this is the case it implies that the convergence is actually C^{∞} in t . We say that \underline{v} is *regular* if the linearization at \underline{v} of the defining equations is surjective. The desirable cases are when all $(\underline{J}, \underline{H})$ -holomorphic quilted cylinders are regular. When this is the case, we say that $(\underline{J}, \underline{H})$ is *regular (for quilted Floer theory)*. We show in Section 4.2 that, given any $\underline{J} = (*)_j$, there always exists a perturbation \underline{H} so that $(\underline{J}, \underline{H})$ is regular. (In Lagrangian intersection Floer theory it is standard that there exists a perturbation H making (J, H) regular [16]. However the existence becomes more subtle in quilted Floer theory because $\underline{H} = (H_j)$ is required to be of *split-type*; see [30, 31] for a similar problem.) Unless otherwise specified, we will always assume \underline{H} has been chosen so that $(\underline{J}, \underline{H})$ is regular. When this is the case, the set of all $(\underline{J}, \underline{H})$ -holomorphic quilted cylinders limiting to \underline{e}^{\pm} forms a smooth manifold. Finally, we say \underline{v} is *isolated* if it belongs to the zero-dimensional component of this manifold.

Next we describe how a $(\underline{J}, \underline{H})$ -holomorphic quilted cylinder $\underline{v} = (\alpha_0, \dots, \alpha_{N-1})$ can be represented as a connection on the bundle $\mathbb{R} \times Q \rightarrow \mathbb{R} \times Y$. In light of the Hodge isomorphism (7), the $(\underline{J}, \underline{H})$ -holomorphic condition is equivalent to the statement that, for all j ,

$$\partial_s \alpha_j + * (\partial_t \alpha_j - X_{t,j}(\alpha_j)) = d_{\alpha_j} \phi_j + * d_{\alpha_j} \psi_j$$

for some $\phi_j, \psi_j : \mathbb{R} \times I \rightarrow \Omega^0(\Sigma_j, P_j(\mathfrak{g}))$. In fact, ϕ_j, ψ_j are uniquely determined by this equation since all flat connections on P_j are irreducible. Furthermore, ϕ_j and ψ_j are as smooth in s, t as the connection α_j . On the other hand, the *Lagrangian seam* condition implies that, for each j , there is some path $a_{j(j+1)} : \mathbb{R} \rightarrow \mathcal{A}_{\text{flat}}(Q_{j(j+1)})$ with

$$(\alpha_j(s, 1), \alpha_{j+1}(s, 0)) = \left(a_{j(j+1)}(s) \Big|_{\Sigma_j}, a_{j(j+1)}(s) \Big|_{\Sigma_{j+1}} \right),$$

and $a_{j(j+1)}$ is the unique connection having this property, up to the action of gauge transformations on $Y_{j(j+1)}$ that restrict to the identity on the boundary. Then the irreducibility of flat connections on $Q_{j(j+1)}$ implies that there is a unique $p_{j(j+1)} : \mathbb{R} \rightarrow \Omega^0(Y_{j(j+1)}, Q_j(\mathfrak{g}))$ such that $\partial_s a_{j(j+1)} - d_{a_{j(j+1)}} p_{j(j+1)} \in$

$H^1_{a_{j(j+1)}}$ and $p_{j(j+1)}$ agrees with ϕ_j and ϕ_{j+1} on the seams. We will write α (resp. a) for the connection on Σ_\bullet (resp. Y_\bullet) that restricts to α_j on Σ_j ($a_{j(j+1)}$ on $Y_{j(j+1)}$). In a similar manner, we define ϕ, ψ , which are forms on Σ_\bullet , and p , which is a form on Y_\bullet . Then $a_{j(j+1)}$ can be chosen so that the data of α, ϕ, ψ, a, p patch together to define a smooth connection A on $\mathbb{R} \times Q$ with

$$A|_{\{s\} \times Y_\bullet} = a(s) + p(s) ds, \quad A|_{\{(s,t)\} \times \Sigma_\bullet} = \alpha(s,t) + \phi(s,t) ds + \psi(s,t) dt.$$

We will refer to such a connection A as a *holomorphic curve representative*.

To summarize, the boundary operator ∂_{symp} counts isolated $\mathcal{G}_0(P_\bullet)$ -equivalence classes of α satisfying

- ((J, \underline{H}) -holomorphic) $\begin{aligned} \partial_s \alpha - d_\alpha \phi + *(\partial_t \alpha - X_t(\alpha) - d_\alpha \psi)|_{\Sigma_\bullet} &= 0 \\ F_\alpha|_{\Sigma_\bullet} &= 0 \end{aligned} \quad (14)$
- (Lagrangian seam) $F_a|_{Y_\bullet} = 0$

with appropriate limits at $\pm\infty$, and where the ϕ, ψ and a are uniquely determined by these conditions.

3 Higher rank instanton Floer theory

Throughout this section, Y will denote a closed, connected, oriented 3-manifold equipped with a principal $\text{PU}(r)$ -bundle $Q \rightarrow Y$. We assume this bundle satisfies the following hypothesis.

(H1) There is an embedding $\iota : \Sigma \rightarrow Y$ of a closed, oriented surface Σ with the property that $d := t_2(Q)|_{\Sigma} \in \mathbb{Z}_r$ is a generator.

Such bundles exist if, for example, Y has positive first Betti number. Note also that (H1) implies that Y is necessarily *not* a homology S^3 . We also fix a subgroup $\mathcal{K} \subseteq \mathcal{G}(Q)$, which we assume consists of connected components of $\mathcal{G}(Q)$. In particular, this means that the identity component $\mathcal{G}_0(Q) \subseteq \mathcal{G}(Q)$ acts freely on \mathcal{K} by left multiplication. Hence $\mathcal{K}/\mathcal{G}_0(Q)$ is well-defined and is the group $\pi_0(\mathcal{K})$ of connected components of \mathcal{K} .

We begin by formulating instanton Floer cohomology, working modulo the gauge subgroup \mathcal{K} . To ensure we have a well-defined theory, we will make various hypotheses along the way. In Section 3.2 we restrict attention to specific subgroups \mathcal{K} and show that these hypotheses are satisfied in each case. In that section we also discuss how the various choices of \mathcal{K} determine various gradings. Section 3.3 specializes the discussion to the case where Y is a broken circle fibration.

3.1 Instanton Floer cohomology for $G = \text{PU}(r)$

Let \mathcal{K} be as above. We will use the same symbol to denote the Sobolev completion of \mathcal{K} in $\mathcal{G}^{k+1,p}(Q)$. Throughout we assume k, p are chosen so that $\mathcal{G}^{k+1,p}(Q)$ is a Lie group that acts smoothly on $\mathcal{A}^{k,p}(Q)$. Fix a smooth \mathcal{K} -invariant function $H : \mathcal{A}^{k,p}(Q) \rightarrow \mathbb{R}$. Then there is a smooth map $X = X_H$ from $\mathcal{A}^{k,p}(Q)$ into the $W^{k,p}$ -completion of $\Omega^2(Y, Q(\mathfrak{g}))$ that represents the differential of H in the sense that $(dH)_a(v) = \int_Y \langle X(a) \wedge v \rangle$ for all $v \in T_a \mathcal{A}^{k,p}(Q)$. By the invariance of H , we have the identity $X(u^*a) = \text{Ad}(u^{-1})X(a)$ for $u \in \mathcal{K}$. We assume $V := X_H$ satisfies conditions (a) and (b) from Theorem 2.5, and when this is the case we call H an *instanton perturbation*.

Next, fix a reference connection $a_0 \in \mathcal{A}_{\text{flat}}^{k,p}(Q)$ and define the *perturbed Chern-Simons functional* $\mathcal{CS}_{H,a_0} : \mathcal{A}^{k,p}(Q) \rightarrow \mathbb{R}$ by setting

$$\mathcal{CS}_{H,a_0}(a) := -H(a) + \frac{1}{2} \int_Y \langle d_{a_0} v \wedge v \rangle + \frac{1}{6} \int_Y \langle [v \wedge v] \wedge v \rangle.$$

where $v := a - a_0$. This map is smooth provided k, p are in an appropriate range for Sobolev multiplication, which we also assume throughout. For example, this is the case if $k = 1, p = 2$. From now on we will drop the Sobolev exponents from the notation, unless they are relevant.

The perturbed Chern-Simons functional only depends on a_0 up to an overall constant, so we will usually write $\mathcal{CS}_H := \mathcal{CS}_{H,a_0}$. The differential of \mathcal{CS}_H at a connection a satisfies

$$(d\mathcal{CS}_H)_a v = \int_Y \langle (F_a - X_H(a)) \wedge v \rangle$$

for all $v \in \Omega^1(Y, Q(\mathfrak{g}))$. It follows from (12) that there is an identity of the form

$$\mathcal{CS}_H(u^*a) - \mathcal{CS}_H(a) = 4\pi^2 r^{-1} \kappa_r \deg(u) \quad (15)$$

for all $a \in \mathcal{A}(Q)$ and $u \in \mathcal{K}$. (Our definition of degree is the negative of the one appearing in [8].) If H is invariant under a subgroup of $\mathcal{G}(Q)$ larger than \mathcal{K} , then (15) holds for all u in this larger subgroup.

Next, we define an instanton Floer cohomology group $HF_{\text{inst}}^\bullet(Q)^\mathcal{K}$. This can be viewed as the Morse cohomology, modulo gauge transformations in \mathcal{K} , of the Chern-Simons functional. That is, when it is defined, $HF_{\text{inst}}^\bullet(Q)^\mathcal{K}$ is the cohomology associated to a chain complex $(CF_{\text{inst}}^\bullet(Q), \partial_{\text{inst}})$. We describe this now.

The chain complex will be generated by \mathcal{K} -equivalence classes of the critical points of \mathcal{CS}_H . These critical points are precisely the *H-flat connections* $a \in \mathcal{A}(Q)$, which are defined by the condition $F_a = X(a)$. We denote the set of *H-flat connections* by $\mathcal{A}_{\text{flat}}(Q, H)$. Note that \mathcal{K} acts on this space, and we set

$$\mathcal{I}_H(Q) := \mathcal{A}_{\text{flat}}^{k,p}(Q, H) / \mathcal{K}.$$

This is independent of the choice of k, p . Given an *H-flat connection* a , we denote its image in $\mathcal{I}_H(Q)$ by $[a]$.

Example 3.1. Suppose $H = 0$. If $\mathcal{K} = \mathcal{G}_0(Q)$, then $\mathcal{I}_0(Q)$ is the moduli space $M(Q) = \mathcal{A}_{\text{flat}}(Q)/\mathcal{G}_0(Q)$ of flat connections. For general \mathcal{K} , $\mathcal{I}_0(Q)$ is a quotient of $M(Q)$ by the discrete group $\mathcal{K}/\mathcal{G}_0(Q)$.

Now we set

$$CF_{\text{inst}}^\bullet(Q) := \bigoplus_{[a] \in \mathcal{I}_H(Q)} \mathbb{Z}_2 \langle [a] \rangle.$$

The desirable cases are when all of the critical points $a \in \mathcal{A}_{\text{flat}}(Q, H)$ are *non-degenerate* in the sense that they are non-degenerate critical points of \mathcal{CS} in the Morse-theoretic sense, taken modulo gauge. Fixing a metric g on Y , this is equivalent to requiring that the extended Hessian

$$\mathcal{D}_a := \begin{pmatrix} *d_a - *dX_a & -d_a \\ -d_a^* & 0 \end{pmatrix} \quad (16)$$

is non-degenerate as an operator $W^{k,p} \rightarrow W^{k-1,p}$. Here, the short-hand $W^{j,p}$ denotes the obvious Sobolev completion of the vector space $\Omega^1(Y, Q(\mathfrak{g})) \oplus \Omega^0(Y, Q(\mathfrak{g}))$. There always exists a perturbation H such that all H -flat connections are non-degenerate [5, Chapter 5]. We discuss this further in Section 4.2. However, for the remainder of this section, we take this as a hypothesis.

(H2) The instanton perturbation H has been chosen so that all H -flat connections are non-degenerate.

It follows immediately from (H2) that all H -flat connections a are irreducible.

One consequence of Hypothesis (H2) is that the chain complex $CF_{\text{inst}}^\bullet(Q)$ admits a relative grading. To define this grading, set $n_{\mathcal{K}} := \inf \{ \deg(u) > 0 \mid u \in \mathcal{K} \}$. By convention, if all elements of \mathcal{K} have degree zero, then $n_{\mathcal{K}} = \infty$.

Proposition 3.2. *Assume (H1-2). Then there is a relative \mathbb{Z} -grading*

$$\mu_{\text{inst}} : \mathcal{A}_{\text{flat}}(Q, H) \times \mathcal{A}_{\text{flat}}(Q, H) \longrightarrow \mathbb{Z},$$

and this satisfies

$$\mu_{\text{inst}}(a_0, u^*a) - \mu_{\text{inst}}(a_0, a) = 4 \deg(u) \quad (17)$$

for all $a_0, a \in \mathcal{A}_{\text{flat}}(Q, H)$ and $u \in \mathcal{K}$. In particular, this induces a relative $\mathbb{Z}_{4n_{\mathcal{K}}}$ -grading on $\mathcal{I}_H(Q)$, and hence on $CF_{\text{inst}}^\bullet(Q)$, where $\mathbb{Z}_{4n_{\mathcal{K}}} := \mathbb{Z}$ if $n_{\mathcal{K}} = \infty$.

We defer the proof of this proposition to the end of this section. Using this relative $\mathbb{Z}_{4n_{\mathcal{K}}}$ -grading, we have a decomposition

$$CF_{\text{inst}}^\bullet(Q) = \bigoplus_{k \in \mathbb{Z}_{4n_{\mathcal{K}}}} CF_{\text{inst}}^k(Q), \quad CF_{\text{inst}}^k(Q) := \bigoplus \mathbb{Z}_2 \langle [a] \rangle,$$

where the sum on the right is over all $[a] \in \mathcal{I}_H(Q)$ with $\mu_{\text{inst}}([a_0], [a]) = k$, and $[a_0] \in \mathcal{I}_H(Q)$ is some fixed reference connection. It follows from Theorem 2.5 that each $CF_{\text{inst}}^k(Q)$ is finitely generated.

It is immediate from (17) that each gauge transformation with non-zero degree acts freely on $\mathcal{A}_{\text{flat}}(Q, H)$. However, it will be convenient to know that the subgroup \mathcal{K} acts freely on the space of H -flat connections. We take this as a hypothesis as well.

(H3) The subgroup \mathcal{K} acts freely on $\mathcal{A}_{\text{flat}}(Q, H)$.

The next step is to introduce the boundary operator

$$\partial_{\text{inst}} : CF_{\text{inst}}^k(Q) \longrightarrow CF_{\text{inst}}^{k+1}(Q).$$

It will take some time to develop the machinery necessary to define this explicitly. However, when we are done it will be given by a mod-2 count of isolated negative gradient trajectories of the perturbed Chern-Simons functional. These gradient trajectories are solutions $a : \mathbb{R} \rightarrow \mathcal{A}(Q)$ to the equation

$$\partial_s a = - * (F_a - X(a)). \quad (18)$$

This equation is plainly invariant under the action of \mathcal{K} . Note that we may alternatively view any path $s \mapsto a(s)$ of connections as defining a single connection $A = a(s)$ on the bundle $\mathbb{R} \times Q \rightarrow \mathbb{R} \times Y$. Conversely, every connection on $\mathbb{R} \times Q$ has the form $A = a(s) + p(s) ds$, where $a : \mathbb{R} \rightarrow \mathcal{A}(Q)$ and $p : \mathbb{R} \rightarrow \Omega^0(Y, \mathfrak{Q}(\mathfrak{g}))$. The curvature decomposes into components as $F_A = F_a + ds \wedge (\partial_s a - d_a p)$. We say that A is (g, H) -ASD or a (g, H) -instanton if it satisfies

$$\partial_s a - d_a p + *(F_a - X(a)) = 0. \quad (19)$$

When $H = 0$, equation (19) is exactly the ASD equation for the connection A (see Section 2.2; the Hodge star is the one on $\mathbb{R} \times Y$ coming from the product metric $ds^2 + g$). Our immediate interest in (19) is that it reduces to (18) when A is in temporal gauge $p = 0$. It follows that solutions to (18) modulo $\text{Maps}(\mathbb{R}, \mathcal{G}_0(Q))$ are identical to solutions of (19) modulo $\mathcal{G}_0(\mathbb{R} \times Q)$, since every connection A can be put into temporal gauge by an element of $\mathcal{G}_0(\mathbb{R} \times Q)$.

The following are equivalent for a smooth (g, H) -instanton $A = a(s) + p(s)$ on $\mathbb{R} \times Q$:

(i) The connection A has finite H -energy:

$$\mathcal{YM}_H(A) := \frac{1}{2} \int_{\mathbb{R} \times Y} \left| F_{a(s)} - X(a(s)) \right|^2 + \left| \partial_s a(s) - d_{a(s)} p(s) \right|^2 < \infty.$$

(ii) The s -derivative $\partial_s a(s) - d_{a(s)} p(s)$ decays exponentially to zero as s approaches $\pm\infty$:

$$\left\| \partial_s a(s) - d_{a(s)} p(s) \right\|_{L^2(Y)} \leq C e^{-\kappa|s|},$$

for some constants $C, \kappa > 0$.

(iii) The connection A converges to H -flat connections $a^\pm \in \mathcal{A}_{\text{flat}}(Q, H)$ at $\pm\infty$:

$$\lim_{s \rightarrow \pm\infty} a(s) = a^\pm, \quad \lim_{s \rightarrow \pm\infty} p(s) = 0;$$

here the convergence is in \mathcal{C}^∞ on Y .

When the last condition holds we will say that A *limits to* a^\pm at $\pm\infty$. The equivalence of (i) and (iii) above imply that the group of gauge transformation in \mathcal{K} act freely on the space of (g, H) -instantons. To state this precisely, we introduce the group

$$\text{Maps}_{\pm\infty}(\mathbb{R}, \mathcal{K}) \subseteq \text{Maps}(\mathbb{R}, \mathcal{G}(Q)) = \mathcal{G}(\mathbb{R} \times Q)$$

of smooth gauge transformation $U = u(\cdot)$ on $\mathbb{R} \times Q$ that converge at $\pm\infty$ to elements of \mathcal{K} . We assume the convergence at $\pm\infty$ is at least in \mathcal{C}^1 on Y .

Proposition 3.3. *Assume (H1-3). Let $A \in \mathcal{A}(\mathbb{R} \times Q)$ be a smooth (g, H) -instanton on $\mathbb{R} \times Q$ that has finite H -energy. Suppose $U \in \text{Maps}_{\pm\infty}(\mathbb{R}, \mathcal{K})$ is a smooth gauge transformation such that $U^*A = A$. Then $U = e$ is the identity.*

Proof. To prove the proposition, it suffices to prove the following:

- (i) Assume (H1). Fix $a_\infty \in \mathcal{A}(Q)$ and suppose this is irreducible. If $A \in \mathcal{A}(\mathbb{R} \times Q)$ is any connection limiting to a_∞ at ∞ (or $-\infty$), then A is also irreducible.
- (ii) Suppose H satisfies (H2), and \mathcal{K} satisfies (H3). Let a_∞ and A be as in (i). If a_∞ is H -flat, then the stabilizer of A in $\text{Maps}_{\pm\infty}(\mathbb{R}, \mathcal{K})$ is trivial.

Note that A is irreducible if and only if there is some $U \in \mathcal{G}(\mathbb{R} \times Q)$ such that U^*A is irreducible. In particular, we may assume $A = a(\cdot)$ is in temporal gauge. First we prove (i). Suppose $d_A R = 0$ for some 0-form R on $\mathbb{R} \times Y$. We want to show $R = 0$. Write $R = r(\cdot)$ as a path of 0-forms on Y . Then writing $d_A R = 0$ in components gives

$$\partial_s r(s) = 0, \quad d_{a(s)} r(s) = 0$$

for each $s \in \mathbb{R}$. The first condition says $r(s) = r_0$ is constant. Taking the limit as s approaches ∞ implies that $d_{a_\infty} r_0 = 0$ and so $r_0 = 0$, as desired.

To prove (ii), let $U = u(\cdot) \in \text{Maps}_{\pm\infty}(\mathbb{R}, \mathcal{K})$ and suppose this fixes $A = a(\cdot)$. This implies $u(s)$ fixes $a(s)$ for each $s \in \mathbb{R}$. By the assumptions on U , the limit $\lim_{s \rightarrow \infty} u(s) = u_\infty$ exists, and it is immediate that u_∞ fixes a_∞ . Since a_∞ is an H -flat connection, it follows from (H3) that $u_\infty = e$. In particular, $u(s) \in \mathcal{G}_0(Q)$ lies in the identity component for all s . By Lemma 3.7 below, it follows that $u(s) = e \in \mathcal{G}(Q)$ is the identity for all sufficiently large s .

To prove that $u(s) = e$ for all $s \in \mathbb{R}$, consider the set of $s \in \mathbb{R}$ such that $u(s) = e$. This set is clearly closed, and we just saw that it is non-empty. It suffices to show this set is also open. It follows from (i) that $a(s)$ is irreducible

for each $s \in \mathbb{R}$. In general, a connection $a \in \mathcal{A}^{k,p}(Q)$ is irreducible if and only if it satisfies a bound of the form

$$\|r\|_{L^p(Y)} \leq C \|d_a r\|_{L^p(Y)} \quad (20)$$

for all 0-forms $r \in \Omega^0(Y, Q(\mathfrak{g}))$; this bound continues to hold for all connections near an irreducible a . Suppose there is some $s_0 \in \mathbb{R}$ such that $u(s_0) = e$. We want to show $u(s) = e$ for all s near s_0 . Using the exponential map for the gauge group, we can write $u(s) = \exp(r(s))$, at least when $|s - s_0|$ is small. Here, $r(s)$ is a path of 0-forms on Y with values in $Q(\mathfrak{g})$. Now differentiate the identity $u(s)^* a(s) = a(s)$ to get $d_{a(s)} r(s) = 0$. By (20), this implies $r(s) = 0$, and hence $u(s) = e$, for all s sufficiently close to s_0 . \square

Fix two H -flat connections $a^\pm \in \mathcal{A}_{\text{flat}}(Q, H)$. Our immediate goal now is to define the moduli space $\mathcal{M}([a^-], [a^+])$ of instanton trajectories from $[a^-]$ to $[a^+]$. This is the primary geometric object used to define the Floer boundary operator ∂_{inst} . Intuitively, $\mathcal{M}([a^-], [a^+])$ can be viewed as the space

$$\left\{ A \in \mathcal{A}(\mathbb{R} \times Q) \mid (19), \lim_{s \rightarrow \pm\infty} A|_{\{s\} \times Y} \in [a^\pm] \right\} / \text{Maps}_{\pm\infty}(\mathbb{R}, \mathcal{K}) \quad (21)$$

where $\mathcal{A}(\mathbb{R} \times Q)$ denotes the space of *smooth* connections, and $\text{Maps}_{\pm\infty}(\mathbb{R}, \mathcal{K})$ is as defined above Proposition 3.3. However, to specify a smooth structure on the space $\mathcal{M}([a^-], [a^+])$ we need to work with suitable function theoretic completions of these spaces. Due to the non-compactness of $\mathbb{R} \times Y$, this is not as straightforward as one might hope. As a consequence, we adopt the following more circuitous approach.

Fix representatives $a^\pm \in [a^\pm]$. Also fix representatives $\mathcal{U} = \{u_0, u_1, \dots\} \subset \mathcal{K}$, one for each element of $\mathcal{K}/\mathcal{G}_0(Q)$. That is, \mathcal{U} is chosen so that the inclusion map $\mathcal{U} \hookrightarrow \mathcal{K}$ descends to a bijection

$$\mathcal{U} \cong \mathcal{K}/\mathcal{G}_0(Q). \quad (22)$$

The bijection (22) imposes a group structure on \mathcal{U} . (The gauge theoretic product of two elements of \mathcal{U} may not be in \mathcal{U} , but this product is always *homotopic* to a unique element of \mathcal{U} ; this is the group structure determined by (22).) Next, recall the Banach manifolds $\mathcal{A}^{k,p}(a, a') \subset \mathcal{A}^{k,p}(\mathbb{R} \times Q)$ and $\mathcal{G}_c^{k+1,p}(\mathbb{R} \times Q) \subset \mathcal{G}_{\text{loc}}^{k+1,p}(\mathbb{R} \times Q)$ from (10) and (11), respectively. For $a, a' \in \mathcal{A}_{\text{flat}}(Q, H)$, define the following auxiliary space

$$\mathcal{N}(a, a') := \left\{ A \in \mathcal{A}^{k,p}(a, a') \mid (19) \right\} / \mathcal{G}_c^{k+1,p}(\mathbb{R} \times Q).$$

We will denote the elements of the space by $[A]_{\mathcal{G}_c}$. Then the group \mathcal{U} acts on the disjoint union

$$\bigsqcup_{u_i, u_j \in \mathcal{U}} \mathcal{N}(u_i^* a^-, u_j^* a^+)$$

via the isomorphism (22), where we are viewing elements of \mathcal{U} as gauge transformations on $\mathbb{R} \times Q$ that are constant in the \mathbb{R} -direction. Moreover, this action is free by Proposition 3.3. Finally, we set

$$\mathcal{M}([a^-], [a^+]) := \left(\bigsqcup_{u_i, u_j \in \mathcal{U}} \mathcal{N}(u_i^* a^-, u_j^* a^+) \right) / \mathcal{U}.$$

At this point, there are three matters that need to be addressed:

- (i) The dependence of $\mathcal{M}([a^-], [a^+])$ on the choices of representatives a^\pm and \mathcal{U} ;
- (ii) The sense in which $\mathcal{M}([a^-], [a^+])$ is a smooth manifold.

We begin with (i). Suppose $b^\pm \in [a^\pm]$ are two other choices of representatives. Likewise, suppose $\mathcal{V} = \{v_k\}_k$ is another set of representatives of $\mathcal{K}/\mathcal{G}_0(Q)$. We want to show that the spaces

$$\left(\bigsqcup_{i,j} \mathcal{N}(u_i^* a^-, u_j^* a^+) \right) / \mathcal{U}, \quad \text{and} \quad \left(\bigsqcup_{k,l} \mathcal{N}(v_k^* b^-, v_l^* b^+) \right) / \mathcal{V}$$

are canonically identified. By assumption, we can write $b^\pm = u_\pm^* a^\pm$ for some $u_\pm \in \mathcal{K}$. Note that the dependence of the space $\bigsqcup_{i,j} \mathcal{N}(v_k^* b^-, v_l^* b^+)$ on the representatives is just through the set $\{v_k^* b^\pm\}_k$. Since $v_k^* b^- = (u_- v_k)^* a^-$, by replacing v_k with $u_- v_k$, we may suppose $u_- = e$ is the identity. Using the actions of \mathcal{U} and \mathcal{V} , it suffices to show that, for each j , there is some l for which the spaces $\mathcal{N}(a^-, u_j^* a^+)$, and $\mathcal{N}(a^-, v_l^* b^+)$ are canonically identified. To see this, suppose we are given some $u_j \in \mathcal{U}$. Then there is a unique $v_l \in \mathcal{V}$ such that u_j and $u_+ v_l$ are in the same component of \mathcal{K} . Let $w : \mathbb{R} \rightarrow \mathcal{G}_0(Q)$ be a path with $w(s) = e$ for $s < -1$ and $w(s) = u_j^{-1} u_+ v_l$ for $s > 1$. This determines a gauge transformation W on $\mathbb{R} \times Q$ in the obvious way, and hence a bijection $\mathcal{N}(a^-, u_j^* a^+) \cong \mathcal{N}(a^-, v_l^* b^+)$ defined by sending $[A]_{\mathcal{G}_c}$ to $[W^* A]_{\mathcal{G}_c}$. This map is independent of the choice of path w since any two choices differ by an element of $\mathcal{G}_c^{k+1,p}(\mathbb{R} \times Q)$, and this group acts by the identity. This resolves (i).

Now we discuss (ii). By assumption, H has been chosen so all H -flat connections are non-degenerate, and it follows that the defining equations for each $\mathcal{N}(u_i^* a^-, u_j^* a^+)$ are Fredholm. Indeed, their linearization at $[A]_{\mathcal{G}_c} \in \mathcal{N}(u_i^* a^-, u_j^* a^+)$ is an operator $\mathcal{D}_{A,g,H}$ obtained by using H to perturb the Fredholm operator $d_A^+ \oplus d_A^*$. This perturbation determined by H is compact (at least when restricted to compact subsets of $\mathbb{R} \times Y$), so many of the Fredholm properties of the operators $\mathcal{D}_{A,g,H}$ and $d_A^+ \oplus d_A^*$ are the same. For example, they have the same index. (This relies on the fact that we considered the instanton equations on the space $\mathcal{A}^{k,p}(a, a')$, rather than, say, the much larger space $\mathcal{A}_{loc}^{k,p}(\mathbb{R} \times Q)$.)

When $\mathcal{D}_{A,g,H}$ is onto, there is a neighborhood of $[A]_{\mathcal{G}_c}$ in $\mathcal{N}(u_i^*a^-, u_j^*a^+)$ that is a smooth manifold of dimension $\text{Ind}(\mathcal{D}_{A,g,H}) = \mu_{\text{inst}}(u_i^*a^-, u_j^*a^+)$; see the proof of Proposition 3.2. We say that the pair (g, H) is *regular (for instanton Floer theory)*, if all H -flat connections are non-degenerate, and if for all i, j , all $[a^\pm] \in \mathcal{I}_H(Q)$, and all $[A]_{\mathcal{G}_c} \in \mathcal{N}(u_i^*a^-, u_j^*a^+)$, the operator $\mathcal{D}_{A,g,H}$ is onto. In Section 4.2 we will show that, given any metric g , there always exists some instanton perturbation H so that (g, H) is regular. For the remainder of this section we assume that this is the case.

(H4) The pair (g, H) is regular for instanton Floer theory.

Assuming (H4), it follows that $\sqcup_{i,j} \mathcal{N}(u_i^*a^-, u_j^*a^+)$ is a smooth manifold. By Proposition 3.3, the group $\mathcal{U} = \mathcal{K}/\mathcal{G}_0(Q)$ acts freely on this space and so the quotient $\mathcal{M}([a^-], [a^+])$ is indeed a smooth manifold with local dimension given by the mod $4n_{\mathcal{K}}$ -reduction of $\mu_{\text{inst}}(a^-, a^+)$. This resolves (ii).

The space $\mathcal{M}([a^-], [a^+])$ admits a free action of \mathbb{R} by translation, and we set

$$\widehat{\mathcal{M}}([a^-], [a^+]) := \mathcal{M}([a^-], [a^+])/\mathbb{R}.$$

These are the (*unparametrized*) *instanton trajectories*. We denote by $\widehat{\mathcal{M}}_0([a^-], [a^+])$ the zero-dimensional component of $\widehat{\mathcal{M}}([a^-], [a^+])$; this may be empty. Theorem 2.5 implies that $\widehat{\mathcal{M}}_0([a^-], [a^+])$ is a finite set. We define $\#_Q(a^-, a^+)$ to be the mod-2 count of its elements. Finally, we define $\partial_{\text{inst}} : CF_{\text{inst}}^\bullet(Q) \rightarrow CF_{\text{inst}}^\bullet(Q)$ to be the degree 1 operator given by

$$\partial_{\text{inst}}\langle [a^-] \rangle = \sum \#_Q(a^-, a^+) \langle [a^+] \rangle,$$

where the sum is over all $[a^+] \in \mathcal{I}_H(Q)$ with $\mu_{\text{inst}}(a^-, a^+) = 1$. Now we can state Floer's main theorem.

Theorem 3.4. *Assume (H1-4). Then $\partial_{\text{inst}}^2 = 0$, and so*

$$HF_{\text{inst}}^\bullet(Q)^{\mathcal{K}} := \frac{\ker \partial_{\text{inst}}}{\text{im } \partial_{\text{inst}}}$$

is well-defined. This abelian group inherits a relative $\mathbb{Z}_{4n_{\mathcal{K}}}$ -grading from the grading on $CF_{\text{inst}}^\bullet(Q)$. Furthermore, $HF_{\text{inst}}^\bullet(Q)^{\mathcal{K}}$ is independent of the choice of regular pair (g, H) , up to isomorphism of relatively $\mathbb{Z}_{4n_{\mathcal{K}}}$ -graded abelian groups.

The proof of Theorem 3.4 follows essentially as in [14], where Floer considers the trivial $SU(2)$ -bundle over a homology 3-sphere. The key technical device needed to carry Floer's proof to our setting are hypotheses (H1) and (H3), which essentially say that there are no reducible flat connections for our bundles (in Floer's setting, the only reducible flat connection is the trivial connection, which can be easily avoided). See also [15, 7] and [5, Chapter 5].

Proof of Proposition 3.2. This proposition is well-known in the case $\text{PU}(2) = \text{SO}(3)$ (see [15, 8, 4]), and essentially the same proof carries over to the more general case of $\text{PU}(r)$. We include the details for convenience. We note that the perturbation H only serves to ensure that the H -flat connections are non-degenerate. To simplify notation we assume $H = 0$ (see also Remark 3.5).

Let \mathcal{D}_A be the operator

$$\mathcal{D}_A = d_A^+ \oplus d_A^* : \Omega^1(\mathbb{R} \times Q, Q(\mathfrak{g})) \longrightarrow \Omega^+(\mathbb{R} \times Q, Q(\mathfrak{g})) \oplus \Omega^0(\mathbb{R} \times Q, Q(\mathfrak{g})).$$

Here $A \in \mathcal{A}(\mathbb{R} \times Q)$ is any connection that converges to non-degenerate flat connections a^\pm at $\pm\infty$, and d_A^+ is the composition

$$d_A^+ : \Omega^1 \xrightarrow{d_A} \Omega^2 \longrightarrow \Omega^+,$$

where the second arrow is the L^2 -orthogonal projection to the space of anti-self dual 2-forms. (It is important here that A is a connection on a pull-back bundle $\mathbb{R} \times Q$ and not on an arbitrary bundle over $\mathbb{R} \times Y$.) The operator \mathcal{D}_A is the linearization of the ASD operator $A \mapsto F_A^+$ coupled with the Coulomb gauge fixing operator $A \mapsto d_A^*(A_0 - A)$. The assumption on the non-degeneracy of the a^\pm implies that, after passing to suitable Banach space completions, \mathcal{D}_A becomes a Fredholm operator. In particular, its index $\text{Ind}(\mathcal{D}_A, a^-, a^+)$ is well-defined.

Claim: The index $\text{Ind}(\mathcal{D}_A, a^-, a^+)$ is independent of the choice of connection A on $\mathbb{R} \times Q$, as well as the choice of metric on Y used to define the splittings above.

We prove independence of the choice of connection; independence of the choice of metric is similar. As a first case, suppose that A and B are connections that agree off of a compact set $K \subset \mathbb{R} \times Y$. Then the difference $\mathcal{D}_A - \mathcal{D}_B$ is an operator given by multiplication by a section of a bundle over $\mathbb{R} \times Y$, and this section has support in K . It follows from the compact embedding $W^{1,2}(K) \hookrightarrow L^2(K)$ that this difference is a compact operator. The index of a Fredholm operator is unchanged under compact perturbations, so $\text{Ind}(\mathcal{D}_A, a^-, a^+) = \text{Ind}(\mathcal{D}_B, a^-, a^+)$ in this case.

More generally, if A and B do not agree on any compact set, then we can reduce to the previous case as follows. Since these connections have the same limits at $\pm\infty$, given any $\epsilon > 0$ one can find a compactly supported 1-form V so that $\mathcal{D}_A - \mathcal{D}_{B+V}$ has norm less than ϵ (multiply the difference $A - B$ by a bump function with suitably large support). The previous case shows that the operators \mathcal{D}_{B+V} and \mathcal{D}_B have the same index. To see that \mathcal{D}_{B+V} has the same index as \mathcal{D}_A , use the fact that if $F : X \rightarrow Y$ is a Fredholm operator, then there is some $\epsilon > 0$ such that if $L : X \rightarrow Y$ is any bounded linear operator with norm less than ϵ , then $F + L$ is Fredholm and has the same index as F . This proves the claim.

Set

$$\mu_{\text{inst}}(a^-, a^+) := \text{Ind}(\mathcal{D}_A, a^-, a^+).$$

Then μ_{inst} is additive in each component, and so is in fact a relative \mathbb{Z} -grading. It therefore suffices to show $\text{Ind}(\mathcal{D}_A, a, u^*a) = 4 \deg(u)$, where $a \in \mathcal{A}_{\text{flat}}(Q)$ is a flat connection, $u \in \mathcal{K}$ is a gauge transformation, and A limits to a and u^*a .

Remark 3.5. (a) Recall we have only assumed that the perturbation H is invariant under the group \mathcal{K} , and not necessarily under the full group $\mathcal{G}(Q)$. In particular, if u is a gauge transformation not in \mathcal{K} , it could be the case that u^*a is not H -flat, even if $a \in \mathcal{A}_{\text{flat}}(Q, H)$. Hence, the extended Hessian at u^*a may be degenerate and this could result in the failure of \mathcal{D}_A to be Fredholm. This being said, the proof we give below will show that $\text{Ind}(\mathcal{D}_A, a, u^*a) = 4 \deg(u)$ for any u in the larger group $\mathcal{G}(Q)$, provided one knows that the operator \mathcal{D}_A is Fredholm when A limits to a and u^*a .

(b) In general, suppose G is a simple Lie group and consider a G -bundle $R \rightarrow X$ over a closed 4-manifold. Then the index of the linearized ASD operator at a connection B on R can be computed using the Atiyah-Singer index formula to give

$$\text{Ind}(\mathcal{D}_B) = c(G)\kappa(R) - \dim(G) (1 - b_1(X) + b^+(X)). \quad (23)$$

See, for example, [6, Equation 7.1.3]. Here $\kappa(R)$ is a characteristic number for R , and $c(G)$ is a normalizing constant depending only on G . For example, when $G = \text{SU}(r)$, one typically takes $\kappa(R) = c_2(R) [X]$ to be the 2nd Chern number. Then one can show $c(\text{SU}(r)) = 4r$.

When $G = \text{PU}(r)$, we will have $\kappa(R) = q_4(R) [X]$ is the generalized Pontryagin number from Section 2.3. That this is the correct choice of κ can be seen as follows: The connection B can be equivalently viewed as a connection on the complexified adjoint bundle $R(\mathfrak{g})_{\mathbb{C}}$. The structure group of this bundle reduces to $\text{SU}(r)$ since $c_1 = 0$ for complexified bundles, so by the previous paragraph we have $\kappa(R(\mathfrak{g})_{\mathbb{C}}) = c_2(R(\mathfrak{g})_{\mathbb{C}}) [X] = q_4(R) [X]$, where the second equality is the definition of q_4 . We will compute $c(\text{PU}(r))$ for this choice of characteristic number in a moment (essentially the same argument can be used to show $c(\text{SU}(r)) = 4r$).

We also note that when B is an irreducible ASD connection, the value (23) recovers the dimension near B of the moduli space of irreducible ASD connections on R .

Now we return to the situation where A is a connection on $\mathbb{R} \times Q$ with flat limits a and u^*a . We may view A as being a connection on the bundle $Q_u := I \times Q / (0, u(q)) \sim (1, q)$ over the closed manifold $S^1 \times Y$. Apply Remark 3.5 (b) to the case

$$G = \text{PU}(r), \quad X = S^1 \times Y, \quad R = Q_u, \quad \text{and} \quad B = A.$$

Then by (23) and the definition of the degree, we have

$$\text{Ind}(\mathcal{D}_A) = c(\text{PU}(r))\kappa(Q_u) = 2c(\text{PU}(r)) \deg(u)$$

since the term involving the Betti numbers vanishes for $X = S^1 \times Y$. We will be done if we can show $c(\text{PU}(r)) = 2$.

To do this, we use (23) again, but with $X = S^4$. Before specifying the $\text{PU}(r)$ -bundle R , we first note that the $\text{SU}(r)$ -bundles over S^4 are classified by their

second Chern class. Let R' denote the $SU(r)$ -bundle with $c_2(R') [S^4] = 2r$. Then we take $R \rightarrow S^4$ to be the adjoint bundle associated to R' . So R is a $PU(r)$ -bundle with $q_4(R) := c_2(\text{End}(R')) = 2rc_2(R') = 4r^2$. Atiyah-Hitchin-Singer show that R' always admits an irreducible ASD connection B [3, Theorem 8.4]. Moreover, in [3, Table 8.1], the authors compute the dimension of the moduli space of irreducible ASD connections on the $SU(r)$ -bundle R' to be $2(4r^2) - r^2 + 1$. This moduli space has the same dimension as the moduli space of irreducible ASD connections on the associated $PU(r)$ -bundle R , and so it recovers the index. Combining this with (23) gives

$$2(4r^2) - r^2 + 1 = \text{Ind}(\mathcal{D}_B) = c(\text{PU}(r))4r^2 - r^2 + 1$$

and so it follows that $c(\text{PU}(r)) = 2$. \square

3.2 Gradings

We continue to assume $Q \rightarrow Y$ is a bundle satisfying hypothesis (H1). In the previous section we defined a cohomology group $HF_{\text{inst}}^\bullet(Q)^\mathcal{K}$ associated to a subgroup $\mathcal{K} \subseteq \mathcal{G}(Q)$ satisfying hypothesis (H3). In this section we work with the specific cases where \mathcal{K} is $\mathcal{G}_0(Q)$, $\ker \eta$ and $\mathcal{G}(\Sigma)$; here $\eta : \mathcal{G}(Q) \rightarrow H^1(Y, \mathbb{Z}_r)$ is the parity operator, and $\mathcal{G}(\Sigma)$ is the subgroup generated by $\mathcal{G}_0(Q)$ and the degree 1 gauge transformation from Proposition 2.10 (b). In each case we will show that \mathcal{K} satisfies hypothesis (H3), and that we obtain a relatively $\mathbb{Z}, \mathbb{Z}_{4r}$ and \mathbb{Z}_4 -graded Floer theory in the respective cases (the latter case recovers the \mathbb{Z}_4 -graded Floer group that is more common in the literature [15, 4, 8]). Along the way, we also compare the associated Floer cohomology groups in the different cases, and show that they contain the same information. Throughout we assume g and H are chosen so hypotheses (H2) and (H4) hold.

3.2.1 A \mathbb{Z} -grading

Here we show that $\mathcal{K} = \mathcal{G}_0(Q)$ acts freely on the space $\mathcal{A}_{\text{flat}}(Q, H)$. Once we have done this, it follows immediately that the induced grading on $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$ is a relative \mathbb{Z} -grading, since each element of $\mathcal{G}_0(Q)$ has degree zero.

It follows from (H2) that every $a \in \mathcal{A}_{\text{flat}}(Q, H)$ is irreducible, and so (20) holds for all $r \in \Omega^0(Y, Q(\mathfrak{g}))$. As a preliminary step, we prove the following.

Lemma 3.6. *Assume (H2) and let $p > 3$. Then there are constants $\epsilon_0, C > 0$ such that (20) holds for all $r \in \Omega^0(Y, Q(\mathfrak{g}))$ and for all connections a on $Q \rightarrow Y$ with $\|F_a - X(a)\|_{L^p(Y)} \leq \epsilon_0$.*

Proof. The result obviously holds for H -flat connections a , with a constant $C = C_a$ possibly depending on a . We first show that the constant C can actually be chosen to be independent of the H -flat connection a . If not, then one could find sequences $\{r_n\}_n$ of 0-forms and $\{a_n\}_n$ of H -flat connections with

$$\|r_n\|_{L^2} = 1, \quad \|d_{a_n} r_n\|_{L^2} \rightarrow 0. \quad (24)$$

The bound (20) is gauge invariant, so by the perturbed version of Uhlenbeck compactness, we may assume that the a_n converge strongly in L^∞ to an H -flat connection a_∞ , after possibly passing to a subsequence. Moreover, this convergence of the a_n combines with the assumptions on the r_n to imply that the r_n are uniformly bounded in $W^{1,2}$ (with the derivatives defined using d_{a_∞}). So by passing to a further subsequence, we may assume the r_n converge weakly in $W^{1,2}$ and hence strongly in L^2 to some r_∞ . This gives $1 = \|r_\infty\|_{L^2}$, and $d_{a_\infty}r_\infty = 0$, which contradicts (20) applied to the flat connection $a = a_\infty$, and so the constant C is independent of the choice of H -flat connection.

Next, we show that the estimate (20) continues to hold for any connection a with $F_a - X_a$ sufficiently L^p -small. This follows easily by the same kind of contradiction argument: If not, then there are a_n and r_n with (24), and $\|F_{a_n} - X(a_n)\|_{L^p} \rightarrow 0$. Then $\{a_n\}_n$ has a subsequence that converges strongly in L^∞ , modulo gauge equivalence, to some limiting H -flat connection, and the result follows exactly as before. (We use $p > 3$ to obtain a compact embedding $W^{1,p} \hookrightarrow L^\infty$.) \square

Hypothesis (H3) for the group $\mathcal{K} = \mathcal{G}_0(Q)$ is a special case of the following.

Lemma 3.7. *Assume (H2) and let $p > 3$. There is some $\epsilon_0 > 0$ such that if a is any connection on $Q \rightarrow Y$ with*

$$\|F_a - X(a)\|_{L^p(Y)} < \epsilon_0, \quad (25)$$

then the stabilizer of a in $\mathcal{G}_0(Q)$ is trivial. Moreover, the set $\mathcal{I}_H(Q)$ is naturally equipped with the structure of a smooth manifold of dimension zero, and it is non-compact whenever it is not empty.

Proof. First we show how the second assertion follows from the first. The curvature is continuous (in fact, smooth) when viewed as a map $\mathcal{A}^{1,p}(Q) \rightarrow L^p(\Omega^2)$, where $L^p(\Omega^2)$ is the L^p -completion of the space $\Omega^2(Y, Q(\mathfrak{g}))$. In particular, the set \mathcal{U} of connections a satisfying (25) is a $\mathcal{G}_0^{2,p}(Q)$ -invariant open set. It follows from the first assertion of the lemma that the set $\mathcal{B} := \mathcal{U}/\mathcal{G}_0^{2,p}$ is naturally a smooth Banach manifold. The tangent space at a $[a] \in \mathcal{B}$ is the $W^{1,p}$ -completion of the space

$$\Omega^1(Y, Q(\mathfrak{g}))/\text{im } d_a. \quad (26)$$

We are interested in the Banach vector bundle $\mathcal{E} \rightarrow \mathcal{B}$ whose fiber over $[a]$ is the L^p -completion of (26). Then the map $a \mapsto *(F_a - X(a))$ descends to a smooth section of \mathcal{E} that is Fredholm with index 0. Moreover, the critical points of this section are exactly the elements of $\mathcal{I}_H(Q)$. By hypothesis (H2) these critical points are all non-degenerate, so by the inverse function theorem $\mathcal{I}_H(Q)$ is a smooth 0-dimensional manifold. The non-compactness follows, for example, from the fact that there is a degree $d \neq 0$ gauge transformation u . Then, by (17), the k -fold iterated action of u on an H -flat connection a yields distinct H -flat connections.

It remains to prove the first assertion in Lemma 3.7. By Lemma 3.6, there is some neighborhood $\mathcal{U}(e)$ of the identity in $\mathcal{G}_0(Q)$ that acts freely on the space of all connections with sufficiently small curvature. We want to show that this neighborhood is all of $\mathcal{G}_0(Q)$. This is another contradiction argument. If it does not hold, then there is a sequence of connections a_n on Q and gauge transformations u_n on Q such that $u_n^* a_n = a_n$ and $\|F_{a_n} - X(a_n)\|_{L^p} \rightarrow 0$. Apply Uhlenbeck's weak compactness theorem to conclude that, after passing to a subsequence, there is a sequence of gauge transformations u'_n such that $(u'_n)^* a_n$ converges weakly in $W^{1,p}$ to some flat connection a_∞ . For simplicity, we relabel $(u'_n)^* a_n$ by a_n . Similarly, replace $u'_n u_n (u'_n)^{-1}$ with u_n , so we still have $u_n^* a_n = a_n$. This relation combines with (8) to give

$$a_n u_n - u_n a_n = du_n.$$

Since the left-hand side is bounded in $W^{1,p}$, it follows that the right is as well. Thus the u_n converge weakly in $W^{2,p}$ to some limiting gauge transformation u_∞ . This satisfies $u_\infty^* a_\infty = a_\infty$. It follows from Lemma 2.9 that $u_\infty = e$ is the identity. To reiterate, the u_n are a sequence of gauge transformation converging to the identity e , and each fixes a connection with small curvature. This contradicts the existence of the neighborhood $\mathcal{U}(e)$ containing the identity. \square

3.2.2 A \mathbb{Z}_{4r} -grading

Here we consider the subgroup $\mathcal{K} = \ker \eta \subseteq \mathcal{G}(Q)$ given by the kernel of the parity operator $\eta : \mathcal{G}(Q) \rightarrow H^1(Y, \mathbb{Z}_r)$. First we want to show that this group acts freely on the space of H -flat connections. To see this, note that any gauge transformation in $\ker \eta$ with non-zero degree automatically acts freely on $\mathcal{A}_{\text{flat}}(Q, H)$ by (17). So it suffices to work with the degree zero elements of $\ker \eta$. However, any gauge transformation u with $\deg(u) = 0$ and $\eta(u) = 0$ is automatically in the identity component $\mathcal{G}_0(Q)$, so the result follows by the analysis of Section 3.2.1.

We saw in Section 2.3 that each element of $\ker \eta$ has degree divisible by r . Moreover, there always exists gauge transformations in $\ker \eta$ of degree r (take u_1^r , where u_1 is the degree 1 gauge transformation from Proposition 2.10 (b)). In particular, the induced grading on $HF_{\text{inst}}^\bullet(Q)^{\ker \eta}$ is a relative \mathbb{Z}_{4r} -grading.

It is interesting to compare the groups $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$ and $HF_{\text{inst}}^\bullet(Q)^{\ker \eta}$. Fix a degree r gauge transformation u_r , and assume that the instanton perturbation H is invariant under the action of u_r . Then pullback by u_r determines a degree $4r$ chain map on the underlying chain complex of $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$. This therefore descends to a degree $4r$ map at the level of homology

$$\iota_{u_r} : HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)} \longrightarrow HF_{\text{inst}}^{\bullet+4r}(Q)^{\mathcal{G}_0(Q)},$$

Moreover, this map is an isomorphism. It follows that all of the data of $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$ is contained in any $4r$ consecutive degrees

$$HF_{\text{inst}}^k(Q) \oplus HF_{\text{inst}}^{k+1}(Q) \oplus \dots \oplus HF_{\text{inst}}^{k+4r-1}(Q).$$

In fact, we can say more. Quotienting $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$ by the action of the map ι_{u_r} recovers the space $HF_{\text{inst}}^\bullet(Q)^{\ker \eta}$

$$HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)} \longrightarrow HF_{\text{inst}}^\bullet(Q)^{\ker \eta}. \quad (27)$$

Fix the k th summand $HF_{\text{inst}}^k(Q)^{\ker \eta}$ of $HF_{\text{inst}}^\bullet(Q)^{\ker \eta}$. The fiber of the quotient map (27) over this summand consists of the countably many summands

$$\dots \oplus HF_{\text{inst}}^{k-4r}(Q)^{\mathcal{G}_0(Q)} \oplus HF_{\text{inst}}^k(Q)^{\mathcal{G}_0(Q)} \oplus HF_{\text{inst}}^{k+4r}(Q)^{\mathcal{G}_0(Q)} \oplus \dots$$

in $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$. Moreover, the quotient map restricts to an isomorphism

$$HF_{\text{inst}}^k(Q)^{\mathcal{G}_0(Q)} \longrightarrow HF_{\text{inst}}^k(Q)^{\ker \eta}$$

for each $k \in \mathbb{Z}$. That is, all of the information of $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$ is contained in $HF_{\text{inst}}^\bullet(Q)^{\ker \eta}$, and vice-versa.

Lastly, we discuss a sense in which the gauge group $\ker \eta$ is a very natural group to consider. Suppose $t_2(Q)$ is the reduction of an integral class. Then there is a $U(r)$ -bundle $\bar{Q} \rightarrow Y$ with $Q = \bar{Q} \times_{U(r)} \text{PU}(r)$ and $c_1(\bar{Q}) \equiv t_2(Q) \pmod{r}$. In [5] and [4], the authors define a $U(2)$ -Floer theory, and their discussion generalizes easily to the group $U(r)$. Their $U(r)$ -Floer chain complex is generated by central curvature connections with fixed determinant, and one works modulo the group of gauge transformations on \bar{Q} with fixed determinant. It turns out that the $U(r)$ -Floer cohomology they describe is exactly the group $HF_{\text{inst}}^\bullet(Q)^{\ker \eta}$ defined here. Indeed, the distinction between their set-up and ours is really just a matter of language. For example, flat connections on Q are identified with central curvature connections on \bar{Q} with fixed determinant. Likewise, the subgroup $\ker \eta \subset \mathcal{G}(Q)$ is identified with the gauge transformations on \bar{Q} that have trivial determinant.¹

3.2.3 A \mathbb{Z}_4 -grading

Let $u_1 \in \mathcal{G}(Q)$ be a degree 1 gauge transformation satisfying the conclusion of Proposition 2.10 (b). Note that the homotopy class of u_1 is uniquely determined by the conditions of the proposition, and hence uniquely determined by the choice of surface $\Sigma \subset Y$ from that proposition. Let $\mathcal{G}(\Sigma) \subseteq \mathcal{G}(Q)$ denote the group generated by u_1 and the identity component $\mathcal{G}_0(Q)$.

Example 3.8. Suppose $d \in \mathbb{Z}_r$ is a generator, and Q is a bundle that is d -compatible with a broken circle fibration $Y \rightarrow S^1$, as in Section 2.4. Take $\Sigma = \Sigma_0$. In that section we saw that $\mathcal{G}(\Sigma_0) = \mathcal{G}_{\Sigma, \bullet}$, where the latter group consists of those gauge transformations that restrict, on Σ_\bullet , to identity-component gauge transformations.

¹However, in [5] and [4] it is pointed out that the issue of orientations for the moduli spaces is more difficult in the $\text{PU}(r)$ theory. By working with \mathbb{Z}_2 -coefficients we have avoided issues with orientations altogether. So from our perspective, $\text{PU}(r)$ -Floer theory on Q is identical to $U(r)$ -Floer theory on \bar{Q} , provided of course that $t_2(Q)$ is induced from an integral class.

It follows that $\mathcal{G}(\Sigma)$ acts freely on $\mathcal{A}_{\text{flat}}(Q, H)$, since each element of $\mathcal{G}(\Sigma)$ either has non-zero degree, or is in $\mathcal{G}_0(Q)$. It is immediate that the instanton Floer cohomology group $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}(\Sigma)}$ admits a relative \mathbb{Z}_4 -grading.

We want to compare the group $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}(\Sigma)}$ with those defined in the previous two sections. We begin with $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$. As in Section 3.2.2, pullback by the degree 1 gauge transformation u_1 determines a degree 4 automorphism ι_{u_1} of $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_0(Q)}$. Quotienting by the action of the group of automorphisms generated by ι_{u_1} recovers $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}(\Sigma)}$. Just as in Section 3.2.2, this says that both groups contain the same information.

In exactly the same way, u_1 determines a degree 4 map ι'_{u_1} on $HF_{\text{inst}}^\bullet(Q)^{\ker \eta}$. Once again, taking the quotient recovers the \mathbb{Z}_4 -graded group $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}(\Sigma)}$.

Remark 3.9. In [4] and [5, Section 5.6], the authors address the map ι'_{u_1} in the case when $t_2(Q)$ is the reduction of an integral class. They define ι'_{u_1} by counting instantons on a $U(r)$ -bundle over $\mathbb{R} \times Y$ determined by Σ .

3.2.4 More General \mathcal{K}

The extension of these results to more general \mathcal{K} is difficult. First, for the analysis of Section 3.1, we must have that \mathcal{K} consists entirely of connected components of $\mathcal{G}(Q)$. To generalize this would require some serious reconstructions of the theory above. We therefore continue to assume that \mathcal{K} consists entirely of connected components of $\mathcal{G}(Q)$. If a gauge transformation has non-zero degree, then it automatically acts freely by (17). So the interesting cases are when we consider gauge transformations u with $\deg(u) = 0$. These are determined, up to homotopy, by their parity. We are therefore considering the action of the finite group $H^1(Y, \mathbb{Z}_r)$. In [5], Donaldson points out that it is not clear how to construct instanton perturbations H for which $H^1(Y, \mathbb{Z}_r)$ acts freely on $\mathcal{A}_{\text{flat}}(Q, H)$. For example, one tends to construct perturbations H satisfying (H2) and (H4) using the holonomy, as in Section 4.2. When this is the case, the gauge transformations coming from elements of $H^1(Y, \mathbb{Z}_r)$ typically will *not* act freely on the space of H -flat connections.

3.3 Instanton Floer cohomology for broken circle fibrations

Now we specialize the discussion to the case relevant to the quilted Atiyah-Floer conjecture. Suppose $f : Y \rightarrow S^1$ is a broken circle fibration as in Section 2, and fix an integer $r \geq 2$. Let $d \in \mathbb{Z}_r$ be a generator and let $Q \rightarrow Y$ be a principal $PU(r)$ -bundle that is d -compatible with f . In particular, $t_2(Q) [\Sigma_0] = d$, where Σ_0 is a regular fiber of f . Since $d \in \mathbb{Z}_r$ is a generator, this clearly satisfies hypothesis (H1) with $\Sigma = \Sigma_0$.

We will work with the relatively \mathbb{Z}_4 -graded group $HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_\Sigma}$ from Section 3.2.3; see also Example 3.8. It follows that this is uniquely determined by the data of $Y, f, t_2(Q), r$ and d (in fact, it follows from the definition that it only depends on f through its homotopy class in $[Y, S^1]$). So in analogy with the

quilted theory, we set

$$HF_{\text{inst}}^\bullet(Y, f, Q)_{r,d} := HF_{\text{inst}}^\bullet(Q)^{\mathcal{G}_{\Sigma_\bullet}}$$

and call this the *instanton Floer cohomology* of $(Y, f, t_2(Q), r, d)$. The goal of this section is to describe the generators and boundary operator for $HF_{\text{inst}}^\bullet(Y, f, Q)_{r,d}$ in a manner analogous to the quilted case in Section 2.5.2.

We begin by restricting attention to a subclass of instanton perturbations that will help tie in the discussion with quilted Floer theory. Suppose for each $j \in \{0, \dots, N-1\}$ we are equipped with a time-dependent function

$$H_{j,t} : \mathcal{A}^{k,p}(P_j) \rightarrow \mathbb{R}$$

that vanishes for $t \notin (0, 1)$, and is invariant under $\mathcal{G}_0(P_j)$ for each t . Assume this is smooth with respect to the $W^{k,p}$ -topology, though we suppress this from the notation. Then we obtain a smooth function $H : \mathcal{A}(Q) \rightarrow \mathbb{R}$ by the formula

$$H(a) := \sum_{j=0}^{N-1} \int_0^1 H_{j,t} \left(a|_{\{t\} \times \Sigma_j} \right), \quad (28)$$

and this is invariant under the gauge subgroup $\mathcal{G}_{\Sigma_\bullet}$. Let $X : \mathcal{A}(Q) \rightarrow \Omega^2(Y, Q(\mathfrak{g}))$ represent the linearization dH as in the beginning of Section 3.1. Similarly, there is a map $X_{j,t} : \mathcal{A}(P_j) \rightarrow \Omega^1(\Sigma_j, P_j(\mathfrak{g}))$ representing the linearization of $H_{j,t}$ (this is the Hamiltonian vector field for $h_{j,t}$). One can check these are related by

$$X(a)|_{I \times \Sigma_j} = dt \wedge X_{j,t}(\alpha(t)),$$

where we have written $a = \alpha(t) + \psi(t) dt$ on $I \times \Sigma_j$. Moreover, $X(a)$ vanishes on Y_\bullet . If, in addition, H is an instanton perturbation (i.e., $V := X$ satisfies conditions (a) and (b) from Theorem 2.5), then we will say that H is *compatible (with quilted Floer theory)*.

It is often convenient to identify H with the split-type Hamiltonian $\underline{H} = \underline{H}_t = (H_{j,t})_j$ on $\mathcal{A}(P_0) \times \dots \times \mathcal{A}(P_{N-1})$. Moreover, restricting $H_{j,t}$ to $\mathcal{A}_{\text{flat}}(P_j)$ and then quotienting by $\mathcal{G}_0(P_j)$, we obtain a Hamiltonian on $M(P_j)$ as in Section 2.5.2. As another point on notation, if α is a connection on Σ_\bullet , then we write

$$X_t(\alpha) \in \Omega^1(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$$

for the t -dependent vector field given over Σ_j by $X_{j,t}(\alpha|_{\Sigma_j})$.

The discussion above carries over directly to the bundle $Q^\epsilon \rightarrow Y^\epsilon$, equipped with the ϵ -smooth structure and metric g_ϵ (see Sections 2.1 and 2.4). Since instanton Floer cohomology depends only on the diffeomorphism type of Y and isomorphism class of Q , it is clear that $HF^\bullet(Y^\epsilon, f, Q^\epsilon)_{r,d} = HF^\bullet(Y, f, Q)_{r,d}$ is independent of ϵ .

On the other hand, there is ϵ -dependence at the chain level. To see this dependence on ϵ explicitly, we pass to certain local coordinates as follows: Let H

be a compatible instanton perturbation that is smooth on $\mathcal{A}(Q)$. The compatibility condition tells us that $H(a)$ only depends on $a|_{I \times \Sigma_\bullet}$, and so H is actually a smooth function on $\mathcal{A}(Q^\epsilon)$ for any $\epsilon > 0$. For example, the 2-form $X(a)$ vanishes on Y_\bullet , and so does not fall victim to the pathological behavior illustrated in Figure 3.

We have seen that every connection A on $\mathbb{R} \times Y$ can be written in the form $A = a(s) + p(s) ds$. Similarly, over $\mathbb{R} \times I \times \Sigma_\bullet$ we can write $A|_{\{(s,t)\} \times \Sigma_\bullet} = \alpha(s,t) + \phi(s,t) ds + \psi(s,t) dt$, for some $\alpha : \mathbb{R} \times I \rightarrow \mathcal{A}(P_\bullet)$ and $\phi, \psi : \mathbb{R} \times I \rightarrow \Omega^0(\Sigma_\bullet, P_\bullet(\mathfrak{g}))$. It follows that $a(s)|_{\{t\} \times \Sigma} = \alpha(s,t) + \psi(s,t) dt$ and $p(s)|_{\{t\} \times \Sigma} = \phi(s,t)$. The curvature F_A on the four-manifold $\mathbb{R} \times I \times \Sigma_\bullet$ can be written in terms of its components as

$$F_A = F_\alpha + (\partial_s \psi - \partial_t \phi - [\psi, \phi]) ds \wedge dt - (\partial_s \alpha - d_\alpha \phi) \wedge ds - (\partial_t \alpha - d_\alpha \psi) \wedge dt.$$

With this notation in place, we can therefore view the instanton boundary operator ∂_{inst} as counting isolated $\mathcal{G}_{\Sigma_\bullet}$ -equivalence classes of connections $A \in \mathcal{A}(\mathbb{R} \times Q^\epsilon)$ satisfying

$$\begin{aligned} \bullet ((g_\epsilon, H)\text{-ASD on } \mathbb{R} \times I \times \Sigma_\bullet) \quad & \partial_s \alpha - d_\alpha \phi + *(\partial_t \alpha - X_t(\alpha) - d_\alpha \psi) = 0 \\ & \partial_s \psi - \partial_t \phi - [\psi, \phi] + \epsilon^{-2} * F_\alpha = 0 \\ \bullet ((g_\epsilon, H)\text{-ASD on } \mathbb{R} \times Y_\bullet) \quad & \partial_s a - d_a p + \epsilon^{-1} * F_a = 0 \end{aligned} \tag{29}$$

with appropriate limits at $\pm\infty$. Here the norms and Hodge stars are all with respect to the fixed metric g . Compare the small ϵ -limit of these equations with (14).

4 The quilted Atiyah-Floer conjecture

In this section we state and discuss the quilted Atiyah-Floer conjecture and its cousins. Throughout we refer to Sections 2.5.2 and 3.3 for notation. We begin by discussing a simple invariance property of Floer cohomology. Fix an integer $r \geq 2$ and a generator $d \in \mathbb{Z}_r$. For $j = 0, 1$, let $f_j : Y_j \rightarrow S^1$ be a broken circle fibration and $Q_j \rightarrow Y_j$ a bundle that is d -compatible with f_j . Suppose there is an orientation-preserving diffeomorphism $Y_0 \rightarrow Y_1$ pulling f_1 back to f_0 and Q_1 back to Q_0 . Then this diffeomorphism induces canonical isomorphisms of \mathbb{Z}_4 -graded abelian groups:

$$\begin{aligned} HF_{\text{symp}}^\bullet(Y_0, f_0, Q_0)_{r,d} &\cong HF_{\text{symp}}^\bullet(Y_1, f_1, Q_1)_{r,d} \\ HF_{\text{inst}}^\bullet(Y_0, f_0, Q_0)_{r,d} &\cong HF_{\text{inst}}^\bullet(Y_1, f_1, Q_1)_{r,d}. \end{aligned} \tag{30}$$

Now we can formally state the conjecture.

Quilted Atiyah-Floer Conjecture. Fix an integer $r \geq 2$ and let $d \in \mathbb{Z}_r$ be a generator. Given any broken circle fibration $f : Y \rightarrow S^1$ and principal $\mathrm{PU}(r)$ -bundle $Q \rightarrow Y$ that is d -compatible with f , there is an isomorphism $\mathrm{HF}_{\mathrm{symp}}^\bullet(Y, f, Q)_{r,d} \cong \mathrm{HF}_{\mathrm{inst}}^\bullet(Y, f, Q)_{r,d}$ of relatively \mathbb{Z}_4 -graded abelian groups. Moreover, this isomorphism is natural in the sense that it intertwines the isomorphisms (30).

There are various stronger versions of this conjecture. For example, one could repeat the discussion of Sections 2.5 and 3, generating the chain complexes over \mathbb{Z} , rather than \mathbb{Z}_2 . This would require coherently orienting the moduli spaces of instantons and J -holomorphic curves in such a way that the boundary operator in each theory still squares to zero.

A second direction in which one could state a stronger conjecture is via a chain level version. The remainder of this section is dedicated to describing such a version. We let $\mathrm{CF}_{\mathrm{symp}}^\bullet(Y, f, Q, \underline{J}, \underline{H})_{r,d}$ (resp. $\mathrm{CF}_{\mathrm{inst}}^\bullet(Y, f, Q, g, H)_{r,d}$) denote the \mathbb{Z}_4 -graded quilted (instanton) Floer chain complex. We are suppressing the boundary operator in the notation. As usual, whenever we write this, we assume the data have been chosen so that these chain complexes are defined. Just as at the homology level, these chain groups satisfy an independence property: Suppose there is an orientation-preserving diffeomorphism from Y_0 to Y_1 that pulls the data on Y_1 back to the data on Y_0 . Then this diffeomorphism determines a canonical isomorphism of the associated \mathbb{Z}_4 -graded chain complexes.

Quilted Atiyah-Floer Conjecture (Chain Level Version). Fix an integer $r \geq 2$ and let $d \in \mathbb{Z}_r$ be a generator. Then given any broken circle fibration $f : Y \rightarrow S^1$, and principal $\mathrm{PU}(r)$ -bundle $Q \rightarrow Y$ that is d -compatible with f , there are choices of $\underline{J}, \underline{H}$ and g, H such that the following conditions hold.

- (a) The chain complexes $\mathrm{CF}_{\mathrm{inst}}^\bullet(Y, f, Q, g, H)_{r,d}$ and $\mathrm{CF}_{\mathrm{symp}}^\bullet(Y, f, Q, \underline{J}, \underline{H})_{r,d}$ are well-defined.
- (b) There is a quasi-isomorphism

$$\Psi : \mathrm{CF}_{\mathrm{inst}}^\bullet(Y, f, Q, g, H)_{r,d} \longrightarrow \mathrm{CF}_{\mathrm{symp}}^\bullet(Y, f, Q, \underline{J}, \underline{H})_{r,d}$$

that preserves the gradings and is natural in the sense that it intertwines the canonical isomorphisms on each chain complex.

4.1 A program for proving the conjecture

Here we introduce a program for proving the chain level Atiyah-Floer conjecture. At the end we discuss briefly why we believe this is a fruitful approach.

To begin, we need to prescribe how the auxiliary data $\underline{J}, \underline{H}$ and g, H can be chosen. Let g be any metric on Y , and fix $\epsilon > 0$. Define the metric g_ϵ as in Section 2.1. It is convenient (but not necessary) to assume that g restricts to a metric on $I \times \Sigma_\bullet$ that is constant in the I -direction. The choice of g_ϵ induces a

split-type compatible almost complex structure \underline{J}_g as in Remark 2.13 (a). Moreover, \underline{J}_g is independent of ϵ because of the conformal invariance of 1-forms on surfaces.

In Section 4.2, we describe how to construct instanton perturbations H that are compatible with quilted Floer theory in the sense of Section 3.3. Any such H induces a split-type Hamiltonian \underline{H}_t as in Section 3.3. Moreover, Proposition 4.4 shows that H can be chosen to be invariant under $\mathcal{G}(Q)$, and so that $CF_{\text{inst}}^\bullet(Y^\epsilon, f^\epsilon, Q^\epsilon, g_\epsilon, H)_{r,d}$ and $CF_{\text{symp}}^\bullet(Y^\epsilon, f^\epsilon, Q^\epsilon, \underline{J}_g, \underline{H}_t)_{r,d}$ are well-defined. This takes care of (a) in the statement of the chain level conjecture.

Remark 4.1. *In the set-up we are describing, we are working with the ϵ -dependent bundle $Q^\epsilon \rightarrow Y^\epsilon$, for a fixed $\epsilon > 0$ small. This ensures that the metric g_ϵ is smooth. However, the associated instanton chain complex only depends on ϵ through the boundary operator ∂_{inst} . That is, the underlying abelian group*

$$CF_{\text{inst}}^\bullet(Y^\epsilon, f^\epsilon, Q^\epsilon, g_\epsilon, H)_{r,d} = \bigoplus_{[a] \in \mathcal{I}_H(Q^\epsilon)} \mathbb{Z}_2 \langle [a] \rangle$$

is independent of ϵ , up to the action of the canonical diffeomorphism (5).

In the case of the quilted Chain complex, the independence of ϵ is even stronger. Since $\underline{J}_{g_\epsilon} = \underline{J}_g$ for all ϵ , it follows that both the abelian group $CF_{\text{symp}}^\bullet(Y^\epsilon, f^\epsilon, Q^\epsilon, \underline{J}_{g_\epsilon}, \underline{H}_t)_{r,d}$ and the boundary operator ∂_{symp} are independent of ϵ . Consequently, we will drop ϵ from the notation, unless it is relevant to the discussion.

As for (b) in the statement of the chain level conjecture, it suffices to define the map Ψ by specifying its value on the generators $\mathcal{I}_H(Q)$ of the domain. Similarly, $\mathcal{I}_{\underline{H}_t}(\underline{L}(Q))$ denotes the generating set of the quilted complex. Given any $[a] \in \mathcal{I}_H(Q)$, one can obtain an element of $\mathcal{I}_{\underline{H}_t}(\underline{L}(Q))$ by restricting $[a]$ to each cobordism in Y_\bullet (to do this one needs to know that H is compatible with quilted Floer theory, which we have assumed is the case):

$$\begin{aligned} \Psi : \mathcal{I}_H(Q) &\longrightarrow \mathcal{I}_{\underline{H}_t}(\underline{L}(Q)) \\ [a] &\longmapsto e_a := \left([a|_{Y_{12}}], [a|_{Y_{23}}], \dots, [a|_{Y_{N1}}] \right). \end{aligned} \quad (31)$$

Here the bracket on the left denotes the $\mathcal{G}_{\Sigma_\bullet}$ -equivalence class, and the brackets on the right denote the \mathcal{G}_0 -equivalence class on the relevant 3-manifold $Y_{j(j+1)}$. Observe that if $u \in \mathcal{G}_{\Sigma_\bullet}$ and $a \in \mathcal{A}_{\text{flat}}(Q, H)$, then the restrictions $a|_{Y_{i(i+1)}}$ and $u^*a|_{Y_{i(i+1)}}$ are $\mathcal{G}_0(Q_{i(i+1)})$ -gauge equivalent (this is essentially the definition of $\mathcal{G}_{\Sigma_\bullet}$). Hence, the map (31) is well-defined.

Theorem 4.2. *The map Ψ defined in (31) is a set bijection. Moreover, Ψ is natural and respects the relative \mathbb{Z}_4 -gradings.*

Proof. The naturality is immediate. That Ψ respects the gradings is proved in [11]. The bijectivity is really just a direct consequence of the definitions, but

we supply a proof for convenience. The idea is that, given any perturbed Lagrangian intersection point \underline{e} , the matching conditions on the boundary imply that we can construct an H -flat connection a that maps to \underline{e} under Ψ . When constructing a there is a choice involved, and the ambiguity is measured by the group $\mathcal{G}_{\Sigma_\bullet}$. However, by Example 3.8, the set $\mathcal{I}_H(Q)$ is exactly the quotient $\mathcal{A}_{\text{flat}}(Q, H)$ by the action of $\mathcal{G}_{\Sigma_\bullet}$, and so the class of a in $\mathcal{I}_H(Q)$ is uniquely determined by the intersection point \underline{e} .

Now we work this out in detail. To simplify the notation we assume $H = 0$, and so $\underline{H}_t = 0$. We will denote by $\iota : \Sigma_i \hookrightarrow Y_{i(i+1)}$ and $\iota' : \Sigma_{i+1} \hookrightarrow Y_{(i-1)i}$ the inclusion of the boundary components. We do not keep track of the index i in the notation of ι and ι' . We first prove that Ψ is surjective. Fix some

$$\left(\left[a_{01} \right], \left[a_{12} \right], \dots, \left[a_{(N-1)0} \right] \right) \in \mathcal{I}(\underline{L}(Q)).$$

We need to show that there is some $a \in \mathcal{A}_{\text{flat}}(Q)$ with $a|_{Y_{i(i+1)}} \in \left[a_{i(i+1)} \right]$. Choose a representative $a_{i(i+1)} \in \mathcal{A}_{\text{flat}}(Q_{i(i+1)})$ for each $\left[a_{i(i+1)} \right]$. These give us an obvious definition for a over the Y_\bullet , however we need to define a over $I \times \Sigma_\bullet$ as well. To do this, let $\alpha_i := a_{(i-1)i}|_{\iota'(\Sigma_i)}$. Then $\alpha_i \in \mathcal{A}_{\text{flat}}(P_i)$ is flat, and therefore so is $\text{proj}^* \alpha_i$ where $\text{proj} : I \times \Sigma_i \rightarrow \Sigma_i$ is the projection (so $\text{proj}^* \alpha_i$ is a flat connection on the 3-manifold $I \times \Sigma_i$). The boundary condition (13) implies that there is some $\mu_i \in \mathcal{G}_0(P_i)$ with $\mu_i^* \alpha_i = a_{i(i+1)}|_{\iota(\Sigma_i)}$. By definition, $\mathcal{G}_0(P_i)$ is path-connected so there is some path $u_i : I \rightarrow \mathcal{G}_0(P_i)$ connecting the identity to μ_i . We can equivalently view the path u_i as an element of $\mathcal{G}(I \times P_i)$, i.e. as a gauge transformation over the cylinder $I \times \Sigma_i$. Then u_i^* ($\text{proj}^* \alpha_i$) is a flat connection on $I \times P_i$ that agrees with $a_{(i-1)i}$ and $a_{i(i+1)}$ at each of the two boundary components. Then we define a to be u_i^* ($\text{proj}^* \alpha_i$) over $I \times \Sigma_i$. It follows that a is continuous, flat and restricts to the desired connections over the $Y_{i(i+1)}$. By choosing the paths $u_i : I \rightarrow \mathcal{G}_0(P_i)$ appropriately, we can also ensure that a is smooth. (For example, choose the u_i so they extend to be smooth maps $\mathbb{R} \rightarrow \mathcal{G}_0(P_i)$ that are constant on the complement of $I \subset \mathbb{R}$.)

For injectivity, suppose $a, a' \in \mathcal{A}_{\text{flat}}(Q)$ are two flat connections that descend to the same Lagrangian intersection point under Ψ . We need to show that $a' = u^* a$ for some gauge transformation $u \in \mathcal{G}_{\Sigma_\bullet}$. The fact that a and a' restrict to the same Lagrangian intersection point implies that these connections restrict (modulo $\mathcal{G}_0(Q_{i(i+1)})$) to the same connection on each $Y_{i(i+1)}$. First we will show that a and a' also restrict (modulo $\mathcal{G}_0(I \times P_i)$) to the same connection on each $I \times \Sigma_i$. This is essentially a consequence of the injectivity of the restriction $M(I \times P) \hookrightarrow M(\{0\} \times P) \times M(\{1\} \times P)$ from Theorem 2.14 for bundles $I \times P \rightarrow I \times \Sigma$ over product cobordisms. Indeed, that a and a' restrict to the same connection (mod gauge) on Y_\bullet , in particular, means that they restrict to the same connection (mod gauge) on each boundary component Σ_i . So $\left[a|_{I \times \Sigma_i} \right]$ and $\left[a'|_{I \times \Sigma_i} \right]$ have the same image in $M(\Sigma_i) \times M(\Sigma_i)$ and

hence these are equal $\left[a|_{I \times \Sigma_i} \right] = \left[a'|_{I \times \Sigma_i} \right]$. So there are gauge transformations $u_i \in \mathcal{G}_0(I \times P_i)$ and $u_{i(i+1)} \in \mathcal{G}_0(Q_{i(i+1)})$ such that

$$u_i^* a|_{I \times \Sigma_i} = a'|_{I \times \Sigma_i} \quad \text{and} \quad u_{i(i+1)}^* a|_{Y_{i(i+1)}} = a'|_{Y_{i(i+1)}}.$$

The data $\left\{ u_i, u_{i(i+1)} \right\}_{i=0, \dots, N-1}$ patch together to form a global (possibly discontinuous) gauge transformation u on $Q \rightarrow Y$. We need to show that u is smooth; it will then automatically be contained in $\mathcal{G}_{\Sigma_\bullet}$.

We first claim that u is continuous. To see this, note that

$$\left(u_i|_{\{1\} \times \Sigma} \right)^* \left(a|_{\{1\} \times \Sigma_i} \right) = a'|_{\{1\} \times \Sigma_i} = a'|_{\iota(\Sigma_i)} = \left(u_{i(i+1)}|_{\iota(\Sigma_i)} \right)^* \left(a|_{\iota(\Sigma_i)} \right).$$

The functoriality of the characteristic classes from Section 2.3 implies that the map $\iota^* : \mathcal{G}(Q_{i(i+1)}) \rightarrow \mathcal{G}(P_i)$ induced by the inclusion map ι restricts to a map of the form $\mathcal{G}_0(Q_{i(i+1)}) \rightarrow \mathcal{G}_0(P_i)$. This shows that

$$\left(u_{i(i+1)}|_{\iota(\Sigma_i)} \right)^{-1} u_i|_{\{1\} \times \Sigma_i} \in \mathcal{G}_0(P_i)$$

and that this fixes $a|_{\{1\} \times \Sigma_i} = a|_{\iota(\Sigma_i)}$. By Lemma 2.9, we must have that this is the identity gauge transformation, and so $u_{i(i+1)}|_{\iota(\Sigma_i)} = u_i|_{\{1\} \times \Sigma_i}$. A similar argument shows $u_{i(i+1)}|_{\iota'(\Sigma_{i+1})} = u_{i+1}|_{\{0\} \times \Sigma_{i+1}}$. This proves the claim.

To see that u is actually smooth, we use the following trick from [6, Chapter 2.3.7] to bootstrap: As in (8), by choosing a faithful matrix representation we can write the action of u on a as $u^* a = u^{-1} a u - u^{-1} d u$, where on the right we are viewing the gauge transformation as a map $u : Q \rightarrow G$, and the concatenation is matrix multiplication. Rearranging this and setting $a' := u^* a$ gives $du = a u + u a'$. The right-hand side is C^0 , so u is of differentiability class C^1 . Repeatedly bootstrapping in this way shows that u is in C^∞ , and hence $u \in \mathcal{G}(Q)$. \square

As we have seen, in order to make each of the Floer chain complexes well-defined, we need to choose perturbations in such a way that the generators are non-degenerate. It turns out that non-degeneracy for one chain complex implies non-degeneracy for the other; see [11] for a proof.

Proposition 4.3. *Suppose H is an instanton perturbation that is compatible with quilted Floer theory, and let \underline{H} be the split-type Hamiltonian perturbation associated to H via the procedure of Section 3.3. Then the elements of $\mathcal{I}_H(Q)$ are non-degenerate if and only if the elements of $\mathcal{I}_{\underline{H}}(\underline{L}(Q))$ are non-degenerate.*

To prove the chain level quilted Atiyah-Floer conjecture, it therefore suffices to show that Ψ intertwines the boundary operators (it will automatically

induce an isomorphism at the level of homology, since it is a bijection). Recall from Remark 4.1 that though the underlying chain groups of each Floer theory is independent of $\epsilon > 0$, the instanton boundary operator ∂_{inst} does depend on ϵ since the instanton moduli spaces do. If one could show that the instanton moduli spaces of dimension zero are in bijective correspondence with the holomorphic curve moduli spaces of dimension zero, then it would be automatic that Ψ is a chain map. When $\epsilon > 0$ is small, this is likely the case since the defining equations for each moduli space are close for small ϵ . In other words, proving the chain level quilted Atiyah-Floer conjecture reduces to showing that each isolated ϵ -instanton is close to a unique isolated holomorphic curve representative, and vice versa. In [13] we show that the first holds. Showing that the ‘vice versa’ holds is currently an active research project.

In the special case where $f : Y \rightarrow S^1$ has no critical points, Dostoglou and Salamon were able to prove the chain level conjecture by exactly the approach we have described here [8, 9, 10]. Our approach to the more general case with critical points was motivated by their work.

4.2 Compatible perturbations

Fix a generator $d \in \mathbb{Z}_r$, and a broken circle fibration $f : Y \rightarrow S^1$. Let $Q \rightarrow Y$ be a principal $\text{PU}(r)$ -bundle that is d -compatible with f , in the sense of Section 2.4. Then Q satisfies hypothesis (H1) from Section 3. In this section we prove the following proposition.

Proposition 4.4. *Fix a countable set of metrics $\{g_n\}_{n \in \mathbb{N}}$ on Y . Then there is a smooth instanton perturbation $H : \mathcal{A}^{k,p}(Q) \rightarrow \mathbb{R}$ satisfying the following:*

- (i) (g_n, H) is regular for instanton Floer theory, for each n ;
- (ii) H is invariant under the full gauge group $\mathcal{G}^{k+1,p}(Q)$;
- (iii) H is compatible with quilted Floer theory;
- (iv) $(\underline{J}_0, \underline{H}_t)$ is regular for quilted Floer theory.

Here, \underline{H}_t is the split-type Hamiltonian associated to H via the procedure of Section 3.3, and \underline{J}_0 is the split-type almost complex structure associated, as in Remark 2.13 (a), to the metric g_0 on Y .

The existence of instanton perturbations satisfying (i) and (ii) is quite standard; our proof will follow [5, Chapter 5.5] closely. The novel feature here is that the perturbation can be chosen to satisfy (iii) and (iv). The proof will effectively show that the set of H satisfying (i-iv) is comeager in a suitable function space. It will also show that H can be taken to be smooth in the $W^{k,p}$ -topology on $\mathcal{A}(Q)$ for all $k \geq 0, p \geq 1$. Moreover, H can be chosen to be arbitrarily small in the following sense: For any $\epsilon > 0, p \geq 1$, and non-negative integers k_0, ℓ_0 , there is a perturbation H satisfying the conclusions of Proposition 4.4 as well as the estimates (35) and (36) with $C_H(k, \ell) = \epsilon$, for all $0 \leq k \leq k_0$ and $0 \leq \ell \leq \ell_0$.

Remark 4.5. Our intended application to the quilted Atiyah-Floer conjecture is as follows (see the discussion of Section 4.1): Fix a metric g on Y , and a sequence $\epsilon_n \rightarrow 0$ of positive numbers with $\epsilon_1 = 1$. Then we set $g_n := g_{\epsilon_n}$. The existence of H from Proposition 4.4 shows that both Floer theories are well-defined for arbitrarily small $\epsilon = \epsilon_n$. Note that, in this case, the g_n are not all smooth with respect to the same smooth structure on Y . This, however, is not an issue since compatible instanton perturbations on Q are ϵ -smooth for all $\epsilon > 0$. The reason we want a single perturbation H for arbitrarily small ϵ is that this allows us to fix H ahead of time, and then carry out the analysis in [11] and [13]. This analysis then tells us how small we need to take ϵ_n .

Proof of Proposition 4.4. Our proof is a combination of the arguments given in [8, Section 7], [5, Section 5.5], and [23, Section 3]. As usual, we drop the Sobolev exponents unless they are relevant to the discussion.

We begin by discussing the class of perturbation H we are considering; we will call these Σ_\bullet -holonomy perturbations. These, in turn, will be constructed using I -dependent holonomy perturbations on the surfaces Σ_i . To describe the latter, fix a component $\Sigma_i \subset \Sigma_\bullet$, and a finite collection of embeddings

$$\gamma_j : \mathbb{R}/\mathbb{Z} \times \mathbb{D}^1 \longrightarrow P_i, \quad j \in \{1, \dots, J\}$$

such that each $\pi \circ \gamma_j$ is orientation-preserving; here, $\pi : P_i \rightarrow \Sigma_i$ is the projection. Assume $\gamma_j(0, \tau)$ is close to a common basepoint in P_i for all j and all $\tau \in \mathbb{D}^1$. Then, for $\tau \in \mathbb{D}^1$, we can consider the map $\rho_\tau : \mathcal{A}(P_i) \rightarrow \mathrm{SU}(r)^J$ given by the based $\mathrm{SU}(r)$ -valued holonomy around each of the J loops $\theta \mapsto \gamma_j(\theta, \tau)$. Fix a word W in J variables, and view this word as a monomial $\mathrm{SU}(r)^J \rightarrow \mathrm{SU}(r)$. Then this data determines a gauge invariant \mathbb{R} -valued map $\mathcal{A}(P_i) \times \mathbb{D}^1 \rightarrow \mathbb{R}$ by the formula $(\alpha, \tau) \mapsto \mathrm{Tr}(W(\rho_\tau(\alpha)))$. Next, integrate this over \mathbb{D}^1 to obtain a smooth $\mathcal{G}(P_i)$ -invariant map $h_{\Sigma_i} : \mathcal{A}(P_i) \rightarrow \mathbb{R}$ given by

$$h_{\Sigma_i}(\alpha) = \int_{\mathbb{D}^1} \mathrm{Tr}(W(\rho_\tau(\alpha))) \beta_0 d\tau, \quad (32)$$

where $\beta_0 : \mathrm{int} \mathbb{D}^1 \rightarrow \mathbb{R}^{\geq 0}$ is a smooth compactly supported bump function with total integral 1. This is a holonomy perturbation for Σ_i . To incorporate the I -dependence (necessary from the definition of *compatible*, and for the proof below), fix another bump function $\beta_1 : I \rightarrow \mathbb{R}^{\geq 0}$ that vanishes at the boundary, and set

$$h_{i,t} := \beta_1(t) h_{\Sigma_i}.$$

We will refer to this as an I -dependent holonomy perturbation for Σ_i , and we think of $h_{i,t} = h_{i,t}(\{\gamma_j\}, W)$ as depending only on the γ_j and the word W (e.g., we think of the bump functions β_0, β_1 as fixed).

Repeating the construction of the previous paragraph for each Σ_i , one can then define $H : \mathcal{A}(Q) \rightarrow \mathbb{R}$ by (28). Suppose $H^1, \dots, H^K : \mathcal{A}(Q) \rightarrow \mathbb{R}$ are constructed in this fashion, and let $\eta : \mathbb{R}^K \rightarrow \mathbb{R}$ be any smooth map. Then we can form the function

$$\mathcal{A}(Q) \longrightarrow \mathbb{R}, \quad a \longmapsto \eta(H^1(a), \dots, H^K(a)). \quad (33)$$

We will refer to any function constructed in this way as a Σ_\bullet -holonomy perturbation.

The arguments of [23, Section 3] (applied to the linearizations of the h_{Σ_i}) show that any such Σ_\bullet -holonomy perturbation satisfies conditions (a) and (b) in Theorem 2.5. In particular, any Σ_\bullet -holonomy perturbation is a compatible instanton perturbation, and hence of the form required by the proposition we are working to prove. (In fact, the perturbation relevant for the theorem will lie in a suitable Banach space of infinite linear combinations of Σ_\bullet -holonomy perturbations.)

Fix a Σ_\bullet -holonomy perturbation H . Since the holonomy takes values in the compact set $SU(r)$, this function is uniformly bounded (each h_{Σ_i} is bounded by the r in $PU(r)$). Likewise, the linearization $X(a) \in \Omega^2(Y, Q(\mathfrak{g}))$ of H is $C^0(Y)$ -bounded with a bound independent of a . For a proof, see the discussion leading up to Lemma 5.12 in [5]; our $X(a)$ corresponds to Donaldson's T_A . The higher derivatives of H are bounded in $W^{k,p}$ -operator norm with bounds depending on r , the curvature of the connection a , and the number of derivatives taken; see [23, Proposition 7].

Rather than working on the full space $\mathcal{A}^{k,p}(Q)$, we find it convenient below to work with the quotient space

$$\mathcal{B} := \mathcal{A}^{k,p}(Q) / \mathcal{G}_{\Sigma_\bullet}^{k+1,p}.$$

In fact, by our assumptions on the bundle Q , most of our work will take place on the (open) subset $\mathcal{B}^* \subset \mathcal{B}$ of irreducible connections. Then \mathcal{B}^* is a smooth Banach manifold. (Note that since we define \mathcal{B} by quotienting only by the subgroup $\mathcal{G}_{\Sigma_\bullet}$, there is still a residual action of the finite group $\mathcal{G}_{\Sigma_\bullet} / \mathcal{G}_0(Q)$ on this space. If we had quotiented by the full group $\mathcal{G}(Q)$, then \mathcal{B}^* would have orbifold singularities.) Given any representative $a \in [a]$, the tangent space $T_{[a]}\mathcal{B}^*$ can be canonically identified with

$$\Omega^1(Y, Q(\mathfrak{g})) / \text{im } d_a.$$

The gauge invariance of any Σ_\bullet -holonomy perturbation H implies that it descends to a well-defined smooth map $\mathcal{B} \rightarrow \mathbb{R}$ on the quotient, and we denote the induced map by the same symbol H .

Now that we have described the class of functions we are interested in, we are ready to prove the proposition. This will be carried out in several stages, each having the following form: Suppose $C \subset \mathcal{B}^*$ is compact, and assume each point has a neighborhood contained in a finite-dimensional submanifold. Fix a compact set $\underline{C} \subset T\mathcal{B}^*|_C$ consisting of vectors over C . Donaldson shows [5, Section 5.5.1] that there are a finite number of perturbations² H^1, \dots, H^K so that the map

$$(H^1, \dots, H^K) : \mathcal{B}^* \rightarrow \mathbb{R}^K \tag{34}$$

²Strictly speaking, the perturbations Donaldson uses differ from ours because they are defined by considering holonomy in all three directions in Y and not just in the Σ_\bullet -directions. In Steps 1-3 below, we will discuss how to extend Donaldson's argument to our Σ_\bullet -holonomy perturbations.

restricts to a \underline{C} -immersion of C in the sense that the linearization restricts to an injection from each fiber of \underline{C} into \mathbb{R}^K . Donaldson's proof of this goes roughly as follows: Each element of \mathcal{B} is determined up to $\mathcal{G}_{\Sigma_\bullet}/\mathcal{G}_0(Q)$ by its holonomy. Likewise, each tangent vector at $[a] \in \mathcal{B}$ is determined by its linearized holonomy (there is no $\mathcal{G}_{\Sigma_\bullet}/\mathcal{G}_0(Q)$ -ambiguity at the tangent vector level since this is a discrete group). Since \underline{C} is compact, only finitely many holonomies are needed to distinguish any two elements in a fiber of \underline{C} , which proves that (34) is a \underline{C} -immersion for suitably chosen H^1, \dots, H^K . We will use a variant of this construction repeatedly, first taking C to be the space of flat connections, then coming from the space of g_n -instantons, and then from the J_0 -holomorphic curve representatives. It will then follow from a standard Sard-Smale argument that the relevant perturbed operators are surjective, which will verify (i) and (iv) in the statement of Proposition 4.4. To verify (ii) and (iii), we need to argue that, in each case, it suffices to take H^1, \dots, H^K in (34) to be Σ_\bullet -holonomy perturbations, as opposed to the perturbations Donaldson considers. This latter part will follow from irreducibility, together with a unique continuation argument.

Step 1: There is some $\epsilon > 0$, and a positive-dimensional and finite-dimensional vector space V consisting of Σ_\bullet -holonomy perturbations H , so that the following holds. Let $B_\epsilon(0) \subset V$ be the ϵ -ball. Then for a dense, open subset of $H \in B_\epsilon(0)$, all H -flat connections are non-degenerate.

To prove this, take C to be the image of the space flat connections in \mathcal{B}^* . Uhlenbeck's compactness theorem implies this is a compact set. We want to eliminate the degenerate flat connections. Note that a flat connection a is non-degenerate if and only if $H_a^1 \in T_{[a]}\mathcal{B}^*$ is zero. We therefore consider the set $\underline{C} \subset T\mathcal{B}^*$ over C whose fiber over $[a] \in C$ is the zero space if a is non-degenerate, and the unit sphere in H_a^1 otherwise. Then \underline{C} is compact as well.

First we argue that Σ_\bullet -holonomy perturbations are sufficient for our purposes. Let $[v] \in \underline{C}$ be non-zero. This can be represented by a vector $v \in \ker d_a \cap \ker d_a^*$ with unit norm, where a is a degenerate flat connections. Since v lies in the kernel of the elliptic operator $d_a \oplus d_a^*$, it satisfies a unique continuation property. In particular, it is uniquely determined by its value on $I \times \Sigma_\bullet \subset Y$. Write $v|_{I \times \Sigma_\bullet} = v + \omega dt$ and $a|_{I \times \Sigma_\bullet} = \alpha + \psi dt$. Then since v is in the kernel of d_a , we have

$$\nabla_t v = d_a \omega.$$

Moreover, since a is flat, it follows that α is as well; in particular, α is irreducible. It follows that ω , and hence all of v is uniquely determined by the value of v on $I \times \Sigma_\bullet$, which we view as a path from I into the space of 1-forms on Σ_\bullet with values in $\ker d_a^*$. Just as connections mod gauge are determined uniquely by their holonomy, 1-forms in the kernel of d_a^* are distinguished by their linearized holonomy. Since v is non-zero, the above observations imply that there is some $t \in (0, 1)$ for which $v(t) \neq 0$. This means that we can find an I -dependent holonomy perturbation $h_{i,t} : \mathcal{A}(P_i) \rightarrow \mathbb{R}$ on some Σ_i , so that the linearization is such that $d_{\alpha(t)} h_{i,t}(v(t))$ is non-negative for all t , and $d_{\alpha(t)} h_{i,t}(v(t)) > 0$ for

some t . Use this to construct a Σ_\bullet -holonomy perturbation $H : \mathcal{A}(Q) \rightarrow \mathbb{R}$ by the formula (28). It follows that

$$d_a H(v) = \int_0^1 d_{\alpha(t)} H(v(t)) dt > 0.$$

Repeating for each non-zero element of \underline{C} , and using compactness of \underline{C} , it follows that one can find a finite number of Σ_\bullet -holonomy perturbations H^1, \dots, H^K for which (34) is a \underline{C} -embedding, as desired.

To wrap up the proof of Step 1, take V to be the span over \mathbb{R} of the H_k . It now follows from a Sard-Smale argument [5, Proposition 5.15] that in any sufficiently small neighborhood of the origin in V , there is a comeager set of $(\epsilon_1, \dots, \epsilon_K) \in \mathbb{R}^K$ such that all H -flat connections are non-degenerate, where we have set $H := \sum_{k=1}^K \epsilon_k H_k$. That the set of such H is open is a consequence of the compactness of C . This completes the proof of Step 1. Note that by Proposition 4.3, it follows that all \underline{H}_t -perturbed Lagrangian intersection points are non-degenerate as well.

Step 2: Suppose H' is a Σ_\bullet -holonomy perturbation so that all H' -flat connections are non-degenerate. Fix $n \in \mathbb{N}$ and $E^{\text{inst}} > 0$. Then there is some $\epsilon > 0$, and a positive-dimensional and finite-dimensional vector space V consisting of Σ_\bullet -holonomy perturbations so that for a dense, open subset of $H'' \in B_\epsilon(0)$, the perturbation $H := H' + H''$ satisfies the following:

- (a) all H -flat connections are non-degenerate,
- (b) all (g_n, H) -instantons with energy bounded by E^{inst} are regular.

Fix H', n and E^{inst} as stated in Step 2. Since non-degeneracy is an open condition, by choosing ϵ to be sufficiently small, it follows that all H -flat connections are non-degenerate whenever $H = H' + H''$ and $H'' \in B_\epsilon(0) \subset V$. This takes care of (a). For simplicity, we may therefore assume $H' = 0$.

To prove the rest of Step 2, follow the argument Donaldson gives, starting on p. 142 in [5]. The only new point here is that we are working with the set of Σ_\bullet -holonomy perturbations, which are different than the set of perturbations considered by Donaldson. To see why this is sufficient, we extend the line of reasoning from Step 1 as follows: First, by unique continuation, it suffices to restrict attention to the end $[S, \infty) \times I \times \Sigma_\bullet$ for any S . Writing instantons here as $A = \alpha + \phi ds + \psi dt$, we may suppose that S is large enough so that $\alpha(s, t)$ is irreducible for all $(s, t) \in [S, \infty) \times I$. The game for the instanton case is to distinguish non-zero elements of H_A^+ , where A is an instanton. The elements of H_A^+ can be identified with the 2-forms W on $\mathbb{R} \times Y$ with $*W = W$ and $d_A W = 0$. These 2-forms satisfy a unique continuation property, so it suffices to distinguish them on $[S, \infty) \times I \times \Sigma_\bullet$. The self-duality condition gives

$$W|_{[S, \infty) \times I \times \Sigma_\bullet} = \omega + ds \wedge \mu + dt \wedge *_\Sigma \mu + ds \wedge dt *_\Sigma \omega,$$

where $*_{\Sigma}$ is the Hodge star on Σ_{\bullet} , and ω (resp. μ) is a map from $[S, \infty) \times I$ into the space of 2-forms (resp. 1-forms) on Σ_{\bullet} . The condition $d_A W = 0$ implies, in particular, that ω, μ are related by

$$\nabla_s \mu + *_{\Sigma} \nabla_t \mu + d_{\alpha}^* \omega = 0.$$

Since α is irreducible, this shows that the values of μ on $[S, \infty) \times I \times \Sigma_{\bullet}$ determines ω uniquely, and hence these values determine W uniquely. Just as in Step 1, we can find an I -dependent holonomy perturbation $h_{i,t}$ on some Σ_i to distinguish μ from 0, provided $W \neq 0$; here we just work, as Donaldson does, at any $s_0 \geq S$. Now use this to construct a Σ_{\bullet} -holonomy perturbation H^1 that distinguishes the 1-form $\phi := \iota_{\partial_s} W|_{s=s_0}$ from zero. Similarly, find further Σ_{\bullet} -holonomy perturbations H^2, \dots, H^K such that the linearization of

$$(H^1, H^2, \dots, H^K) : \mathcal{B}^* \longrightarrow \mathbb{R}^K$$

sends ϕ and $\partial_s \alpha$ to linearly independent vectors in \mathbb{R}^K . At this point, Donaldson's proof carries through without modification, starting at the bottom of p. 143, and this completes Step 2.

Step 3: Suppose H' is a Σ_{\bullet} -holonomy perturbation so that all H' -flat connections are non-degenerate. Fix $E^{\text{symp}} > 0$. Then there is some $\epsilon > 0$, and a positive-dimensional and finite-dimensional vector space V consisting of Σ_{\bullet} -holonomy perturbations so that for a dense, open subset of $H'' \in B_{\epsilon}(0)$, the perturbation $H := H' + H''$ satisfies the following:

- (a) all \underline{H}_t -Lagrangian intersection points are non-degenerate,
- (b) all $(\underline{J}_0, \underline{H}_t)$ -holomorphic strips with energy bounded by E^{symp} are regular.

Here \underline{H}_t is the split-type Hamiltonian associated to H .

Just as in the previous step, we will automatically have (a) provided H'' is small. The idea for obtaining the rest of the claims in Step 3 is to view each holomorphic strip in terms of representatives, i.e., connections on $\mathbb{R} \times Y$. This, in turn, can then be viewed as maps $\mathbb{R} \rightarrow \mathcal{B}^*$, and by unique continuation one could replace \mathbb{R} with an interval of the form $[S, \infty)$. Here one wants to perturb so the elements of the cokernel of each linearized Cauchy-Riemann operator vanishes. Just as with the elements of H_A^+ in the previous step, any element in a cokernel can be viewed as a path ϕ from \mathbb{R} into $T\mathcal{B}^*$; once again, unique continuation allows us to replace the domain of ϕ with $[S, \infty)$, if we wish. Write $\phi|_{I \times \Sigma_{\bullet}} = \mu + dt \wedge \eta$, where μ (resp. η) maps into the space of 1-forms (resp. 0-forms) on Σ_{\bullet} . Then irreducibility, together with the defining equations for the elements of the cokernel imply that the values of μ determine ϕ uniquely. This is analogous to what we saw with elements of H_A^+ . Now the argument follows as in Step 2.

To put this all together, let H_0 be as in Step 1. Now recursively define a sequence $\{H_k\}$ of Σ_{\bullet} -holonomy perturbations as follows: Fix sequences

$$\left\{ E_M^{\text{inst}} \right\}_{M \in \mathbb{N}'} \quad \left\{ E_N^{\text{symp}} \right\}_{N \in \mathbb{N}}$$

of real numbers increasing to ∞ . Fix an enumeration

$$\varphi : \mathbb{N}_{\geq 1} \longrightarrow \left\{ (g_n, E_M^{\text{inst}}, E_N^{\text{symp}}) \right\}_{n, M, N \in \mathbb{N}'}$$

where g_n are the metrics from the statement of the proposition. For each $\ell \geq 1$, fix a Σ_\bullet -holonomy perturbation H_ℓ as follows:

- If $\varphi(\ell) = (g_n, E_N^{\text{inst}})$ for some n, N , then take define $H_\ell := H''$, where H'' satisfies the conclusions of Step 2 with $H' := H_{\ell-1}$ and $E^{\text{inst}} := E_N^{\text{inst}}$.
- If $\varphi(\ell) = E_M^{\text{symp}}$ for some M , then take define $H_\ell := H''$, where H'' satisfies the conclusions of Step 3 with $H' := H_{\ell-1}$ and $E^{\text{symp}} := E_M^{\text{symp}}$.

At each stage, refine the choice of H_ℓ as follows: Given a Σ_\bullet -holonomy perturbation H and non-negative integers k_0, ℓ_0 , define

$$\|H\|_{k_0, \ell_0} := \sum_{k' \leq k_0, \ell' \leq \ell_0} C_H(k', \ell')$$

where $C_H(k, \ell)$ is the minimum constant satisfying the conclusion of Lemma 4.6, below. By scaling H by a small constant, the value $\|H\|_{k_0, \ell_0}$ can be made as small as we wish. Then we refine the choice of H_ℓ by demanding that

$$\|H_\ell\|_{\ell, \ell} \leq \min \left\{ 2^{-\ell}, \epsilon_\ell \right\},$$

where ϵ_ℓ is one third of the distance from $H_{\ell-1}$ to the nearest instanton perturbation not satisfying the conditions (a) or (b) from the appropriate Step. This refinement is possible at each stage because of two reasons: (i) in each of Steps 2 and 3, each $\|\cdot\|_{k_0, \ell_0}$ restricts to a well-defined norm on the vector space V , and (ii) in each Step, the set of H'' satisfying (a) and (b) is open.

Finally, consider the series $H = \sum_{\ell \in \mathbb{N}} H_\ell$. Then for each k, p , this series converges in C^∞ on bounded subsets of the space of maps $\mathcal{A}^{k,p}(Q) \rightarrow \mathbb{R}$. Moreover, H satisfies the conclusions of the proposition. \square

Lemma 4.6. *Let H be any Σ_\bullet -holonomy perturbation as in (33). Fix a non-negative integer ℓ . Then there is a constant $C_H(0, \ell)$ so that*

$$\begin{aligned} & \|d^\ell X_a(v_1, v_2, \dots, v_\ell)\|_{L^p(Y)} \\ & \leq C_H(0, \ell) \|v_1\|_{L^p(Y)} \|v_2\|_{L^p(Y)} \cdots \|v_\ell\|_{L^p(Y)}, \end{aligned} \quad (35)$$

for all connections $a \in \mathcal{A}(Q)$, 1-forms $v_1, \dots, v_\ell \in \Omega^1(Y, \mathcal{Q}(\mathfrak{g}))$, and $1 \leq p \leq \infty$. Here $d^\ell X_a : \Omega^1 \times \Omega^1 \times \dots \times \Omega^1 \rightarrow \Omega^2$ is the ℓ th derivative of $a \mapsto X(a)$ with respect to, say, the $W^{1,2}$ -topology. When $\ell = 0$, this says that $X(a)$ is uniformly bounded in L^p .

More generally, if k is a positive integer, then there is a constant $C_H(k, \ell)$ so that

$$\begin{aligned} & \|d^\ell X_a(v_1, v_2, \dots, v_\ell)\|_{W^{k,p}(Y)} \\ & \leq C_H(k, \ell) \left(1 + \|F_a\|_{W^{k-1,p}(Y)}^{k-1} \right) \|v_1\|_{W^{k,p}(Y)} \|v_2\|_{W^{k,p}(Y)} \cdots \|v_\ell\|_{W^{k,p}(Y)}. \end{aligned} \quad (36)$$

Proof. The L^p -estimates were proved in [23, Lemma 8], and reduce to the fact that the holonomy maps into the compact set $SU(r)$. In fact, Kronheimer shows that for each positive integer ℓ , the map H is bounded in \mathcal{C}^ℓ as a map $\mathcal{A}(Q) \rightarrow \mathbb{R}$ with the L^p -topology on $\mathcal{A}(Q)$. It therefore suffices to verify the $W^{k,p}$ -estimates for $k \geq 1$. We will focus on the $W^{1,p}$ -estimate for $X(a)$; the ones for the derivatives $d^\ell X_a$ are similar.

Recall that H is constructed from finitely many functions $h = h_{\Sigma_i} : \mathcal{A}(P_i) \rightarrow \mathbb{R}$ of the form (32). To establish the desired bound for X_a , it suffices to show there is some constant C so that

$$\|x(\alpha)\|_{W^{1,p}(\Sigma_i)} \leq C(1 + \|F_\alpha\|_{L^p(\Sigma_i)})$$

for all connections $\alpha \in \mathcal{A}(P_i)$; here $x : \mathcal{A}(P_i) \rightarrow \Omega^1(\Sigma_i, P_i(\mathfrak{g}))$ is the Hamiltonian vector field for h . Throughout we will liberally refer to the notation γ_j , etc., leading up to (32).

For simplicity, we assume $J = 1$ and so the word $W : SU(r) \rightarrow SU(r)$ is just the identity map. This gives

$$h(\alpha) = \int_{\mathbb{D}^1} \text{Tr}(\rho_\tau(\alpha)) \beta_0 d\tau.$$

This implies that $x(\alpha)$ a 1-form supported in the image of the thickened loop $\pi \circ \gamma_1 : \mathbb{R}/\mathbb{Z} \times \mathbb{D}^1 \rightarrow \Sigma_i$. By trivializing the bundle over this image and pulling back, we find it convenient to think of $x(\alpha)$ as a \mathfrak{g} -valued 1-form on $\mathbb{R}/\mathbb{Z} \times \mathbb{D}^1$, in which case we have

$$x(\alpha) = \text{proj } \rho_\tau(\alpha) \beta_0 d\tau, \tag{37}$$

where proj is the orthogonal projection from the Lie group $SU(r)$ to its Lie algebra, as in [23, Section 3(i)] and [5, p. 135]. This is covariantly constant in the \mathbb{R}/\mathbb{Z} -direction since this is the direction where the holonomy occurs. It follows that to estimate the $W^{1,p}$ -norm of $x(\alpha)$ we need only estimate the L^p -norm of

$$\partial_\tau x(\alpha),$$

where τ is the variable in \mathbb{D}^1 , as above.

By (37), it suffices to bound the L^p norm of $\partial_\tau \rho_\tau(\alpha)$. Let $\alpha(\tau)$ be the restriction of α to $\mathbb{R}/\mathbb{Z} \times \{\tau\}$, which we can view as a τ -dependent connection on $\mathbb{R}/\mathbb{Z} \times \{0\}$. Then, up to an overall constant depending only on the derivative $\partial_\tau \gamma_1$, the norm of $\partial_\tau \rho_\tau(\alpha) \in \mathfrak{g}$ is controlled by the norm of $\partial_\tau \text{hol}(\alpha(\tau)) \in \mathfrak{g}$, where hol is the holonomy around $\mathbb{R}/\mathbb{Z} \times \{0\}$. Letting $d\text{hol}_\alpha$ denote the linearization at α of the holonomy map, we have

$$\partial_\tau \text{hol}(\alpha(\tau)) = d\text{hol}_{\alpha(\tau)}(\partial_\tau \alpha(\tau)).$$

Now $\partial_\tau \alpha$ is, up to gauge, the $d\tau$ -component of F_α . Combining this with the uniform L^p bounds from [23, Lemma 8], we have

$$|\partial_\tau \text{hol}(\alpha(\tau))| \leq C \|F_\alpha\|_{L^p(\mathbb{R}/\mathbb{Z} \times \{\tau\})}.$$

Raising both sides to the power of p , and integrating τ over \mathbb{D}^1 gives

$$\|\partial_\tau x(\alpha)\|_{L^p(\Sigma_i)} = \|\partial_\tau x(\alpha)\|_{L^p(\mathbb{R}/\mathbb{Z} \times \mathbb{D}^1)} \leq C \|F_\alpha\|_{L^p(\mathbb{R}/\mathbb{Z} \times \mathbb{D}^1)},$$

as desired. □

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] M. F. Atiyah, N. J. Hitchin and I. M. Singer. Proceedings of the Royal Society of London. *Series A, Mathematical and Physical Sciences* Vol. 362, No. 1711 (Sep. 12, 1978), pp. 425-461
- [4] 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.
- [5] S. Donaldson. *Floer homology groups in Yang-Mills theory*. Cambridge University Press, Cambridge, 2002. With the assistance of M. Furuta and D. Kotschick.
- [6] S. Donaldson, P.B. Kronheimer. *The Geometry of Four-Manifolds*, Clarendon Press, Oxford 1990.
- [7] 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.
- [8] 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.
- [9] S. Dostoglou, D. Salamon. Self-dual instantons and holomorphic curves. *Ann. of Math.* (2) 139 (1994), no. 3, 581-640.
- [10] S. Dostoglou, D. Salamon. Corrigendum: Self-dual instantons and holomorphic curves. *Ann. of Math.* 165 (2007), 665-673.
- [11] D. Duncan. An index relation for the quilted Atiyah-Floer conjecture. arXiv:1311.4004.
- [12] D. Duncan. On the components of the gauge group for $PU(r)$ -bundles. arXiv:1311.5611.
- [13] D. Duncan. Compactness results for the quilted Atiyah-Floer conjecture. arXiv:1212.1547v3.
- [14] A. Floer. An instanton-invariant for 3-manifolds. *Comm. Math. Phys.*, 118(2):215-240, 1988.
- [15] 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.

- [16] A. Floer, H. Hofer, D. Salamon. Transversality in elliptic Morse theory for the symplectic action. *Duke Mathematical Journal* (C) October 1995 80, No. 1.
- [17] K. Fukaya. Floer homology for 3 manifold with boundary I. 1997. Unpublished manuscript, available at <http://www.math.kyoto-u.ac.jp/~fukaya/fukaya.html>.
- [18] K. Fukaya. Anti-self-dual equation on 4-manifolds with degenerate metric. *Geom. Funct. Anal.* 8, 466-528 1998.
- [19] D. Gay, R. Kirby. Indefinite Morse 2-functions: broken fibrations and generalizations, 2011. arXiv:1102.0750.
- [20] D. Gay, K. Wehrheim, C. Woodward. Connected Cerf theory. In preparation.
- [21] L. Hörmander. *The Analysis of Linear Partial Differential Operators III*.
- [22] L. Jeffrey. Extended moduli spaces of flat connections on Riemann surfaces. *Math. Ann.* 298, 667-692 (1994).
- [23] P. Kronheimer. Four-manifold invariants from higher-rank bundles.
- [24] Y. Lekili. Heegaard Floer homology of broken fibrations over the circle. arXiv:0903.1773.
- [25] J. Milnor. *Lectures on the h-Cobordism Theorem. Notes by L. Siebenmann and J. Sondow*. Princeton University Press, Princeton, N.J., 1965.
- [26] D. Salamon, K. Wehrheim. Instanton Floer homology with Lagrangian boundary conditions. arXiv:math/0607318.
- [27] F. Warner. *Foundations of differentiable manifolds and Lie groups*. Scott-Foreman, Glenview, Illinois, 1971.
- [28] K. Wehrheim. *Uhlenbeck compactness*. EMS Series of Lectures in Mathematics, 2004.
- [29] K. Wehrheim, C. Woodward. Floer field theory. Preprint.
- [30] K. Wehrheim, C. Woodward. Quilted Floer cohomology. *Geometry and Topology*, 14:833-902, 2010. arXiv:0905.1370.
- [31] K. Wehrheim, C. Woodward. Erratum: Quilted Floer cohomology. *Geometry and Topology* 16 (2012) 127-154.
- [32] 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.