

COBORDISM MAPS AND ATIYAH–FLOER ISOMORPHISMS FOR FRAMED INSTANTON HOMOLOGY II

BERND JOHANNES WUEBBEN

ABSTRACT. We prove that the mixed isomorphisms of Daemi, Fukaya and Lipyanskiy type between framed instanton homology and the Lagrangian Floer homology of a Heegaard splitting intertwine the integral instanton cobordism maps, oriented by almost-complex structures, with signed symplectic cobordism maps, up to chain homotopy and one overall sign for each cobordism. This holds for every connected cobordism between connected three-manifolds carrying a framed arc and an $\mathrm{SO}(3)$ bundle trivialized along the ends and the arc. Where the integral framed instanton homology is nonzero, the symplectic maps are signed counts of holomorphic triangles, with the Lagrangian correspondence of a compression body as one boundary condition for one- and three-handles. At three-manifolds whose integral framed instanton homology vanishes, the complexes are contractible and the maps are zero. The orientations are transported from the gauge theory through the mixed anti-self-duality and holomorphic curve equation. Consequently these maps form a projective functor on homology, independent of the handle decomposition up to sign. The unsigned analysis is that of the comparison modulo two; the new input is that two degenerations of the mixed equation, the collapse of a compression body and excision along a torus, preserve orientations. Where this homology is nonzero, we also prove that the sign by which the orientations change when a sphere is attached to a strip is constant under small perturbations and invariant under reversing the splitting. The relation with homology-oriented integral instanton maps is left open.

Date: First version: 30 September 2026. This version: 19 October 2026.

2020 *Mathematics Subject Classification.* 57K31, 57R58 (primary); 53D40, 58J05 (secondary).

Key words and phrases. Atiyah–Floer conjecture, framed instanton homology, coherent orientations, determinant lines, Lagrangian Floer homology, cobordism maps.

CONTENTS

1. Introduction	3
2. Geometric data and analytic foundations	8
3. Complexes, graded signs and determinant lines	17
4. The integral gauge convention	20
5. Determinant lines on mixed configuration spaces	25
6. The mixed isomorphism over the integers	30
7. Triangles and orientations under collapse	35
8. Torus excision over the integers	45
9. The compression cup before reflection	59
10. Sphere characters under small holonomy perturbations	63
11. Reflection of the transported orientations	68
12. Continuation and the diagonal compression cup	73
13. Small data, naturality of the mixed maps and handle squares	77
14. Descent to cobordisms and proof of the main theorem	81
Appendix A. Analytic foundations	83
Appendix B. Relative preparation and shifting	95
Appendix C. A sign calculation for projective composition	101
Appendix D. Ordinary integral excision	102
References	109

1. INTRODUCTION

1.1. The integral comparison. Let M be a closed connected oriented three-manifold. The framed instanton construction of Kronheimer and Mrowka [KM11a, Section 4.1] replaces M by $M^\# = M \# T^3$ and uses an $\mathrm{SO}(3)$ bundle which is trivial on M and odd on a torus in the added summand. A Heegaard splitting determines two Lagrangians in the moduli space of flat connections on the stabilized surface. Daemi, Fukaya and Lipyan-skiy construct an isomorphism between the resulting gauge and symplectic Floer homologies [Flo88a, Flo88b], an instance of the Atiyah–Floer conjecture [Ati88], by counting solutions of a mixed equation [DFL21]. The mixed equation couples the anti-self-duality equation on a four-dimensional region to a holomorphic curve equation on a surface; the two solutions have matching flat boundary values.

Floer homologies over $\mathbb{Z}$ carry information that is invisible modulo two: every count is a signed count, and the signs come from coherent orientations of determinant lines. The construction of [DFL21] gives an integral map for each three-manifold. Functoriality with respect to four-dimensional cobordisms is a separate question. The companion paper [Wue26a], referred to as Paper I, answers it over the field $\mathbb{F}_2$ of two elements: isomorphisms of the type of [DFL21] intertwine the Kronheimer–Mrowka cobordism maps with maps defined by holomorphic curves alone. These maps count holomorphic triangles for two-handles, and for one- and three-handles holomorphic triangles with the Lagrangian correspondence of a compression body as one boundary condition; we call the latter compression cups and caps. The analysis of Paper I rests on two degenerations of the mixed equation: a compression body collapsing next to a holomorphic strip, and excision along a torus.

This paper proves the integral version, up to one overall sign for each cobordism. The symplectic maps are the same geometric counts, now with orientations; they are not defined by conjugating instanton maps through the isomorphism. Three orientation questions arise.

- (i) The action of the lifting class used in framed instanton homology, the class in $H^1(M^\#; \mathbb{Z}/2)$ whose action by bundle automorphisms is divided out in [KM11b, Section 5.1], must preserve orientations, so that the quotient complex is defined over $\mathbb{Z}$. This is settled by the invariance of the almost-complex orientations under bundle automorphisms (Section 4), by Proposition 4.2 for the conversions, and by Proposition 5.4 on the symplectic side.
- (ii) The two degenerations must preserve the orientations of individual solutions. This is the content of Theorems B and C.
- (iii) The orientation system transported through a reflected mixed domain must be compared with the dual of the original system. The compression cup supplies a tensor, and this comparison turns it into a map $CF(M, p) \rightarrow CF(M^+, q)$, the form in which Paper I states the cup. This is Theorem D.

An unspecified sign on each generator would not suffice: the ambiguity in a cobordism map must be a single sign.

On the gauge side we use the almost-complex orientations of Kronheimer–Mrowka [KM11b, Section 5.1], which are invariant under bundle automorphisms. The almost-complex structure on a cobordism is fixed near the odd torus, and its restrictions to the ends need not be the reference structures of the objects. We therefore compose with the isomorphism of orientation lines between the two structures, called a *conversion*. When the index of the operator interpolating between the two structures is odd, the position of the conversion relative to the generator orientation line changes the map by the grading operator rather than by a sign (Section 4). We write $I_J^\#$ for the resulting gauge functor.

1.2. Statement. The category **B** is described in Section 2. An object is a closed connected oriented three-manifold with the framing data needed for its torus summand. A morphism is a connected oriented cobordism W equipped with a normally framed arc ν joining its ends and an $\mathrm{SO}(3)$ bundle E trivialized at the ends and along the arc. A presentation p consists of a Heegaard splitting, metric data, small primary holonomy perturbations and regularizing data. A presentation is *small* if it is regular and its primary perturbation has norm below a radius that depends only on the splitting and the metric data (Definition 2.4); small decompositions of a cobordism are words of elementary moves through small presentations (Section 13.1).

Write $C_G(M, p)$ for the integral gauge complex with the almost-complex orientation convention, and $CF(M, p)$ for the symplectic complex with transported orientations. The graded homotopy category $\mathbf{K}(\mathbb{Z})$ and the notation $F \simeq \pm G$ are defined in Section 3. For a small decomposition d of a cobordism, $S_d(W)$ is the composite of signed holomorphic triangle maps, compression cups, compression caps, continuation maps, and the pushforwards and conversions induced by frame changes, that is, changes of the trivialization of the bundle at an end (Section 2), and by other isomorphisms of decorated data (Section 14).

Theorem A (Integral comparison). *Orient the gauge complexes by almost-complex structures and the symplectic complexes by transport through the mixed equation, and let the presentations be small. Then the mixed maps*

$$N_p : C_G(M, p) \longrightarrow CF(M, p)$$

are isomorphisms of finite free integral complexes. For every morphism $(W, E, \nu) : M_0 \rightarrow M_1$, small endpoint presentations p_0, p_1 , and small decomposition d , one has

$$(1.1) \quad N_{p_1} I_J^\#(W, E, \nu) \simeq \pm S_d(W, E, \nu) N_{p_0} \quad \text{in } \mathbf{K}(\mathbb{Z}).$$

The sign is one sign for the entire map. Consequently the symplectic maps are independent of the decomposition up to sign and chain homotopy. On homology they define a projective functor and $H(N)$ is a projective natural isomorphism from $I_J^\#$ to that functor. Each symplectic homology map has the parity of the corresponding gauge map.

At objects with nonzero integral gauge homology the symplectic maps are the geometric maps of Paper I with the orientations constructed here. If the integral gauge homology vanishes, the complex is contractible; maps incident to it represent the zero morphism in $\mathbf{K}(\mathbb{Z})$, with the convention of Section 12.5. Reduction modulo two recovers the maps of Paper I, up to chain homotopy in the contractible case.

The Euler characteristic of the framed instanton homology of a rational homology sphere M is $|H_1(M; \mathbb{Z})|$ [Sca15, Corollary 1.4], so the integral homology is nonzero for every rational homology sphere; it is also nonzero for $\#^k(S^1 \times S^2)$ [Sca15, Section 7.8]. Whether framed instanton homology is nonzero for every closed three-manifold is an open question [Yan23, Question 1.5]. If the integral gauge homology of an object vanishes, so does its framed instanton homology with complex coefficients; such an object would answer that question in the negative. The second paragraph of Theorem A covers this case.

The qualifier “projective” has the usual meaning: maps are identified with their negatives. Multiplying individual elementary maps by signs is permitted, whereas multiplying different matrix entries by unrelated signs is not; the composition identities hold up to these signs.

1.3. The orientation comparisons. Theorem A rests on three further statements. The first two say that the two degenerations of Paper I preserve orientations; the third compares the transported orientations with their reflections.

For a two-handle cobordism the mixed domain of the comparison contains a compression body C next to a holomorphic strip. Paper I shows that, as the metric on C collapses, the regular index-zero solutions correspond bijectively to those of a limiting problem in which C is replaced by the Lagrangian boundary condition $L(C)$; stretching ordinary necks then separates the holomorphic triangles with boundary on the two Heegaard Lagrangians and on $L(C)$. The limiting solutions carry the orientation induced from the mixed orientations through the collapse (Definition 7.2).

Theorem B (Orientations under collapse). *The orientation of Definition 7.2 is well defined and coherent at breaks, and the collapse bijection of Paper I between regular index-zero mixed solutions at small collapse scale and regular index-zero solutions of the limiting problem is orientation preserving. This holds compatibly with ordinary gauge and strip gluing and in compact parameter families.*

Together with the integral identification of the pieces into which the limiting domain splits (Lemma 7.6) and of the gauge two-handle map with a triangle map (Lemma 7.7), Theorem B gives the two-handle square (Theorem 7.8): for a two-handle cobordism the gauge two-handle map and the holomorphic triangle map are intertwined by the mixed maps in $\mathcal{K}(\mathbb{Z})$, up to one sign.

For one-handles the comparison passes through a mixed domain whose gauge input is a connected sum along a torus. On the gauge side such a manifold is split by Floer’s excision theorem, in the form of [KM11b, Theorem 5.6]. Paper I proves the corresponding excision for the mixed equation over $\mathbb{F}_2$ by stretching a torus collar. Let M^+ be obtained from M by k one-handles. The two sides are the torus-summed objects $(-M)^\#$ and $(M^+)^\#$, and the closed cut is $Z \cong (-M)^\# M^+ \# T^3$ (Section 2.5). The folded pair has product Lagrangians

$$A_i = \tilde{P}_i^- \times R_i \subset \mathbb{X} = \mathbf{M}_h^- \times \mathbf{M}_g.$$

Write $N_\pm$ for the mixed maps of the two sides and N_Z for the mixed map of (A_0, A_1) , with gauge input Z . Let $\bar{E}$ be the Kronheimer–Mrowka reverse excision map from the tensor product of the two gauge complexes to the gauge complex of Z . The output complex of the folded pair is a complex in $\mathbb{X}$. Its identification $\mathcal{K}^r$ with the tensor product of the symplectic complexes has degree $|\Theta|$, the parity of the conversion Θ at the cut. It is canonical up to one sign and absorbs the conversion’s orientation line (Section 8). Here and below $\otimes^r$ uses the Koszul rule of Section 3, in which signs arise from passing the right-hand factor.

Theorem C (Oriented torus excision). *After the conversion at the cut, the mixed map N_Z of the folded pair composed with reverse excision is chain homotopic over $\mathbb{Z}$, up to one sign, to the tensor product of the mixed maps of the two sides:*

$$\mathcal{K}^r N_Z \bar{E} \simeq \pm (N_- \otimes^r N_+).$$

The chain homotopy is assembled from signed counts of isolated solutions in explicit one-parameter families, composed or tensored with closed maps (Theorem 8.15).

The third statement concerns the *sphere character*: the sign by which a coherent orientation of strips changes when a sphere representing the generator of $\pi_2(\mathbf{M}_h)$ is attached (Definition 10.1). For orientation systems induced by relative spin structures, the effect of sphere bubbling is governed by the background class [FOOO09, Chapter 8], [WW15a]. The transported system is defined through gauge theory and is not known to be of this form, so the sphere character has no a priori topological formula; Theorem D uses only its constancy and its invariance under reflection. The character matters because a coherent continuation orientation system with prescribed ends exists exactly when the sphere characters at the

two ends agree (Lemma 12.1). Reflection interchanges the two compression bodies; $\bar{p}$ is the resulting presentation of $-M$.

Theorem D (Reflection and the diagonal cup). *Let M have nonzero integral framed instanton homology, and fix a Heegaard splitting, metric data and the reference almost-complex structure. For primary perturbations in the small ball (Definition 2.4), the sphere character does not depend on the perturbation, and it is unchanged when the two compression bodies are exchanged. Consequently there is an isomorphism*

$$\Phi_p : CF(-M, \bar{p}) \longrightarrow CF(M, p)^\vee,$$

unique up to one overall sign and compatible with the orientations of the determinant lines over all strip configurations. With Φ_p , the compression cup of the trivial compression body ($k = 0$), whose correspondence is the diagonal, is chain homotopic, up to one sign, to the symplectic continuation map.

Theorem D needs only that the sphere character is constant in the perturbation and agrees across a fold; its value is not computed, and no relative spin structure on the compression correspondence is chosen.

1.4. Strategy. The main new ideas are the constancy and reflection invariance of the sphere character (Sections 10 and 11), which make the reflection pairing and the identification of the diagonal cup with continuation possible; the comparison of orientations under collapse through a two-parameter family of operators (Section 7); and, for excision, the identification $\mathcal{K}^r$ of degree $|\Theta|$, which absorbs an odd conversion at the cut and is read off at a single constant configuration (Section 8). Because orientations must be transported along configurations that are not solutions, about a third of the paper is linear analysis on configuration spaces (Section 5 and Appendices A and B).

The mixed isomorphism (Section 6). The integral mixed map has a unit diagonal in the filtration used in Paper I. The diagonal configuration is a constant pair, and its orientation is the canonical orientation of an invertible operator. An odd integer on the diagonal would give a mod-two isomorphism but not an integral one; the unit diagonal gives the integral isomorphism.

Collapse (Section 7). Determinant lines are transported through the collapse family, from the preglued configurations to the solutions. The continuity needed is continuity over smooth configurations with all derivatives controlled on compact sets and prescribed asymptotics; strong convergence of solutions alone does not identify the operator families. The preglued configuration is not a solution, and the chart in which the contraction runs is not the chart that defines the determinant bundle. Both are interpolated in a two-parameter family of Fredholm operators that is invertible on three edges of a square (Proposition 7.4), so transport carries the canonical orientation at the preglued configuration to the canonical orientation at the solution (Lemma 7.3). The collapse plays the role, for the mixed equation, that strip shrinking and the composition of seam conditions play in quilted Floer theory, where the corresponding orientation comparison is made in [WW15a].

Excision (Section 8). On the joined torus domains every solution carries two orientations. The closed-cut orientation combines the gauge orientation of the reverse excision cobordism, the conversion at the cut and the orientation of the closed-cut piece; along a long closed neck its counts become $N_Z \bar{E}$. The product orientation combines the orientations of the two separate domains across torus completion; at large width its counts are $N_- \otimes^r N_+$. The two are related by one isomorphism of lines for each output, read off at the joined constant pair,

and these isomorphisms assemble into the closed map $\mathcal{K}^r$. The one-parameter families of Paper I supply the homotopies.

Reflection and the cup (Sections 9–12). A folded triangle in $M_h^- \times M_g$, the folded form of a quilted triangle with seam condition $L(C)$, gives a tensor in the product of the reflected source complex and the target complex. Excision identifies this tensor with a relative invariant of the one-handle filling, which gives a cup formula before any reflection pairing is inserted. It remains to identify the pairing with the one for the symplectic continuation maps. We prove that the sphere character is constant under small changes of primary data, and then compare it across a fold. The doubled folded triangle is a disc with two outgoing strip ends, modelled on the reflected lower presentation and on the upper one. A sphere attached near one end can be moved through the unfolded disc to the other end, at zero collapse scale, where the image of the diagonal is an interior line of the unfolded disc; since node gluing and the transported orientation are continuous along the move, the two sphere characters agree (Proposition 11.2 and Figure 2). Nonvanishing homology enters here: a generic path of primary perturbations has a finite subdivision on each piece of which some intersection point persists (Lemma 10.5), and the character is read off at these points. Lemma 12.1 then identifies the diagonal cup with continuation. The cup formula gives naturality under small changes of primary data, the inverse relation between the reflected mixed map and N_p , and the one-handle square. Reflection gives the three-handle square.

Descent (Sections 13 and 14). Over $\mathbb{Z}$, equality on homology does not imply that two chain maps are homotopic. The proof therefore keeps the homotopies supplied by the one-dimensional moduli spaces and uses the invertibility of the mixed maps only after the elementary squares have been established.

1.5. Organization. Section 2 sets up the geometric data, states the analytic foundations (Propositions 2.1–2.3), defines admissibility and lists the results used from Paper I. Sections 3 and 4 fix the sign conventions and the integral gauge functor. Section 5 constructs the transported orientations, and Section 6 proves that the mixed maps are integral isomorphisms. Theorem B is proved in Section 7, Theorem C in Section 8, and Theorem D in Sections 9–12. Sections 13 and 14 prove Theorem A. The appendices prove the analytic foundations, the relative preparation and shifting, a sign calculation, and the ordinary integral excision of Kronheimer and Mrowka, with lifting groups, that the mixed excision uses.

1.6. Scope. The almost-complex convention used for $I_j^\#$ is part of the statement. We do not identify it with Scaduto’s homology-oriented integral cobordism maps [Sca15]: an odd conversion can introduce a grading operator rather than an overall sign, and Section 4 describes the additional grading information a comparison would need. We likewise do not identify N_p with the integral map of [DFL21]; their reductions modulo two are compared in Paper I. Comparison of the transported orientations with the relative-spin orientations of quilted Floer theory [WW15a, WW10] is a separate question: the curves and correspondences are those of Paper I, but an equality of mod-two counts does not supply an integral orientation comparison. Disconnected objects, units, counits, closed pairings and nontrivial bundles on the objects are outside the present statement; bundles on cobordisms may be nontrivial relative to the trivialized ends.

1.7. Notation. The following symbols are used throughout; the references give the place of definition.

$M^\#, \phi_M$	torus-summed object and its lifting group (Section 2)
$I_J^\#$	integral gauge functor with almost-complex orientations (Section 4)
$C_G(M, p), CF(M, p)$	integral gauge and symplectic complexes of a presentation p
$N_p, N_p^\dagger$	mixed map of p ; transpose of the reflected mixed map (Section 13)
$K(\mathbb{Z}), \simeq \pm$	graded homotopy category; equality up to chain homotopy and one sign (Section 3)
$\otimes^r$	tensor product of maps with the right sign rule (Section 3)
$\text{gr}_C, \Theta, \Theta $	grading operator of a complex; conversion between almost-complex structures and its parity (Section 4)
Σ_h, M_h	stabilized Heegaard surface of genus h and its moduli space of odd flat connections
$\mathfrak{h}, L(C)$	primary perturbation; Lagrangian of a compression body C
$U_i, V_i, C_{h,g}, L_C$	compression bodies and correspondence of the one-handle geometry (Section 2.5)
$t, \zeta_i, R_i, \tilde{P}_i$	height, height functions, height Lagrangians, lower Lagrangians (Section 2.5)
$\mathbb{X}, A_i, D_i, Z$	product target, folded pair, gauge pieces, closed cut
$W_{\pm, R}, W_{\pm, \infty}, W_\lambda^J$	collared, completed and joined torus domains
$E, \bar{E}$	excision and reverse excision maps of Kronheimer and Mrowka
$\mathcal{K}, \mathcal{K}^r$	identification of the output lines of the folded pair with those of the two sides, absorbing the cut conversion; $\mathcal{K}^r = \mathcal{K}^{\text{gr} \Theta }$ is the closed form
ϵ	collapse scale (Section 7)
$L_G(a), L_S(x), \mathfrak{o}(L)$	orientation lines of a gauge generator and of an intersection point; group of orientations of a line L
Γ_c	transported orientation of a configuration c (Section 5)
$\mathfrak{B}, r_*, r_\chi$	ball of primary data on which the sphere character is constant, its radius, and the small radius (Sections 10 and 13.1)
$\theta_G, \theta_S, S_{W_2}$	two-handle top cycles and triangle map (Section 7.1)
$U_C, V_C, \check{U}_C, n_p$	compression cup and cap, unpaired cup, transposed reflected mixed map (Sections 9 and 12)
χ^p, κ_p, Φ_p	sphere character, reflection ratio, reflection isomorphism (Sections 10–11)
S_d	symplectic map of a small decomposition d (Section 14)
(I1)–(I15)	principal results used from Paper I (Section 2.7)

2. GEOMETRIC DATA AND ANALYTIC FOUNDATIONS

2.1. Framed objects and cobordisms. Fix an oriented ball in M and form $M^\# = M \# T^3$ along that ball. Write $T^\circ = T^3 \setminus \text{int } B^3$. Choose a torus $T_0 \subset T^\circ$, a neighbourhood $N(T_0)$, and an odd bundle $P_T \rightarrow T^3$ with $\langle w_2(P_T), [T_0] \rangle = 1$. The bundle $P_M \rightarrow M^\#$ is obtained by joining P_T to the trivial bundle on M . The lifting class

$$\alpha_M = \text{PD}[T_0] \pmod{2} \in H^1(M^\#; \mathbb{Z}/2)$$

acts on the determinant-one configuration space by bundle automorphisms that do not lift to determinant-one gauge transformations; the framed complex uses the quotient by $\phi_M = \langle \alpha_M \rangle$, as in [KM11b, Section 5.1]. A class in $H^1(\cdot; \mathbb{Z}/2)$ acting in this way is called a *lifting class*, and a group of such classes a *lifting group*.

For (W, E, ν) as in the introduction, remove a tube around the normally framed arc and insert $T^\circ \times [0, 1]$. Denote the result by $(W^\#, E^\#)$. The lifting subgroup ϕ_W is the one induced by this torus tube, with the endpoint restrictions ϕ_{M_i} . At an intermediate disconnected cut

we retain the full subgroup specified by the tube; it need not be the product of the groups on the components. This distinction is essential in the excision calculation in Section 8.

The category identifies isomorphic decorated cobordisms and uses gluing along the fixed endpoint data for composition. The connected source and target condition permits a handle decomposition with one-, two-, and three-handles, together with changes of presentation. As in Paper I, elementary words also include metric changes, Heegaard moves and frame changes. A *frame change* of an object M is a map $u: M \rightarrow \mathrm{SO}(3)$ equal to 1 near the fixed ball; it changes the trivialization of the bundle at an end and defines a morphism $[u]: M \rightarrow M$ (Paper I, Section 4.5). On the symplectic side it acts through the symplectomorphisms of M_h induced by bundle automorphisms. A *filling frame* of a two-handle cobordism trivializes the bundle over the pieces of the standard decomposition, with a transition of the bundle's lifting class over one of them (Paper I, Definition 4.6); the *global frames* are the trivializations over the two ends obtained from it (Paper I, Section 4.4). The presentation graph is connected and the gauge cobordism map is independent of such a word. These are the geometric inputs to the descent at the end of the paper.

2.2. Compression bodies and the symplectic complex. If $M = H_0 \cup_\Sigma H_1$ is a Heegaard splitting, adjoin $[0, 1] \times T^2$ to each handlebody. The resulting torus compression bodies C_i have boundary $\Sigma_h \sqcup T^2$, where Σ_h is the stabilized Heegaard surface and h is its genus, one more than that of Σ . Let M_h be the moduli space of flat connections on the odd $\mathrm{SO}(3)$ bundle over Σ_h , with the lift convention of [DFL21, Section 2]. Its symplectic form is the usual Atiyah–Bott form. Restriction of flat connections on C_i gives a Lagrangian $L(C_i) \subset M_h$.

Primary perturbations $\mathfrak{h}_i$ are finite sums of averaged holonomy functions supported in C_i° , away from the fixed collars; they move the Lagrangians. Their critical equations are

$$(2.1) \quad *_3 F_B + \nabla_B \mathfrak{h}_i = 0.$$

We use a sign-sensitive trace convention for fixed-determinant $\mathrm{U}(2)$ holonomy: on each solid torus choose a continuous square root of the determinant holonomy and take the trace of the resulting $\mathrm{SU}(2)$ part. The choices are made once for the finite family. Equation (2.1) determines perturbed Lagrangians $P_i = L_{\mathfrak{h}_i}(C_i)$. At transverse data, their intersection points correspond to the gauge critical orbits. *Secondary perturbations* are holonomy perturbations supported in the interior of a mixed domain, away from the matching boundary and the collars, or, on a gauge cobordism, in compact collars of its ends as in [KM11b, Section 3.7]; they make the moduli spaces regular. *End perturbations* are the perturbations of the Floer equation on the gauge ends. The *regularizing data* of a presentation are its end perturbation, the cutoff of the end perturbation in the mixed equation, and its secondary data; as in Paper I, Section 3.5, the secondary data may include a compact deformation of the almost-complex structure of the section away from the matching line.

For each intersection point x , let a_x be the corresponding gauge critical point and $L_G(a_x)$ its real determinant line. For a real line L , write $\mathfrak{o}(L)$ for the rank-one free abelian group generated by its two orientations with the relation $[-o] = -[o]$. The integral symplectic chain group is

$$(2.2) \quad CF(M, p) = \bigoplus_{x \in P_0 \cap P_1} \mathfrak{o}(L_G(a_x)).$$

The differential counts rigid strips with the transported orientation system of Section 6. The mod-two reduction of (2.2) is the chain group of Paper I.

2.3. Mixed configurations. A mixed domain consists of a gauge region and a symplectic surface meeting along a matching boundary. On the gauge side a connection satisfies a perturbed anti-self-duality equation. On the symplectic side a map to M_h , or to a product of such spaces, satisfies a perturbed Cauchy–Riemann equation. Along a matching boundary the surface restriction of the connection is flat and represents the value of the map. A piece, end or configuration is *ordinary* if it is purely gauge-theoretic or purely symplectic, as opposed to mixed. At ordinary ends we prescribe irreducible nondegenerate gauge limits or transverse Lagrangian intersection points. At a mixed end the asymptotic values range over a critical manifold; we keep them as finite-dimensional *stationary coordinates*, separate from the decaying part, as in the extended weighted spaces of [DFL21, Definition 3.9].

The mixed domain on which N_p is defined is the special quintuple of [DFL21, Section 3.1]; a *domain of DFL type* has the same pieces and ends, up to compactly supported changes of the seam placement and of the neck lengths and to the translation-invariant data on the mixed ends, with the metrics and data allowed in (I1). The mixed domains used below are the domain defining N_p , mixed and folded triangles, their strip gluings, the closed-cut mixed domain, the finite-width torus domains, their completed limits, and the collapse families. We also use ordinary strips, continuation strips, and the disks obtained by doubling and capping a folded triangle. Statements about determinant lines apply on configuration spaces, not just on their solution subsets.

We use the smooth topology on these spaces: in fixed end trivializations all derivatives converge on compact sets, the asymptotic limits converge with their parameters, and the decaying parts converge in the chosen weighted norms at every derivative order. On a completed torus domain of Paper I, Definition 10.1(b), the torus tail $T \times [0, \infty)_r \times \mathbb{R}_\theta$ is one of these ends, with limit the odd flat connection α : in the fixed product trivialization of the tail, the difference of the connection from α is a decaying part, measured in H_δ^m with the weight $\tau_0(\theta)$ of Paper I, Definition 10.1. A continuous compact family means a family continuous in this topology. Sobolev completions are used locally to construct Fredholm operators and determinant lines. They do not replace the smooth topology when a holonomy family is differentiated repeatedly.

2.4. Analytic foundations and shifting. The unsigned moduli spaces, their energy-index estimates, and the low-dimensional compactifications are those of Paper I. The linear analysis used in addition is collected in the next proposition. It is proved in Appendix A from the local theory of [DFL21, DFL24], the end models of Paper I, and standard results on pseudoholomorphic curves. The comparison at affine centres and the independence of the end weights, described after the proposition, are Propositions 7.4 and 5.3, proved in Section 7 and Appendix A.10.

Proposition 2.1 (Analytic foundations). *For each domain listed above, and each compact family of smooth configurations and data used below, the following hold. The families may vary the almost-complex structures and perturbations smoothly, the metrics among those which near every matching seam are products up to a conformal factor depending on the base coordinates, the finite torus lengths and widths in compact ranges, and the collapse scale in compact intervals of positive scales; on the completed torus domains they are compact in the topology above.*

- (a) *Every smooth configuration with fixed asymptotics has local matching-boundary charts. The realized operators of Definition 5.1 and their formal adjoints on the dual weighted spaces satisfy local estimates and compact-remainder estimates uniform on the family. They are Fredholm, their kernels and cokernels consist of smooth elements, and they*

form continuous families. Changes of chart and the permitted changes of realization give the same determinant line.

- (b) At ordinary ends the weights lie in common spectral gaps; at mixed ends the stationary coordinates are separated from the decaying part. The scalar Neumann operator $d_A^* d_A$ is invertible on the weighted spaces at every configuration whose restriction to an odd surface or torus collar is irreducible, with bounds uniform on the family. Regular zero-dimensional solutions and the constant reference pairs with no secondary or inhomogeneous data have invertible realized operators, including the constant pairs on the finite, completed and joined torus domains and on the closed-cut piece; by openness the same holds for sufficiently small data. For the preglued families of Lemma 8.5 the estimates of (a) are uniform as the torus lengths tend to infinity, with the completed domains as limits. At positive collapse scale the estimates are uniform on compact intervals of scales bounded away from zero. The zero-scale limit domains are included.
- (c) The pure symplectic continuation, cap and strip-gluing moduli spaces have regular low-dimensional compactifications, with no additional bubbling boundary in the dimensions used here, and along the gluing paths from preglued configurations to regular solutions determinant transport preserves the canonical orientations. Let a disc contain an interior node and be separated from the matching and boundary collars. Then the real Cauchy–Riemann part of a mixed or pure operator glues at the node: the determinant lines form a continuous line bundle in the plumbing topology, the node contributes its complex fibre, canonically oriented and even, and node gluing is associative and commutes with ordinary strip gluing.

The weights lie in the signed spectral gaps of the ends, the adjoints use the opposite signed weights, and every stationary and parameter column lies in the weighted target. The fixed-positive-scale comparison at affine centres (Proposition 7.4) uses one sufficiently small weight in the indicial gaps; independence of the weight within the full spectral gap is proved for solutions and their counts (Proposition 5.3). The constants are uniform on each fixed compact family. As the collapse scale tends to zero the estimates of (b) are not uniform; Lemma 7.1 works with the logarithmic loss of Paper I.

The second input concerns relative preparation and shifting in families. Daemi, Fukaya and Lipyanskiy state the one-input shift on their special quintuple, the model mixed domain of [DFL21, Section 3.1] on which their isomorphism is defined, and sketch its family version [DFL21, Lemmas 4.62 and 4.64]; the domains, topology and families used here, and the several-input form, are not covered there. The following proposition is proved in Appendix B.

Proposition 2.2 (Relative preparation and shifting). *Let a compact family of smooth configurations be given on one of the domains below, with fixed asymptotic labels, where the data and geometry may vary as in Proposition 2.1.*

- (a) On a one-input mixed domain (the domain defining the mixed map, a separate finite or completed torus domain, or the closed-cut piece with its ordinary input Z), relative preparation makes the original tails stationary beyond a common cutoff, preserving members already stationary there. The subsequent shift deforms the family to gauge paths followed by a constant mixed pair C^c , where c is a representative of the critical point over the output label (Appendix B), keeping members which are already of this form in that subspace.
- (b) On the triangles and folded triangles at positive collapse scale, and on the finite-width joined torus domains, the deformation of Definition 5.5 retains the ordinary cobordism and reaches the reference configuration described there, taken for the torus domains

at the closed cut of (I2); its relative form preserves the subspace of configurations already of that form.

These deformations are continuous in the topology of this section jointly in configurations and data, and coherent with ordinary gluing at the gauge input.

Coherence of the orientations at the output end is proved in Proposition 5.4 for the domains of (a), after Proposition 5.4 for the triangles, and in Lemma 8.3(ii) on the joined torus domains.

Limiting triangle orientations are not obtained by shifting to a constant triangle; they are defined by inverse collapse (Definition 7.2).

Section 8 uses Propositions 2.1 and 2.2 on the separate, completed and joined torus domains and on the closed-cut piece; its opening paragraph describes how they enter.

Proposition 2.3 (Ordinary determinant-one self-gluing). *Consider the nonsingular gauge cobordisms and designated closed cuts in the torus-excision construction of Section 8. Fix the ordinary Floer data on the original ends, including their trajectory perturbations. The end critical sets are finite, irreducible and nondegenerate, the low-dimensional trajectories are regular, and the restricted lifting groups act freely on these sets. There are compatible data on the new torus ends and invariant secondary perturbations, supported in fixed compact end collars, for which the cut moduli spaces and the required parameter families are regular. For these choices the following hold.*

- (a) *For fixed regular cut data and all sufficiently large neck lengths, the index-zero solutions are exactly the gluings of matched determinant-one cut solutions on the prescribed bundle. The gluing fibre is the torsor of Lemma D.1, and the neck-length threshold can be chosen uniformly over all index-zero solutions.*
- (b) *Determinant gluing uses the graded symmetry of Section 3. Its signs are determined by the ordered degrees and the fixed orientation choices. In the nonseparating torus case the determinant-one count is twice the graded partial trace, before the larger lifting-group quotient is taken.*
- (c) *Relative regular families give the integral composition and change-of-data homotopies. In a family the data parameter is shared by the pieces and occurs once in the linearization.*

The proof is in Appendix D. It applies the ordinary instanton compactification of [KM11b, Section 3.9] and its signed self-gluing formula [KM11b, Proposition 5.5]. Regular cut data are essential: nondegeneracy of the end critical points alone does not regularize a cut solution. Part (c) is a statement about parametrized moduli spaces, rather than a pointwise gluing bijection in every fibre of a family that is regular as a parametrized family. Neither Proposition 2.1 nor Proposition 2.2 is used in this proposition.

2.5. The one-handle geometry. The one- and three-handle comparisons, and the excision of Section 8, use the following geometry from Paper I, Section 10.1.

Compression bodies. Let a presentation p of M consist of torus compression bodies U_0, U_1 with boundary $\Sigma_h \sqcup T^2$, $h \geq 2$. For $k \geq 0$ let $C = C_{h,g}$, $g = h + k$, be a compression body from Σ_h to Σ_g which adds k one-handles, and put $V_i = C \circ U_i$. Then $V_0 \cup (-V_1)$ is the torus-summed object of $M^+ = M \#^k(S^1 \times S^2)$. Restriction of flat connections on C gives an embedded Lagrangian correspondence $L_C \subset M_h^- \times M_g$.

Height Lagrangians. The clean intersection $L(V_i) \cap (L(U_i) \circ L_C)$ is a copy of $SU(2)^{g-1}$. Choose a primary perturbation ζ_i of V_i , the *height function*, whose restriction to it is a

product Morse function with 2^{g-1} critical points. For small $t > 0$ perturb the primary data of V_i by $t\zeta_i$; the resulting Lagrangians $R_i \subset M_g$ are the *height Lagrangians at height t* . The lower data are perturbed Lagrangians $\tilde{P}_i \subset M_h$ of U_i close to $L(U_i)$.

The folded pair. In $\mathbb{X} = M_h^- \times M_g$ put $A_i = \tilde{P}_i^- \times R_i$. The *incidence conditions* require the three pairs among (A_0, L_C, A_1) to be transverse and their triple intersection to be empty. Holomorphic triangles in $\mathbb{X}$ with boundary on (A_0, L_C, A_1) are the folded form of quilted triangles with seam condition L_C ; the compression cups and caps count them. The gauge pieces are $D_i = (-U_i) \cup_{T^2} V_i$, with $\partial D_i = (-\Sigma_h) \sqcup \Sigma_g$; flat connections on D_i restrict bijectively to pairs of flat connections on U_i and V_i , so $L(D_i) = L(U_i)^- \times L(V_i)$, perturbed to A_i . The *closed cut* of the folded pair is

$$Z = D_0 \cup (-D_1) \cong (-M) \# M^+ \# T^3.$$

The torus domains (Figure 1). Near the torus component of its matching surface, the special quintuple X_D of a pair of torus compression bodies contains a flat product region $N_X = T \times [-1, 0]_s \times \mathbb{R}_\theta$. Since the moduli space of the odd torus is a point, the matching condition on its face only requires the restriction of the connection to each torus fibre to be flat. Attaching $T \times [0, R]_r \times \mathbb{R}_\theta$ along the face gives the *collared domain* W_R , with the same flat-fibre condition at $r = R$; for $R = \infty$ it is the *completed domain* W_∞ , whose torus tail is an end with limit the odd flat connection α . For the two sides, the reflected presentation of M (subscript $-$) and the enlarged presentation of M^+ (subscript $+$), the *joined domain* W_λ^J glues the collars of width λ along their torus faces; the glued collar $T \times [-\lambda, \lambda] \times \mathbb{R}_\theta$ is the *bridge*. The domains $W_{\pm, R}$ and $W_{\pm, \infty}$ are called *separate*. The joined domain has gauge inputs $(-M)^\#$ and $(M^+)^\#$ and strip output at $A_0 \cap A_1$. It contains a closed hypersurface $Z_{\text{cut}} \cong Z$ in its gauge part, which cuts it into the reverse excision cobordism $\bar{\mathcal{E}}_\lambda$ of Kronheimer and Mrowka [KM11b, Section 5.3] and a domain M_λ of DFL type for the folded pair with gauge input Z ; see (I2).

2.6. Small presentations.

Definition 2.4. Fix a Heegaard splitting with its metric data. Its *small ball* consists of the primary perturbations in the span of a finite family of sign-sensitive holonomy functions, containing the fixed family F_0 of Section 10 and subject to the conditions of Paper I, Definition 3.3 (in particular containing the family F_* of (I5)), whose norm (10.2), a Lipschitz norm of the gradient, is less than a radius $r_\chi > 0$. The radius depends only on the splitting and its metric data, and is transported by isomorphisms of decorated data. A presentation is *small* if it is regular and its primary perturbation lies in the small ball of its splitting and metric data. Small presentations are admissible in the sense of Paper I (see (I5)).

The radius r_χ is fixed in Section 13.1, so that the perturbed Lagrangians lie in the neighbourhoods of (I5) and of Section 9. Small fold pairs, compatible handle data and small decompositions are defined there in terms of small presentations.

2.7. Principal results used from Paper I. The following results of Paper I are stated there over $\mathbb{F}_2$; only their unsigned content is used. In addition, Sections 7 and 8 and Appendix A adapt arguments of Paper I, Sections 7 and 10, to stabilized families of configurations that are not solutions; each adaptation states which steps change.

- (I1) *Deformation invariance of the mixed map* (Paper I, Lemma 9.7). Let D be a domain of DFL type for a pair of perturbed compression-body Lagrangians, transverse or clean as in Paper I, with the end types of the special quintuple and regular invariant




FIGURE 1. The torus region, drawn in the base coordinates (normal coordinate horizontally, the end coordinate θ vertically); every point carries a torus T with the flat odd connection. Left: the collared domain W_R of Section 2.5, with the flat product region N_X of X_D , the collar and the flat face at $r = R$; for $|\theta| \geq 2$ (dotted lines) the region lies in the mixed ends. The completed domain W_∞ has $R = \infty$. Right: the joined domain W_λ^J , in which the collars of the two sides are glued along their flat faces into a bridge of width 2λ . (After Paper I, Figure 5.)

data, whose metric is conformal to a product near each seam and on each mixed end, and whose almost-complex structure on the section is $*_\Sigma$ near the seams. Then the index-zero count of D is chain homotopic over $\mathbb{F}_2$ to the mixed map of that pair. The translation-invariant data on the mixed ends may differ from those of the quintuple.

- (I2) *The closed cut* (Paper I, Lemma 10.18). The joined torus domain W_λ^J contains a closed hypersurface Z_{cut} , identified with $Z = D_0 \cup (-D_1)$, which separates it into the reverse excision cobordism and a domain of DFL type for the folded pair (A_0, A_1) with gauge input Z .
- (I3) *Torus completion* (Paper I, Theorem 10.32, with Propositions 10.27 and 10.30). For large collar length and width, gluing gives bijections between the index-zero solutions on the completed torus domains and those on the collared and joined domains. The linearized operators at the preglued configurations are invertible with uniform bounds, and each glued solution is the unique zero of the equation in a ball of fixed radius in the gluing chart.
- (I4) *The folded triangle modulo two* (Paper I, Theorem 11.8). One has $N_Z \circ \mu_{\mathcal{W}} \simeq \mu_{\mathbb{X}} \circ (N_0 \otimes N_1)$ over $\mathbb{F}_2$, where $\mu_{\mathcal{W}}$ is the gauge triangle of the pieces D_0 , C , D_1 and $\mu_{\mathbb{X}}$ the holomorphic triangle map in $\mathbb{X} = M_h^- \times M_g$; here N_0, N_1 are the mixed maps of the pairs (A_0, L_C) and (L_C, A_1) , and N_Z that of (A_0, A_1) .
- (I5) *Admissible neighbourhoods* (Paper I, Definition 11.12). Paper I fixes, with each Heegaard splitting, a finite family F_* of primary perturbations containing the height functions and having the submersion property of Paper I, Section 3.2; the generic corrections of the height Lagrangians are taken from its span.

For each sufficiently small height $t > 0$ there is a C^1 -neighbourhood $\mathcal{N}(t)$ of the pair of compression-body Lagrangians and of the reflected pair, depending only on

the splitting, metric data and height functions. For the product compression body ($k = 0$), every lower pair in $\mathcal{N}(t)$ and every height Lagrangian at height t , with correction $g \in \text{span } F_*$, $|g| \leq \delta(t)$, satisfy:

- (a) the unfolding statements of Paper I;
- (b) the auxiliary generators are the nondegenerate critical points of a Morse function C^2 -close to t times the height function, with the same indices;
- (c) their critical actions satisfy $2 \max_{a,x} |\mathcal{A}_i(a) - \mathcal{A}_i(x)| < 2\pi^2$.

For a finite family $F \supseteq F_*$ the set $\mathcal{U}_F(t)$ is the path component of zero in the set of primary data in $\text{span } F$ whose Lagrangians, and those of their reflections, lie in $\mathcal{N}(t)$; the mod-two folded triangle and the mod-two cup formula of Paper I hold at every regular point of $\mathcal{U}_F(t)$, for a height Lagrangian of that family satisfying the incidence conditions, by (I14)(ii), first for end and secondary data below its threshold and then, on homology, for the prescribed data by Theorem 3.8(ii) and Paper I, Section 11.5.

The admissible neighbourhood $\mathcal{U}(F)$ is the path component of zero in $\mathcal{U}_F(t_*) \cap \mathcal{U}_F(t'_*)$ for two heights $t'_* \leq t_*$ fixed with the splitting. It is an open neighbourhood of zero in $\text{span } F$, and $\mathcal{U}(F) \subseteq \mathcal{U}(F')$ for $F \subseteq F'$. Admissibility is a condition on the primary data through the path joining them to zero, not on their Lagrangians alone.

- (I6) *Uniform stationary gap* (Paper I, Proposition 10.7). The stationary operators on the slices of the mixed ends of the torus domains are self-adjoint with kernel the tangent space of the critical manifold, and their nonzero spectrum avoids a fixed interval about zero, uniformly in the collar length and, on the joined slices, in the width $\lambda \geq \lambda_*$.
- (I7) *Neck gluing and families* (Paper I, Theorems 5.1 and 5.4). For regular pieces joined along closed or ordinary ends, the index-zero solutions at large neck length are the gluings of matched index-zero solutions of the pieces, and the count of the joined problem is the composite of the counts of the pieces; the same holds with one index-one trajectory or one index -1 member of a family. A family of data on one domain, with fixed ends and critical data, a compactly supported parameter derivative and endpoints regular through index one, gives after a small generic perturbation invariant under the lifting group a chain homotopy between the endpoint counts, counted by its index -1 solutions, with one boundary branch for each such solution and each index-one trajectory at an end.
- (I8) *Collapse inverse and contraction* (Paper I, Propositions 7.15, 8.1 and 8.3). Let V be a regular index-zero solution of the limiting problem. For small collapse scale ϵ the linearization of the mixed equation at the preglued configuration $W_\epsilon = W_\epsilon(V)$ is invertible, with inverse bounded by a constant times $1 + \log(1/\epsilon)$, and in the strong norms of Paper I, Section 8.1, by a constant times $(1 + \log(1/\epsilon))^{m+1}$. In the matching chart the nonlinear part of the equation has derivative bounded by a constant times the strong norm on a ball of fixed radius. The contraction gives a unique regular solution $W_\epsilon + z_*$ in that ball, with $\|z_*\|$ bounded by a power of $1 + \log(1/\epsilon)$ times the pregluing error e_ϵ , which tends to zero as a positive power of ϵ .
- (I9) *Weighted parameter column* (Paper I, Lemma 10.15). The conclusions of the families statement in (I7) hold for families whose variation is translation invariant on the mixed ends rather than compactly supported, provided the critical set on each mixed end is connected and, with a family of stationary references solving the equation, independent of the parameter, the variation vanishes on the gauge and ordinary ends, and the stationary operators keep a uniform gap and an energy decay at a rate larger than the weight.

- (I10) *Invariant regularity* (Paper I, Theorem 5.6). Under the compactness and gluing hypotheses of (I7), generic small holonomy perturbations invariant under the lifting group, from a finite-dimensional space of directions supported near the fixed perturbation supports, with loops of zero lifting class joined to an original end whose limit is flat on a collar of an odd surface, and with supports and loop regions in a fixed compact set, make all moduli spaces of index at most one regular, relative to strata that are already regular.
- (I11) *Product collar at a closed cut* (Paper I, Lemma 9.12 with stages (A) and (B) of its proof, the corresponding stages in the proof of Theorem 11.8, and Lemma 10.18(v)). On the triangle domain of a two-handle cobordism at a positive collapse scale, the closed three-manifold which separates the two gauge inputs from the output (Paper I, Lemma 9.11) receives a product collar along a path of metrics: (A) the compression scale over the collapsing region is raised to one; (B) the metric is interpolated linearly to one which is a product on a collar of the cut. Along the path the scales stay in a compact positive range, the parameter derivative is supported in a compact set disjoint from the ends, and the data are invariant under the lifting group. The folded triangle, with the cut Z , has the corresponding path. On the joined torus domain W_λ^J , a change of metric supported near the two arcs of the cut, away from the bridge ends, the arms and the matching lines, gives Z_{cut} a neighbourhood $(-1, 1) \times Z_{\text{cut}}$ with the product metric.
- (I12) *Torus-fibre slice* (Paper I, Lemma 10.9). There are $e_T > 0$ and C_0 , depending only on T , with the following property. Let $\Omega \subset \mathbb{R}^2$ be simply connected, and let A be a connection on $T \times \Omega$ which near every fibre is, after an even gauge transformation, within e_T of α in C^1 . Then there is a unique even gauge transformation g of $T \times \Omega$ such that $g^*A = \alpha + b + \phi_1 dz_1 + \phi_2 dz_2$, where $b(z)$ is a one-form on the fibre with $d_\alpha^* b(z) = 0$ and $\|b(z)\|_{C^0(T)} \leq C_0 e_T$, and the $\phi_j(z)$ are sections of ad on the fibre. The connection g^*A is the *torus-fibre slice* of A . If in addition the curvature is small in C^{k+1} , the fibrewise Sobolev norms of b , ϕ and their base derivatives are bounded by those of the curvature on the fibre.
- (I13) *The auxiliary grading* (Paper I, Lemma 9.1 and Section 11.3, including Lemma 11.7). At the height data, the intersection points of the auxiliary pairs (A_0, L_C) and (L_C, A_1) are the critical points of the height function on the clean intersection $Q \cong \text{SU}(2)^r$, $r = g - 1$, of Section 2.5, a product Morse function. On their complexes a relative integer grading m is defined by $m(a) - m(x) = \text{ind}(u) - 8\kappa_{\text{aux}}(u)$, where the relative class $\kappa_{\text{aux}} \in \frac{1}{2}\mathbb{Z}$ is the topological energy corrected by the critical actions at the endpoints; the same relation holds for every gauge, mixed or strip configuration u from a to x . For small height, m is the Morse index of the height function for (A_0, L_C) and of its negative for (L_C, A_1) , and there is a unique generator q_i^* of maximal m ; at uncorrected height data it is the maximum or the minimum of the height function according to the order of the pair, and under the corrections of the height Lagrangians of Paper I, Section 11.6, the generators and q_i^* are their continuations. For small height, end perturbation and secondary perturbations, $\kappa_{\text{aux}} \geq 0$ on every counted configuration (Paper I, Lemma 9.1(iii)); at corrected height data and at changed lower data the same follows from (I14).
- (I14) *Changes of the end data* (Paper I, Lemma 11.7). Fix the height t and a finite family $F \supseteq F_*$. The end data of an auxiliary pair are its lower primary datum, in $\text{span } F$, and the primary datum $th_i + g$ of its height Lagrangian, in $\text{span } F_*$; the height data are $(0, th_i)$.

- (i) For $|g| \leq \delta(t)$ there is a ball $\mathcal{N}_{C,R}$ about zero in $\text{span } F$, depending on C , on R and on F , such that, for end and secondary data below the threshold of (ii) with $c = 2\pi^2$, under changes of the lower primary data within $\mathcal{N}_{C,R}$ the relative classes and the grading of (I13) are transported from the height data, with the endpoint values of the height function replaced by the critical actions $\mathcal{A}_i$ of the changed pairs, and the coefficient of q_i^* in $N_i\theta_i$, and in the filling invariant, remains 1 over $\mathbb{F}_2$; the representing solutions need not remain flat.
- (ii) Along a smooth path of end data starting at the height data, on which the generators of the auxiliary pairs remain nondegenerate and continue those at the height data with their gradings and relative classes, and twice the maximal difference of their critical actions stays below a constant $c < 4\pi^2$, the following hold at the endpoint, for end perturbations along the path and secondary perturbations at the endpoint satisfying

$$4 \sup \|f\|_\infty + 2 \sup \|\eta\|_{L^2}^2 < 4\pi^2 - c$$

and the thresholds of Paper I, Theorem 3.8. Every counted gauge trajectory, mixed solution and strip has $\kappa_{\text{aux}} \geq 0$; the differential of the auxiliary gauge complex vanishes, its 2^r generators matching the rank of its free homology; the filling invariant is the generator θ_i of maximal m ; and $N_i\theta_i = q_i^*$, both on chains over $\mathbb{F}_2$.

Paper I proves (ii) by transporting the filling class along a fine subdivision of the path by gauge continuation maps, which are isomorphisms triangular with respect to m , and by identifying the mixed coefficient at the endpoint through the isomorphism of Theorem 3.8(i), triangular with respect to m by the index formula, and Lemma 9.2; (i) is the case of the path that first corrects the height datum and then moves the lower datum along a segment. The hypotheses of (ii) hold with $c = 2\pi^2$ along every smooth path of lower data from zero whose Lagrangians stay in the neighbourhood $\mathcal{N}(t)$ of (I5), preceded by the correction $th_i \mapsto th_i + g$, $g \in \text{span } F_*$, $|g| \leq \delta(t)$, of the height datum.

- (I15) *Generation and connectivity* (Paper I, Propositions 12.1 and 13.2). Any two admissible presentations of an object are joined by presentation moves: changes of disk system, stabilization and destabilization away from the ball, isotopies fixed near the ball, and changes of the metric and primary data. Every morphism $(W, E, \nu): M \rightarrow M'$, with any admissible presentations p of M and q of M' , is the composite of a word of elementary moves from (M, p) to (M', q) : presentation moves, frame changes, one-handle attachments, two-handle cobordisms between the presentations of a subordinate triple with their global frames, and three-handle attachments. Each handle move may be taken between presentations satisfying the height and neighbourhood conditions of Paper I, Section 13.2.

3. COMPLEXES, GRADED SIGNS AND DETERMINANT LINES

3.1. The homotopy category. A complex is a finite free $\mathbb{Z}/2$ -graded abelian group $C = C^0 \oplus C^1$ with an odd differential d_C . For a homogeneous map $F: C \rightarrow D$ of degree f , set

$$(3.1) \quad \delta F = d_D F - (-1)^f F d_C.$$

A closed map satisfies $\delta F = 0$. Thus an odd closed map anticommutes with the differentials. We allow sums of the two degrees when discussing homotopies of the composites below.

The category $K(\mathbb{Z})$ has these complexes as objects and closed maps modulo boundaries as morphisms. In particular

$$F \simeq G \iff F - G = \delta K.$$

For homogeneous F, G of degree f , the homotopy K can be chosen of degree $f + 1$. The notation $F \simeq \pm G$ means that one of $F - G, F + G$ is a boundary, with a single choice of sign.

If A and B are closed and B has degree b , then

$$(3.2) \quad (\delta K)A = \delta(KA), \quad B\delta K = (-1)^b \delta(BK).$$

These identities justify composition of every homotopy used below. They also show that projectivizing by the relation $[F] = [-F]$ is compatible with composition.

Lemma 3.1. *An acyclic finite free $\mathbb{Z}/2$ -graded complex over $\mathbb{Z}$ is contractible. An even quasi-isomorphism between finite free complexes is a chain homotopy equivalence.*

Proof. Let $B^i = \text{im}(d : C^{i-1} \rightarrow C^i)$. Each B^i is free as a subgroup of a finitely generated free group. If the complex is acyclic, the exact sequence $0 \rightarrow B^i \rightarrow C^i \xrightarrow{d} B^{i+1} \rightarrow 0$ splits. Relative to these splittings the differential identifies the complementary summand of C^i with B^{i+1} . Its inverse on this summand, and zero on the complement, defines a contracting homotopy, denoted s , with $ds + sd = 1$. Apply this to the mapping cone of a quasi-isomorphism to obtain the second assertion. $\square$

Over $\mathbb{Z}$ the lemma is applied only to acyclic complexes and to quasi-isomorphisms; a map that induces zero on homology need not be null-homotopic.

3.2. The tensor convention. The orientation lines of the generators precede the determinant lines of trajectories and cobordisms. With this ordering the tensor differential is

$$(3.3) \quad d(x \otimes y) = (-1)^{|y|} dx \otimes y + x \otimes dy.$$

For homogeneous maps F, G we use

$$(3.4) \quad (F \otimes^r G)(x \otimes y) = (-1)^{|F||y|} Fx \otimes Gy.$$

The superscript in $\otimes^r$ records that signs arise from passing the right-hand factor, as for F^r and Θ^r . Closed maps tensor to a closed map of degree $|F| + |G|$ with this rule. The usual left tensor differential is related to (3.3) by the even chain isomorphism

$$T(x \otimes y) = (-1)^{|x||y|} x \otimes y.$$

Thus a formula imported in a left tensor convention is conjugated by T at each tensor end before it is used here.

Lemma 3.2. *For a disconnected gauge end $Y_1 \sqcup Y_2$, with the orientation lines of the generators ordered as $L(a_1) \otimes L(a_2)$, gluing index-one trajectories on the two components gives (3.3).*

Proof. A trajectory on Y_1 has an odd determinant line. To place it next to $L(a_1)$, move it past $L(a_2)$; this gives $(-1)^{|a_2|}$. A trajectory on Y_2 is already adjacent to the orientation line of its generator. Summing the two contributions gives the formula. $\square$

Put $\text{gr}_C x = (-1)^{|x|} x$. If a homogeneous map F of degree f commutes strictly with the differentials, the map $F^r = F \text{gr}_C^f$ is closed. For strictly commuting maps F, G of degrees f, g , respectively,

$$(3.5) \quad (FG)^r = (-1)^{fg} F^r G^r.$$

The discrepancy is one sign per composite. By contrast the factor gr_C is not generally an overall sign.

For example, on $C^0 = C^1 = \mathbb{Z}$ with zero differential, the map interchanging the two generators and its product with gr_C are not equal up to sign. The difference persists on homology. This elementary example explains why an integral comparison must specify the convention for odd cobordism maps.

3.3. Duals, evaluation and tensors of maps. The graded dual has

$$d^\vee f = -(-1)^{|f|} f \circ d, \quad F^\vee(f) = (-1)^{|F||f|} f \circ F.$$

The canonical even double-dual map is

$$\text{ev}_C : C \longrightarrow C^{\vee\vee}, \quad \text{ev}_C(x)(f) = (-1)^{|x||f|} f(x).$$

The sign is necessary: the unsigned double-dual identification does not commute with the differential in this convention. There is an unsigned chain isomorphism

$$(3.6) \quad \beta : C^\vee \otimes^r D \longrightarrow \text{Hom}(C, D), \quad \beta(f \otimes y)(x) = f(x)y.$$

Indeed, substituting the dual differential in (3.3) gives exactly (3.1) after evaluation on a homogeneous x . For a tensor in $C \otimes^r D$ the *bending map* is

$$\text{bend}_C = \beta \circ (\text{ev}_C \otimes^r 1) : C \otimes^r D \longrightarrow \text{Hom}(C^\vee, D).$$

If e_i is a homogeneous basis of C , the tensor corresponding to a homogeneous map $F : C \rightarrow D$ is

$$\text{ten}(F) = \sum_i e_i^\vee \otimes F(e_i).$$

This expression is independent of the chosen basis. It is a cycle if and only if F is closed. For $F = 1$, it is the coevaluation tensor in the present convention.

3.4. Determinant lines. For a real Fredholm operator $D : X \rightarrow Y$ set

$$\det D = \Lambda^{\max} \ker D \otimes (\Lambda^{\max} \text{coker } D)^*, \quad |\det D| = \text{ind } D \pmod{2}.$$

We use the graded determinant convention of [Zin16]. Stabilization by a finite-dimensional map $V \rightarrow Y$ with $D \oplus V$ surjective gives the local trivializations. Enlarging the stabilization gives compatible transition maps. This construction applies to continuous families of Fredholm operators without requiring constant kernel dimension.

For direct sums the determinant lines are ordered in the order of the summands. Interchanging two summands gives $(-1)^{\text{ind } D_1 \text{ind } D_2}$. The dual of a tensor reverses the order of its factors and uses nested evaluation. An invertible operator has the canonical determinant orientation of $\mathbb{R}$; a continuous path of invertible operators preserves this orientation.

Lemma 3.3 (Ordinary gluing). *Suppose two Fredholm problems have matching ordinary ends with invertible asymptotic operator. Gluing with a sufficiently long neck gives a determinant-line isomorphism*

$$\det D_1 \otimes \det D_2 \longrightarrow \det(D_1 \# D_2)$$

well defined up to a positive scalar. It is continuous in compact parameter families, natural under bundle automorphisms, and associative for disjoint necks. On invertible operators it is the canonical identification.

Proof. Choose finite-dimensional stabilizations supported away from the ends so that the two stabilized operators are surjective. Patch their right inverses using cutoffs. The error tends to zero with the neck length, so a Neumann series corrects the patched operator to a right inverse. Projection of the patched kernels to the kernel of the glued stabilized operator gives an isomorphism for long necks. Its top exterior power and the stabilization isomorphisms define the asserted map.

Enlarging stabilizations commutes with this construction. Changing cutoffs, projections or neck length gives a path of kernel isomorphisms, so the sign does not change. For two disjoint necks use stabilizations supported away from both; the two successive patching maps are homotopic to the direct three-piece patching map. This proves associativity, with only the graded signs caused by reordering the factors. When both operators are invertible, no stabilization is needed and the map is $\mathbb{R} \otimes \mathbb{R} \rightarrow \mathbb{R}$. The gauge version is the determinant-line construction accompanying ordinary Floer gluing [Don02, Proposition 5.11]; the strip version uses the same stabilization argument for real Cauchy–Riemann operators as for the cylinders of [FH93], with the gluing of determinant lines along strip-like ends of [WW15a, Proposition 2.4.3]. $\square$

Gluing a constant cylinder is the positive identity: its operator is invertible, and the patched kernel map converges to the identity as the neck grows. Weighted gluing is treated by conjugating with the exponential weight. The resulting asymptotic operator must be invertible, and matched ends must use compatible weights. At a mixed end, ordinary gluing is applied together with the stationary coordinates, which are kept as separate finite-dimensional factors.

4. THE INTEGRAL GAUGE CONVENTION

4.1. Almost-complex orientations with free ends. Choose an S^1 -invariant almost-complex germ J_T near $S^1 \times N(T_0)$, with T_0 an almost-complex torus in its specified orientation. On an object, such a structure is equivalently specified by a nonvanishing vector field on $M^\#$, normalized near T_0 . Choose such a structure J_M^0 for each object, called the *reference almost-complex structure* of M . On a cobordism require J_W to agree with the fixed germ near $N(T_0) \times [0, 1]$ and to be translation invariant on its end collars. Its end restrictions J'_0, J'_1 are not prescribed in advance.

Lemma 4.1. *Every decorated cobordism admits an almost-complex structure of this kind.*

Proof. The fibre of the bundle of compatible oriented almost-complex structures is $\mathrm{SO}(4)/\mathrm{U}(2) \cong S^2$. Let $A = N(T_0) \times [0, 1]$. The primary relative obstruction lies in $H^3(W^\#, A; \mathbb{Z})$ and maps to $W_3(W^\#) = 0$. Here an oriented four-manifold admits a spin^c structure; equivalently its second Stiefel–Whitney class has an integral lift. The remaining relative class is in the image of the connecting map from $H^2(A; \mathbb{Z})$. The restriction $H^2(W^\#; \mathbb{Z}) \rightarrow H^2(A; \mathbb{Z})$ is surjective: the torus summand supplies a class pairing once with T_0 , and this class extends by the Mayer–Vietoris sequence across $S^2 \times [0, 1]$, along which the tube is attached, since it restricts to zero on the boundary sphere of the punctured torus summand. Thus the connecting map vanishes.

The secondary obstruction lies in $H^4(W^\#, A; \mathbb{Z})$. Since $A \simeq T^2$, its third cohomology vanishes; since $W^\#$ is connected with nonempty boundary, $H^4(W^\#; \mathbb{Z}) = 0$. The relative group therefore vanishes. Homotopy extension makes the structure equal to the prescribed germ, and a collar reparametrization makes it translation invariant near the ends. $\square$

The end restrictions are left free in Lemma 4.1, since prescribing both of them would change the obstruction problem. The passage to the reference structures of the objects is made by the conversion maps of Section 4.3.

The almost-complex deformation of the anti-self-duality operator has a complex-linear endpoint. The complex orientation of its kernel and cokernel orients its determinant line. This construction is invariant under the full group of bundle automorphisms, which act complex linearly at that endpoint [KM11b, Section 5.1]. Ordinary determinant gluing preserves these orientations. In particular the action of ϕ_M on the orientation lines has no additional sign twist.

4.2. Comparison of two almost-complex orientations. Let v, v' be the normalized vector fields corresponding to J, J' on $S^1 \times Y$, where $Y = M^\#$. Their difference class $\delta(v, v') \in H^2(Y; \mathbb{Z})$ satisfies

$$2\delta(v, v') = c_1(\xi_{v'}) - c_1(\xi_v).$$

The difference class is the obstruction to homotoping v to v' over the two-skeleton. When $H^2(Y; \mathbb{Z})$ has two-torsion, it is not determined by the displayed Chern-class difference.

Proposition 4.2. *If $\alpha \in H^1(Y; \mathbb{Z}/2)$ is the reduction of an integral class, the comparison of the J - and J' -orientations changes around a loop with lifting class α by*

$$(4.1) \quad \xi_{J, J'}(\alpha) = (-1)^{\langle \delta(v, v'), \text{PD } \alpha \rangle}.$$

In particular it is 1 on ϕ_M when J and J' agree near T_0 .

Proof. Represent the loop by a path of connections whose endpoints are identified by a bundle automorphism of class α . Closing up gives a bundle P_α on $S^1 \times Y$ with

$$w_2(P_\alpha) = \text{pr}_Y^* w_2(P) + [dt] \smile \alpha.$$

The determinant-line comparison for this closed operator is divided by the comparison for the constant loop. The P -independent orientation factor therefore cancels.

For completeness, the bundle-dependent comparison can be computed on a reducible model $\mathbb{R} \oplus L$. The anti-self-dual operator on the L -summand has both the complex structure induced by J and that of L . On the summand where these structures are opposite, conjugation changes the determinant orientation by $(-1)^{\text{ind}_\mathbb{C}}$. The index formula gives, for two almost-complex structures, the ratio

$$(-1)^{\langle c_1(L) \smile \delta(v, v'), [S^1 \times Y] \rangle}.$$

This is the relative form of the usual almost-complex orientation calculation [Don87, Section 3(c) and the proof of Proposition 3.25]. Changing the instanton number does not change the ratio: instanton addition preserves both complex orientations.

Because α has an integral lift, so does the displayed Stiefel–Whitney class. The reducible model therefore computes the desired ratio. Only the term $[dt] \smile \alpha$ survives after division by the constant-loop comparison; integration over the circle gives (4.1). The closing-up identification is made in pair form $\det D \otimes (\det D_{\text{const}})^*$, so it introduces no choice of orientation at the closed end. It is the real excision identification, independent of J and J' . Finally α_M is dual to T_0 , and the difference class restricts to zero there when the structures have the same germ. $\square$

We call $\xi_{J, J'}$ the *lifting-class sign* of the conversion from J to J' (Section 4.3): it is the sign by which the orientation of the interpolating operator changes around a loop with lifting class α . The lifting classes of the framed objects are reductions of integral classes, so Proposition 4.2 applies to every lifting class used below.

4.3. Conversion maps between almost-complex structures. A *conversion* is the isomorphism $C_G(Y, J) \rightarrow C_G(Y, J')$ between the gauge complexes oriented by two almost-complex structures, obtained by gluing to the orientation line of each generator an oriented operator that interpolates the two complex-linear operators; it compares the orientation lines of [KM11b, Section 5.1] for J and J' . For a critical point a , let Q_a interpolate the complex operators for J and J' . Choose a weight for which this operator is Fredholm and choose an orientation ϑ_a of its determinant line. Transport ϑ_a along paths of connections. Since J and J' agree near T_0 , the lifting-class sign $\xi_{J,J'}$ is trivial on ϕ_M by Proposition 4.2, so this transport is consistent on the framed quotient. It is unique up to one sign after a choice at one reference connection.

Let $D_{J,a}$ be the operator on $\mathbb{R} \times Y$ that interpolates the complex-linear operator for J and the operator at a ; by [KM11b, Section 5.1], the generator orientation line $L(a)$ of a in $C_G(Y, J)$ is $\det D_{J,a}$. Let

$$\gamma : \det D_{J,a} \otimes \det Q_a \longrightarrow \det D_{J',a}$$

be the determinant-line gluing of Lemma 3.3 along the complex-linear end, with the factors in the displayed order. We place ϑ_a after the generator orientation line and obtain

$$(4.2) \quad \Theta_{J \rightarrow J'}^r : C_G(Y, J) \longrightarrow C_G(Y, J'), \quad x \longmapsto \gamma(x \otimes \vartheta_a).$$

Its parity, written $|\Theta| = \text{ind } Q_a \pmod{2}$, is independent of a . Moving the odd trajectory line past ϑ_a shows that

$$d\Theta^r = (-1)^{|\Theta|} \Theta^r d.$$

Thus it is a closed isomorphism of degree $|\Theta|$. If Θ^1 denotes the convention with ϑ_a placed before the generator orientation line, then

$$(4.3) \quad \Theta^r = \Theta^1 \text{gr}^{|\Theta|}.$$

Composition of conversion maps agrees with the direct conversion up to one sign. The sign is the same for all generators, because the orientations ϑ_a are transported over the connected space of connections.

Definition 4.3. For (W, E, ν) choose a structure J_W as in Lemma 4.1, with end restrictions J'_0, J'_1 . Define

$$(4.4) \quad I_J^\#(W) = \Theta_{J'_1 \rightarrow J_{M_1}^0}^r I(W^\#, E^\#, J_W)^{\phi_W} \Theta_{J_{M_0}^0 \rightarrow J'_0}^r.$$

At the induced end structures the middle map uses the almost-complex orientations and the stated ordering of its determinant line. Tensor ends use (3.4).

Proposition 4.4. *The homotopy class of (4.4) is independent of J_W up to one sign. These classes define a projective gauge functor, and their reduction modulo two is the framed gauge functor of Paper I.*

Proof. Compare two choices of J_W after inserting their endpoint conversions. They orient the same real determinant line over the space of connections with fixed endpoint limits. Their ratio is continuous and takes values in $\{1, -1\}$, so it is constant on that space. Gluing a path at either endpoint shows that the ratio does not change when the limiting critical point is changed. All structures involved have the fixed germ near the odd torus, so the lifting-class sign of Proposition 4.2 is trivial on the lifting groups of the ends; this makes the assertion invariant under changing the lift of an end critical point. Thus the ratio is one sign for all entries of the map.

Change the metric and the secondary perturbations on $W^\#$ along a path, varying J_W continuously among compatible structures with the fixed germ. The almost-complex orientations then orient the parametrized moduli spaces of [KM11b, Section 3.9] continuously, and the oriented boundary of the one-dimensional family gives a chain homotopy between the two maps, up to one sign. A change of the Floer data on an object is the case of a product cobordism.

For composition choose an almost-complex structure on the composite as in Lemma 4.1, and make it cylindrical at the intermediate cut. The restrictions then have identical structures there, and the ordinary gauge composition theorem applies [KM11b, Proposition 5.4]. Its hypothesis is that the subgroup of $H^1(W; \mathbb{Z}/2)$ used on the composite contains the image of the Mayer–Vietoris connecting map $H^0(Y; \mathbb{Z}/2) \rightarrow H^1(W; \mathbb{Z}/2)$ of the cut Y . This image vanishes when Y is connected, and composition in $\mathbf{B}$ takes place at connected cuts. Insert the conversions to the reference structures of the objects. The two right-placed conversions at the common end are inverse conversions between the same two structures, so their composite is ± 1 . Independence of J_W gives the result for the original choices. The product cobordism gives the identity class. All orientation signs disappear modulo two. $\square$

4.4. Reflection, transposition and other orientation conventions. Choose the reference structures of M and $-M$ in pairs:

$$J_{-M}^0 = \Psi_* \rho_* J_M^0.$$

Here ρ reverses the cylindrical direction and Ψ is the fixed identification of the reflected torus-summed presentation with $(-M)^\#$. Take $T_0 = \{\theta_0\} \times T^2$ and the product germ there; require Ψ to preserve the odd torus, its germ, and its lifting class, and choose the pairing involutive, so that it returns J_M^0 when applied to J_{-M}^0 . Define $\Phi_{G,M}$ by the bent product cylinder: the product cobordism $[0, 1] \times M^\#$, with both ends bent to incoming ends, is a cobordism from $(-M)^\# \sqcup M^\#$ to the empty set. Since $H^1([0, 1] \times M^\#; \mathbb{Z}/2) = H^1(M^\#; \mathbb{Z}/2)$, its lifting group, generated by the class of the torus tube, restricts to the diagonal subgroup of $\phi_{-M} \times \phi_M$. Its count on $C_G(-M) \otimes C_G(M)$ is therefore the orbit count (D.2), summed over all lifts of the two end orbits. Its gauge map, with the ordered determinant gluing and the almost-complex orientations of the paired reference structures, is a pairing $\langle \cdot, \cdot \rangle_M : C_G(-M) \otimes C_G(M) \rightarrow \mathbb{Z}$, and we put $\Phi_{G,M}(y)(x) = (-1)^{|y|} \langle y, x \rangle_M$. The sign makes $\Phi_{G,M}$ a chain map: the pairing is closed for the right tensor differential (3.3), so $\langle y, dx \rangle_M = -(-1)^{|x|} \langle dy, x \rangle_M$, and with the dual differential this gives $d^\vee \Phi_{G,M} = \Phi_{G,M} d$. Its index-zero solutions are the constant ones, which pair a critical point with its reflection; their operators are invertible and their end lines even. Over a pair of orbits $([\bar{a}], [a])$ the two diagonal classes are those of $(\bar{a}, a)$ and $(\bar{a}, \alpha_M a)$. Since ϕ_M acts freely on the critical points and commutes with reflection, only the first contains a constant solution. In the bases of generators, matched by reflection, each entry of the pairing is therefore 0 or ± 1 , with exactly one nonzero entry in each row and column. Hence

$$(4.5) \quad \Phi_{G,M} : C_G(-M) \longrightarrow C_G(M)^\vee, \quad \Phi_{G,-M} = \Phi_{G,M}^\vee \text{ ev},$$

is an even isomorphism. For the identity, the bent cylinder of $-M$ is that of M with its two incoming ends exchanged, since the pairing of reference structures is involutive; exchanging the input lines of y and x costs $(-1)^{|x||y|}$, so $\langle x, y \rangle_{-M} = (-1)^{|x||y|} \langle y, x \rangle_M$. Since the pairing vanishes unless $|x| = |y|$, the two signs $(-1)^{|x|}$ and $(-1)^{|y|}$ in the definitions cancel, and this is the displayed identity.

For a morphism $(W, E, \nu) : M_0 \rightarrow M_1$, the *reversed morphism* $\overline{W} : -M_1 \rightarrow -M_0$ is the same oriented decorated manifold, read in the opposite direction.

Lemma 4.5 (Gauge transpose). *For every morphism $(W, E, \nu) : M_0 \rightarrow M_1$,*

$$I_J^\#(W)^\vee \Phi_{G, M_1} \simeq \pm \Phi_{G, M_0} I_J^\#(\overline{W}) \quad \text{in } \mathbf{K}(\mathbb{Z}).$$

The same identity holds for gauge continuation maps, with W the product cobordism carrying the two sets of data.

Proof. By Proposition 4.4, $I_J^\#(\overline{W})$ may be computed with the same J_W . Its end restrictions are then $\Psi_* \rho_* J'_1$ and $\Psi_* \rho_* J'_0$, and its conversions are the reflections of those of W . Reflection carries each operator Q_a to an operator of the same index, so the conversions of $\overline{W}$ have the parities $|\Theta_0|, |\Theta_1|$ of the conversions Θ_0, Θ_1 of W at its incoming and outgoing ends, and $I_J^\#(W)$ and $I_J^\#(\overline{W})$ both have degree $|\Theta_0| + |\Theta_1|$.

Let $y \in C_G(-M_1)$ and $x \in C_G(M_0)$. Gluing the bent cylinder of M_1 to $1 \otimes^r I_J^\#(W)$, and that of M_0 to $I_J^\#(\overline{W}) \otimes^r 1$, gives the same bent cobordism. It carries J_W , the paired reference structures at its ends, and a conversion at each end. The first composite orders the lines as $L(y), L(x), \vartheta_0, \det, \vartheta_1$ and the second as $L(y), L(x), \bar{\vartheta}_1, \det, \bar{\vartheta}_0$, where $\det$ is the line of the cobordism operator, even at the counted solutions, and $\vartheta_i, \bar{\vartheta}_i$ are the orientations of the conversion lines. The two orders differ by the constant $(-1)^{|\Theta_0|+|\Theta_1|}$, and each $\bar{\vartheta}_i$ differs from the transported ϑ_i by one fixed sign; the constant solutions of the bent cylinders have even lines. The rule (3.4) contributes $(-1)^{(|\Theta_0|+|\Theta_1|)|x|}$ to $(I_J^\#(\overline{W}) \otimes^r 1)(y \otimes x)$ and no sign to $(1 \otimes^r I_J^\#(W))(y \otimes x)$. Hence

$$\langle y, I_J^\#(W)x \rangle_{M_1} = \pm (-1)^{(|\Theta_0|+|\Theta_1|)|x|} \langle I_J^\#(\overline{W})y, x \rangle_{M_0},$$

with one sign for all x and y . Both sides vanish unless $|x| = |y| + |\Theta_0| + |\Theta_1|$. Evaluated on y and then on x , the two sides of the identity in the lemma give

$$(-1)^{(|\Theta_0|+|\Theta_1|)|y|+|y|} \langle y, I_J^\#(W)x \rangle_{M_1} \quad \text{and} \quad (-1)^{|y|+|\Theta_0|+|\Theta_1|} \langle I_J^\#(\overline{W})y, x \rangle_{M_0}.$$

These agree up to the one sign, because whenever the pairing is nonzero

$$(|\Theta_0| + |\Theta_1|)(|x| + |y|) \equiv |\Theta_0| + |\Theta_1| \pmod{2}.$$

Thus the degree-dependent reordering of an odd conversion line is absorbed by the graded dual and by the rule (3.4), and no grading operator remains. The product cobordism with varying data gives the identity for gauge continuation. $\square$

If the conversion lines are instead placed on the left, put $I_1(W) = \Theta_1^1 I(W) \Theta_0^1$. In the induced end gradings the unconverted gauge map $I(W)$ is even. If the two conversions have parities $|\Theta_0|, |\Theta_1|$, then

$$(4.6) \quad I_J^\#(W) = (-1)^{|\Theta_0|+|\Theta_1|} I_1(W) \text{gr}^{|\Theta_0|+|\Theta_1|}.$$

This is an identity of maps in the stated conventions. The factor $\text{gr}^{|\Theta_0|+|\Theta_1|}$ depends on the degree, so on an object whose homology has both parities the two maps differ by a grading operator and not by an overall sign.

A $\mathbb{Z}/4$ grading compatible with all the orientation data would provide a further regrading $\eta(x) = (-1)^{g(x)(g(x)-1)/2} x$. The elementary identity

$$\eta(g+d)\eta(g)^{-1} = (-1)^{dg+d(d-1)/2}$$

explains how an odd grading operator can then be converted into a conjugation and one sign per cobordism map. The comparison of such a lift with the grading data of Scaduto [Sca15] is left open. Theorem A is therefore stated for $I_J^\#$, and the mod-two equality of the two gauge functors is not used to identify the integral maps.

5. DETERMINANT LINES ON MIXED CONFIGURATION SPACES

This section orients the mixed moduli spaces over the integers. The orientation of a configuration is obtained by deforming it, with the shifting of Proposition 2.2, to a gauge path followed by an invertible constant pair, and transporting back the almost-complex orientation of the path tensored with the canonical orientation of the constant pair. Proposition 5.4 shows that this is independent of the deformation for domains with one gauge input; Definition 5.5 gives the corresponding construction for triangles and joined torus domains. The first two subsections fix the operators whose determinant lines are oriented and show that the orientations of solutions do not depend on the choices made in realizing them, including the end weights.

5.1. Realizations of the linearized mixed operator. Let c be a smooth configuration on a mixed domain, with its specified limits. Denote the mixed equation by $\mathcal{F}(c) = 0$. A chart Φ centred at c incorporates the exact matching condition and the normal seam condition. Its differential identifies a fixed linear boundary problem with the tangent space to the configuration space.

Definition 5.1. A permitted realization of the linearized problem at c is an operator of the form

$$(5.1) \quad L_c = (\mathcal{S} d\mathcal{F}_c, \lambda^2 d_{A'}^{*,g'} \pi_A) d\Phi_0.$$

It acts on the prescribed weighted spaces with stationary mixed-end coordinates included. Here π_A denotes the connection component of a tangent vector to the configuration space. At mixed ends we use spatially Coulomb stationary tangent representatives and matching lifts whose remaining divergence decays, so that the scalar row takes values in the weighted target. Both rows are transported by the chart differential. A change of stationary lifts likewise transports their finite scalar columns as well as the equation columns. The map $\mathcal{S}$ is a continuously chosen invertible target identification, equal to the identity outside a compact set. The slice centre A' is a smooth connection differing from the connection in c only on a compact core and sufficiently little in C^0 . The metric is $g' = f^2 g$, with $f = 1$ off the collapse region and f depending only on the coordinates of the planar factor near a seam, and $\lambda > 0$ equals one outside a compact set. The chart has the same end and seam conventions as the original mixed problem. Changes of all these choices are taken inside their fixed components. In a comparison of weights, the definition also requires that the full stationary and parameter columns belong continuously to every intervening weighted target.

The last condition is part of the definition because a configuration which is not a solution may decay more slowly than the weights of a spectral gap allow. The restriction on the target identification prevents an arbitrary reversal of orientation from being included as a change of realization. The reference identification is the one induced by the geometric charts, and its permitted component is fixed before orientations are transported.

For the collapse contraction we also use the extended realization of Proposition 7.4. Its centre and geometric target identification can have stationary limits different from those of the configuration and the identity, respectively. Its source representatives are Coulomb for the

varying centre, both rows are transported, and differentiation of the moving equation target contributes the residual term specified there. These extended realizations form a separate family at a fixed positive scale; they are not compactly supported changes in the sense of Definition 5.1.

Lemma 5.2 (Realization comparison). *Under the Fredholm and continuity statements of Proposition 2.1(a), permitted realizations at a fixed configuration have canonically identified orientation double covers. The identification is continuous in compact families and commutes with ordinary gluing. At a regular zero of index zero it preserves the canonical determinant orientation.*

Proof. Changes of charts act on the entire operator in (5.1). Together with target identifications, they are bounded invertible changes of source and target and induce the standard naturality maps on determinant lines. An interval of weights permitted by Definition 5.1 gives a continuous Fredholm path after conjugation to fixed spaces, and hence an identification of orientation covers. At a solution, improved asymptotic decay identifies the kernels, cokernels and common stabilized determinant lines throughout the permitted gap. For a change of slice centre at a solution, split the tangent space into infinitesimal gauge directions and a fixed complement. Gauge invariance gives $d\mathcal{F}_c dA_c = 0$. The gauge-fixed operator is then block triangular. Its gauge block is invertible by the Poincaré estimate and the smallness of $A' - A_c$; its other block is the intrinsic linearized equation on the quotient. Thus invertibility at a regular zero of index zero is unchanged.

Conformal changes in the specified unperturbed region change neither the anti-self-duality equation in four dimensions nor the surface Cauchy–Riemann equation or the normal seam subspace. The same block argument applies to the gauge row. Positive row multiplication preserves orientations. These paths give the comparison. A homotopy between two such paths gives an orientation cover over a square, so its endpoint identification is unchanged. The gluing statement follows by choosing the changes away from the ordinary gluing collars and then using continuity. At a regular zero of index zero every operator in these paths is invertible, so the canonical orientation is preserved. $\square$

The lemma uses the Fredholm families of Proposition 2.1(a), including those on the completed domains. Two situations are treated separately. At a configuration which is not a solution, the block-triangular argument above does not apply to a change of Coulomb centre; where such a configuration occurs, in the collapse argument, the two-parameter determinant-line comparison of Proposition 7.4 is used instead. A reference path in a small weighted space may fail to lie in a larger weighted space. For weight comparisons of its endpoint orientations, truncate its decaying coordinates on a common far end while retaining the exact stationary family. The truncation homotopy is taken at the initial small weight. The truncated path, together with the endpoint truncation segments at solutions or stationary references, admits the larger weights. Determinant transport around this homotopy gives the comparison; the decay of the original path is left unchanged.

5.2. Independence of the end weights. Reverse incoming ends to outward coordinates. Write an autonomous end operator as $\partial_s - \mathcal{A}_e$. Its stationary kernel is the primary tangent space on a mixed end; at an ordinary end it is zero. Choose $\gamma_e, \gamma_e^{\text{opp}} > 0$ so that the negative spectrum is at most $-\gamma_e$ and the positive spectrum is at least γ_e^{opp} ; they need not be equal. The permitted positive primal weights satisfy $r_e < \gamma_e$, and the formal adjoint is paired at the dual weight $-r_e$. Mixed-end stationary coordinates remain a separate finite summand.

Proposition 5.3. *Consider the mixed domains and ordinary strips used here, with a weight in the open gap above at each autonomous end, or the weight zero at an ordinary end. Solutions, the gauge transformations between them with prescribed stationary limits, and the determinant lines of the linearized operators at solutions do not depend on the choice of weights. The identifications are transitive, respect the lifting actions and ordered gluing, and give the same oriented counts and homotopies whenever their small-weight compactifications have been established.*

The proof, in Appendix A.10, compares decay rates, kernels, cokernels and stabilized determinant lines across the weights by the standard arguments for asymptotically translation-invariant operators with a spectral gap. Two points are specific to mixed ends. The component of a solution in the stationary kernel of $\mathcal{A}_e$ is integrated from infinity and uses the prescribed limit in the stationary critical family, which is carried as a separate finite summand at every weight. The component in the positive spectral subspace is also integrated from infinity, so decay at every rate below γ_e holds without any restriction from the opposite gap γ_e^{opp} .

Proposition 5.3 concerns solutions. A configuration which is not a solution and decays slowly need not lie in a space with a larger weight; reference paths are compared by the truncation described after Lemma 5.2.

5.3. Orientations transported from constant pairs. A constant pair C_x joins the gauge critical point a_x to the symplectic intersection point x . Its operator is invertible. Let $L_S(x) = L_G(a_x)$. More generally, a constant pair C^c may be built from any representative c of the critical point over x (Appendix B). It is equivalent to C_x under the lifting group, by a gauge transformation when c lies in the orbit of a_x . Its operator is invertible at the data in use, which are small, by the invertibility clause of Proposition 2.1(b). Its positive element corresponds to that of C_x under conjugation, and the almost-complex orientation identifies $L_G(c)$ with $L_G(a_x)$, compatibly with transport (Appendix B). For a mixed configuration c with gauge input a and symplectic output x , shifting and gluing identify it, through a path of configurations, with a configuration of the form $A\#C^{c'}$, where A is a gauge path from a to a representative c' . The configurations of this form make up the *terminal subspace* $\mathcal{T}_x$ of Appendix B: they agree with a constant pair outside the gauge input end.

Proposition 5.4. *Let c be a configuration on a one-input domain of Proposition 2.2(a), with gauge input a and output x . There is a relative orientation isomorphism*

$$(5.2) \quad \Gamma_c : L_G(a) \otimes \det D_c \longrightarrow L_S(x),$$

defined up to a positive scalar, which at a configuration $A\#C^{c'}$ of the terminal subspace is the almost-complex orientation of A followed by the canonical orientation of the invertible constant pair. It is independent of the shifting path, natural under the lifting group, and coherent under ordinary gluing of gauge paths at the input and of strips at the output.

Proof. The construction follows [DFL21, Section 4.7], where the orientation is fixed by gluing to an invertible constant pair. Here it is carried out on configurations rather than gauge classes, relative to the terminal subspace, using Propositions 2.1 and 2.2. On $\mathcal{T}_x$ a common gluing length identifies the determinant line of $A\#C^c$ with the gauge line of A tensored with the determinant line of the invertible constant pair. The gauge factor is oriented by the almost-complex construction and the second factor canonically, so this terminal orientation is continuous on compact families in $\mathcal{T}_x$.

Let p_1, p_2 be two shifting paths from c , ending at terminal configurations t_1, t_2 . Apply the relative family shifting of Proposition 2.2(a) to the compact one-parameter family $p_1^{-1}p_2$ from t_1 to t_2 . Its relative clause keeps the tracks of the two endpoints in the terminal subspace, and the family ends in it. This gives a square whose top edge and two sides lie in the terminal subspace, where the terminal orientation is continuous. Since the square is contractible, transport along p_1 and along p_2 gives the same orientation at c . In particular t_1 and t_2 are joined in the terminal subspace. For a loop ℓ at c , taking $p_2 = \ell \cdot p_1$ shows that the orientation monodromy of a loop with fixed limits is trivial. This proves path independence of the orientation transported from the constant reference.

For a strip attached at the output, first shift the original mixed configuration to $A\#C^c$. Then shift C^c together with the attached strip to another gauge path followed by $C^{c'}$, with c' a representative over y . Associativity of determinant gluing identifies both orders of this construction with gauge gluing at the final reference. Continuity transports that equality back to the original configuration. The gauge-end statement is the same argument with the extra path attached before shifting.

The lifting automorphism preserves gauge orientations, conjugates the constant-pair operator, and commutes with gluing. Conjugation preserves the canonical orientation of an invertible operator. Hence the construction descends to the framed quotient. If shifting reaches the other lift of a_x , the two end lines are identified by this same action. $\square$

The pure strip orientation is defined by applying (5.2) to $C_x\#u$ and cancelling the constant pair. This is the transported orientation construction of [DFL21, Section 4.7].

Domains with two gauge inputs are oriented in the same way, after the output region has been separated from the inputs by a closed cut. For a mixed triangle, the closed three-manifold that separates the two gauge inputs from the output (Paper I, Lemma 9.11) writes its gauge region as an ordinary two-input cobordism followed by a truncated output DFL domain. For the folded triangle the cut is Z and the compression-body arms of the output domain are D_0, D_1 (Section 2.5). A configuration is not required to be critical at this finite cut. On a compact family at positive collapse scales, the cut first receives a product collar. This is done by a change of metric at fixed configuration, on compact ranges of positive scales, which raises the compression scale near the cut to one and makes the metric a product there (I11).

The finite-width joined torus domains W_λ^J of Section 2.5 have the same structure, with no collapse scale. By (I2), the closed cut $Z_{\text{cut}} \cong Z$ decomposes W_λ^J into the reverse excision cobordism, which is the ordinary two-input cobordism, and the closed-cut piece M_λ , a truncated output DFL region for the pair (A_0, A_1) with compression-body arms $D_{0,\lambda}, D_{1,\lambda}$. After the compactly supported metric change of (I11) near the arcs of the cut away from the bridge ends and the matching lines, the cut has a product collar. The configuration spaces do not depend on the metric (Appendix B.1), so the change is a transport along a path of data. The construction below applies verbatim with the output region M_λ and the cut Z_{cut} , on compact ranges of widths.

Prepare the output DFL region and extend it positively, as in the proof of [DFL21, Lemma 4.62]. Then shift the output region far along the cut collar until it becomes a constant mixed pair, leaving the two-input gauge cobordism unchanged; Appendix B.5 gives the formulas. We call this deformation the *localized shift*. It ends at a configuration $A_E\#C^c$: a constant mixed pair C^c , with c a representative over the output label y , glued at a parallel copy of the cut inside its product collar to a configuration A_E on the ordinary two-input cobordism with the original inputs.

At the end of the deformation remove the compact nonprimary terms whose dependence or output crosses the cut. On the constant mixed factor use its zero-secondary reference datum, retaining the fixed compatible end structures. These are homotopies of Fredholm data supported in a compact set, used only to transport orientations. Ordinary determinant gluing then gives the ordinary cobordism line tensored with the canonically oriented constant-pair line.

Definition 5.5. Let c be a configuration on a mixed triangle or a folded triangle at a positive collapse scale, or on a finite-width joined torus domain W_λ^J . A configuration is *terminal* if it agrees with a constant pair $C^{c'}$ beyond a parallel copy of the cut in its product collar and on a two-sided collar of that copy; the configurations $A_E \# C^{c'}$ above are terminal. The *reference orientation* at a terminal configuration is the almost-complex orientation of the ordinary cobordism line of A_E , with the input lines in cyclic order, tensored with the canonical orientation of the invertible constant-pair line through ordinary determinant gluing at the copy of the cut, after the compact data homotopy above. The *orientation at positive scale* of c is the isomorphism obtained by transporting the reference orientation back to c along the localized shift. Its source is the tensor product, in this order, of the input lines, the conversion line at the cut when one is present, and $\det D_c$; this placement of the conversion line is Convention 8.1, stated at the beginning of Section 8. Its target is the output line. On the joined torus domains the reference orientation is given explicitly by (8.3).

The relative form of the localized shift (Proposition 2.2(b)) keeps terminal configurations terminal. The square argument in the proof of Proposition 5.4 therefore shows that the orientation of Definition 5.5 is independent of the deformation and coherent with ordinary gluing. Associativity is applied before any factor is reordered.

5.4. Boundary orientations of one-parameter families. A parametrized family is oriented by placing the oriented parameter space before the determinant line of the vertical operator. At a regular point this induces the usual orientation of the finite-dimensional zero set. To see continuity when the vertical kernel dimension changes, choose one stabilization on a neighbourhood in the parameter space. The orientation is then the exterior-power orientation of the stabilized kernel with the fixed stabilization factor cancelled. Transition maps between such neighbourhoods are the determinant-bundle transition maps, so no basis of the moving kernel is chosen.

For a one-parameter regular family, the boundary orientation is outward normal first. At a strip break the translation vector is the positive vector $\partial_s u$. Moving its odd line past the preceding generator factors gives exactly the signs in (3.3). At a parameter endpoint the two outward normals are opposite. It follows that a compactified one-dimensional family yields an identity of the form

$$(5.3) \quad F_1 - F_0 = \delta K,$$

with the Koszul signs already included in δ . If an orientation for the whole family has been chosen only up to sign, the corresponding endpoint identity has one overall sign, not independently chosen signs for its separate matrix entries.

The compatibility of a geometric boundary collar with this determinant-line orientation rests on ordinary linear gluing at the end where the break occurs (Lemma 3.3) and on the direction of the translation vector along the gluing collar; the computation is carried out in Lemma 8.12. This comparison of collars, and not only an orientation of the interior of the moduli space, is what (5.3) requires.

6. THE MIXED ISOMORPHISM OVER THE INTEGERS

6.1. Differentials and the mixed map. For an index-one gauge trajectory or holomorphic strip, its sign is positive when the relative orientation agrees with its translation vector. For a regular index-zero mixed solution, its sign is the comparison of (5.2) with the canonical element of its invertible determinant line. Set

$$(6.1) \quad \begin{aligned} d_G o_a &= \sum_b \sum_{[A] \in \mathcal{M}_G^0(a,b)} \text{sign}(A) o_b, \\ d_S o_x &= \sum_y \sum_{[u] \in \mathcal{M}_S^0(x,y)} \text{sign}(u) o_y, \\ N_p o_a &= \sum_x \sum_{[c] \in \mathcal{M}_N^0(a,x)} \text{sign}(c) o_x. \end{aligned}$$

The superscript 0 denotes dimension after quotienting translation when that action is present. All sums are finite by the compactness and energy-index bounds of Paper I.

Proposition 6.1. *The transported orientations give $d_S^2 = 0$ and $d_S N_p = N_p d_G$ over $\mathbb{Z}$.*

Proof. The compactified one-dimensional strip space has two-level broken strips as its boundary. Coherence in Proposition 5.4 identifies their boundary orientation with the product of the two zero-dimensional signs, with the sign from the translation factor. The oriented boundary count is zero, yielding $d_S^2 = 0$. For the mixed domain defining N_p , the two types of boundary are a gauge trajectory followed by a mixed solution and a mixed solution followed by a strip. The outward-normal orientations of the two ends give opposite signs. Their sum is $d_S N_p - N_p d_G$. The index and energy restrictions exclude additional boundary strata in this dimension. $\square$

6.2. The unit diagonal. The filtration used in Paper I orders the finite generator set so that every index-zero mixed solution goes from a generator to a generator of no larger filtration. At equal filtration the only solutions are the unique small deformations of the constant pairs. This is a statement about the set of solutions, before orienting it.

Lemma 6.2. *Let F be an integral endomorphism of a finite filtered free module. (For the mixed map, the source and target bases are identified through $L_S(x) = L_G(a_x)$.) Suppose its associated graded map is diagonal with each diagonal entry in $\{1, -1\}$. Then F is an integral isomorphism. Its inverse preserves the filtration.*

Proof. Write $F = D(1+U)$, with D the diagonal sign matrix and U strictly filtration-lowering. If there are r filtration steps then $U^r = 0$, and

$$F^{-1} = (1 - U + U^2 - \dots + (-1)^{r-1} U^{r-1}) D^{-1}.$$

All coefficients are integral. $\square$

Each diagonal solution is one regular point and contributes 1 or -1 , so Lemma 6.2 already gives the invertibility of N_p . With the normalization by constant pairs the diagonal entries are $+1$.

Proposition 6.3. *The map N_p is a chain isomorphism. Its diagonal entries are $+1$ in the constant-pair normalization. Its homotopy class is independent of regularizing data, after the corresponding gauge continuation is inserted; for a change of the ordinary strip structure the symplectic continuation map is inserted at the output.*

Proof. Take the segment from zero regularizing data to the chosen small data. At a constant pair the linearization is invertible. The chart implicit-function theorem therefore supplies a nearby branch of regular solutions. The uniqueness statement in the triangularity argument identifies this branch with the diagonal solution. The charts and the continuous Fredholm families of Proposition 2.1(a) make the branch, which is continuous in the chart norm, continuous for determinant-line transport; the constant pair is invertible by Lemma A.2. Its sign is locally constant and begins at the canonical sign $+1$.

The filtration now gives an inverse integral matrix. Since N_p commutes with the differentials, its inverse does too. At fixed ordinary ends, a regular compact-data homotopy gives (5.3). A change of the gauge-end function is treated by the inserted-cylinder comparison in Proposition 6.4(i), which retains the ordinary continuation map, and a change of the ordinary strip structure by part (ii) of that proposition. This proves the stated independence. $\square$

The normalization uses two properties of the segment. The secondary perturbations are supported away from the strip end, so the strip operator is constant along it. The segment lies in a star-shaped neighbourhood of zero on which all reference constant-pair operators are invertible. A general path of perturbations has neither property, which is why the segment is used.

6.3. Changes of the ordinary end data. A change of a far-end equation is not compactly supported on the mixed domain. It is compared by inserting the ordinary continuation domain.

Proposition 6.4. *Fix primary data and the mixed domain defining a mixed map, such as the domain defining N_p , or the domain of the pair (A_0, A_1) with gauge input Z at a fixed positive width. Its mixed maps satisfy the following integral chain homotopies.*

- (i) *If f_0, f_1 are two sufficiently small invariant gauge trajectory perturbations vanishing near the whole primary critical set, with the same symplectic end datum, then*

$$N_1 \text{cont}_G(f_0, f_1) - N_0 = d_S H_G + H_G d_{G,0}.$$

- (ii) *If the gauge end is fixed and J_0, J_1 are regular ordinary strip structures on the same ordered primary pair, then*

$$\text{cont}_S(J_0, J_1) N_0 - N_1 = d_{S,1} H_S + H_S d_G.$$

For the pair (A_0, A_1) the output pair is the product pair and the structures and continuation are split.

The homotopies H_G and H_S count the isolated points of parametrized moduli spaces of vertical index -1 .

The proposition concerns changes of the ordinary end data at fixed primary data. Changes of the primary data, for presentations in one small ball, are treated in Proposition 13.1.

Proof. This is the usual neck-stretching argument for continuation maps; we record the points specific to the mixed equation. The argument uses Propositions 2.1 and 2.2, and the neck gluing and families results recalled in (I7).

The stretched domains. For part (i), let β rise from zero to one on $[-1, 1]$, and let χ be the cutoff of the end perturbation in the mixed equation, equal to one for $s \leq -3$ and zero for $s \geq -2$. On the incoming gauge end use the end perturbation

$$q_T(s) = \chi(s)(f_0 + \beta(s+T)(f_1 - f_0)), \quad T \geq 5.$$

It equals f_0 on the far end and f_1 on the long cylinder between the transition band and the mixed piece. Translate a fixed regular ordinary continuation perturbation with the β band, and keep the secondary data of N_1 on the mixed piece. For large T the supports of these two perturbations, and the regions on which they depend, are disjoint. The continuation piece is then a fixed nonautonomous cylinder, so a rigid continuation trajectory has index zero and is not divided by translation. For part (ii), choose a continuation structure $J_{01}(\sigma, t)$ from J_0 to J_1 , with zero Hamiltonian and constant end collars. On the outgoing strip of N_0 , keep its original structure until $s = T - 1$ and use $J_{01}(s - T, t)$ afterwards. The far output equation is J_1 for every finite T , and the gauge end and the stationary mixed ends do not change. In both parts the far-end equations are fixed, and the derivative of the data in T is supported in the translated transition band. In coordinates in which the two pieces are fixed, the stretched domain is their gluing along a neck of length about T that carries the fixed autonomous equation; this is the setting of (I7).

Step 1: energy and critical points. The smallness neighbourhood for f_0, f_1 is chosen so that no new critical points occur: otherwise compactness of the perturbed critical equation would give a critical point of the functional with the primary perturbation outside the neighbourhood of the critical set on which both functions vanish. The two end perturbations therefore have the same critical points and Hessians. In part (i) the energy identity for the modified curvature is

$$E_T = 8\pi^2\kappa + 2\|\eta_T(A)\|^2 + 2 \int \partial_s q_T(s)(A_s) ds, \quad \partial_s q_T = \beta'(s + T)(f_1 - f_0) + \chi'(s)f_1.$$

where κ is the topological energy and η_T the secondary perturbation of the stretched domain. The last integral is bounded in absolute value by $\|f_1 - f_0\|_\infty + \|f_1\|_\infty$, and the translated secondary perturbations have uniformly bounded norm. The gradient of the autonomous function f_1 on the growing middle cylinder is part of the modified curvature, so it is not estimated as a compactly supported source. In part (ii) the Hamiltonian is zero, so the energy of a section is twice its symplectic area even for domain-dependent J , with no term involving the derivative of J , and the transition adds no curvature or boundary term. In both parts, configurations with the same intermediate critical lift have cancelling actions at the neck, and ordinary determinant gluing makes the index additive. The relation between index and topological energy therefore gives an energy bound and a lower bound for the index that do not depend on T .

Step 2: transversality. Regularity is needed for the stretched domains and for the interpolating families used below. In these families the gauge perturbations vary with fixed compactly supported dependence and output, and the strip structures vary on a compact set. The far ends and the collars of the genus matching surfaces stay fixed. Coefficients at an ordinary Lagrangian boundary component vary through a compact family of compatible structures with the same totally real boundary condition, and on a product target the structures stay split. A change of the principal coefficients of the section operator on a compact set is not a relatively compact perturbation. Instead, local elliptic estimates, a finite cover of the parameter space and the fixed parametrices on the ends give Fredholm estimates with a compact remainder. The local estimates for the operator and for its weak adjoint vary continuously with the coefficients, and both rows and all stationary columns of the operator are transported.

For mixed solutions, with the strip structure prescribed, invariant holonomy perturbations supported in a fixed compact set detect the whole cokernel, as in Paper I, Theorem 5.6(ii). The scalar component of a cokernel element vanishes because the scalar Neumann operator

is invertible. The gauge component vanishes on the support region, because it pairs to zero with the holonomy perturbations along loops of zero lifting class supported there. Every perturbation term has its output in that region, so the adjoint equation is local outside it, and unique continuation on the connected gauge region gives vanishing everywhere. The section component vanishes by Green’s formula on a genus matching arc, extension by zero and unique continuation, applied on each component of the section. This uses that every component of the section meets a genus matching arc and that the gauge region is connected, which holds for the domains of the proposition.

For continuation strips with domain-dependent J in the class above, let u be a nonconstant solution. The set where $u_t \neq 0$ is dense in any region where J may vary: if u_t vanished on an open set, so would u_s , and unique continuation would make u constant. If a cokernel element of the parametrized operator is nonzero, choose a point of that set at which it is also nonzero (its zero set has empty interior, by unique continuation for the adjoint equation). There the variations $Y \mapsto Yu_t$ of compatible structures span the tangent space of the target, contradicting the pairing with the cokernel element. Because the variations are localized in the domain, the strip need not be injective. At a constant solution, the energy identity for the linearized operator on the fixed tangent Lagrangians makes the decaying kernel zero. Its index is zero by deformation to an autonomous constant strip, during which the asymptotic operators stay invertible by transversality. For product targets use independent variations on each nonconstant factor; each constant factor is invertible. This direct transversality argument for the linearized strip operator supplies the continuation-strip input in the proof of Proposition 2.1(c).

We choose perturbations inductively in the index, starting from the lowest index allowed by the energy bound. At each stage the solutions of smaller index are already regular, those of negative index, counting the parameter, are absent, and outside a neighbourhood of the already regular broken configurations the solutions of the next index form a compact set, and a finite-dimensional space of perturbations makes it regular. This space is then fixed, and later choices are small enough to preserve the earlier regularity and the end collars. The Sard–Smale theorem for parametrized families then gives data that are regular for parametrized families through index one, with the data parameter counted once. In the parametrized operator $(v, \tau) \mapsto D_b v + \tau C_b$ the single column C_b contains $(\partial_b J_b)u_t$, the transport of the target and the variation of the gauge perturbations; ordinary gluing copies this column once, onto the piece that carries it. Monotonicity and the exclusion of bubbles, each of which costs index four, make the regular continuation spaces compact up to breaking through index one.

Step 3: compactness and gluing. For large T , the results of (I7) apply to the two pieces: the continuation piece and the mixed piece of N_1 in part (i), and the mixed piece of N_0 and the continuation strip in part (ii). A limit of solutions of index d consists of the two pieces and broken trajectories W_j along the neck, with

$$d - I_1 - I_2 - \sum_j \text{ind } D_{W_j} = 4n, \quad E_{\text{lost}} = 4\pi^2 n, \quad n \geq 0,$$

where I_1, I_2 are the indices on the two pieces; the continuation piece contributes its full index. Both pieces are regular, so their indices are nonnegative, and every nonconstant trajectory on the neck has positive index. Hence for large T there are no solutions of negative index, simultaneously for all labels by the uniform lower index bound. At $d = 0$ both pieces have index zero, there are no broken trajectories and $n = 0$. The continuation pieces keep their full transition data in the limit, and no energy term involving the derivative of J arises. By (I7),

gluing is then a bijection from the matched pairs of rigid solutions on the two pieces onto the rigid solutions of the stretched domain, with one length threshold for the finitely many pairs.

The quotient by the lifting group is kept in the gauge construction. The connected cylinder adds no lifting class, since the connecting image from H^0 of the connected cut vanishes [KM11b, Proposition 5.4]; its bundle automorphisms identify the lifting groups of its two ends, and the fibre product over the common critical orbit gives exactly the multiplicity of matrix composition. A strip adds no gauge region and no lifting factor. In both parts the orientations are compared by ordinary determinant gluing (Lemma 3.3) before the line of the intermediate generator is cancelled. For a nonautonomous continuation, the full stabilized operator is kept on its compact part, with its two prescribed end models, and the autonomous inverse is used only on the matched tail. The gluing exact sequence and its independence of the cutoffs then apply with both far ends fixed, without replacing the continuation by an autonomous trajectory. Relative mixed shifting (Propositions 2.2 and 5.4) and associativity preserve the constant-pair normalization, and the constant strip carries its canonical orientation. All maps here are even. Thus for large T the count F_T of rigid solutions on the stretched domain is

$$F_T = N_1 \text{cont}_G(f_0, f_1) \quad \text{in (i),} \quad F_T = \text{cont}_S(J_0, J_1)N_0 \quad \text{in (ii),}$$

with no independent sign for a component or a matrix entry.

The homotopy. Fix such a length T_* . If necessary, make an arbitrarily small compactly supported perturbation of the endpoint data that makes index one regular; by compactness and local uniqueness it does not change the regular rigid count. In part (i), interpolate the inserted data to those of N_0 by replacing β with $b\beta$, $b \in [0, 1]$, and interpolating the compactly supported perturbations; the far input stays f_0 . In part (ii), interpolate to N_1 through compactly supported changes of the strip structure and of the gauge perturbations, with far output J_1 . Along either path the far ends are fixed and the data change in a compact set, so the parametrized moduli spaces are regular through index one by the transversality above. Compactness holds with the coefficient family on the compact part, with index loss in nonnegative multiples of four and positive indices of ordinary breaks. As in (I7), the compactified one-dimensional parametrized space has as boundary only the two endpoint counts and breaks along index-one trajectories at ordinary ends, with collars given by gluing at the ordinary end. With the parameter oriented first, (5.3) gives the displayed identities, with H_G or H_S the count of parametrized solutions of vertical index -1 ; this holds also when the two data differ in the strip structure on a compact set. Finally Proposition 5.3 extends these counts from their common small weight to all permitted weights; slow reference paths are prepared before their weights are compared. $\square$

6.4. Orientations in families of data. The same reasoning applies to compact parameter families of primary and geometric data whenever their asymptotic intersections remain transverse. Locally choose fixed-centre charts and stabilizations. A nearby regular solution has small coordinates in such a chart. The line orientation pulled back from Proposition 5.4 varies continuously, by Proposition 2.2 for families with varying data, while the canonical orientation of an invertible operator is continuous on the invertible locus. Their ratio is consequently locally constant. This fact is used in the collapse and closed-cut comparisons; it is a local statement in the parameter and uses only these charts.

7. TRIANGLES AND ORIENTATIONS UNDER COLLAPSE

7.1. Two-handle data and the triangle product. Let $W_2: M_0 \rightarrow M_1$ be a two-handle cobordism, attached along a framed link with n components, with the bundle trivial at both ends and a filling frame as in Paper I, Section 4.4. Following Paper I, Section 9.1, choose a Heegaard triple $(\Sigma; H_0, H_1, H_2)$ of genus g subordinate to the link. Then $H_0 \cup H_1$ and $H_0 \cup H_2$ are Heegaard splittings of M_0 and M_1 , and $Y_{12} = H_1 \cup H_2$ is a connected sum of $r = g - n$ copies of $S^1 \times S^2$. Let C_i be the torus compression body of H_i (Section 2), so that $h = g + 1$, and write $Y_{ij}^\# = C_i \cup (-C_j)$; thus $Y_{01}^\# = M_0^\#$ and $Y_{02}^\# = M_1^\#$. The three Lagrangians in $\mathbf{M}_h$ are

$$P_0, \quad P_1 = L(C_1), \quad P_2 = L_{\mathfrak{h}_2}(C_2),$$

where P_0 is the image of $L_{\mathfrak{h}_0}(C_0)$ under the symplectomorphism of $\mathbf{M}_h$ induced by the filling frame. The primary perturbation of C_1 is zero. The Lagrangians $L(C_1)$ and $L(C_2)$ meet cleanly in $Q \cong \mathrm{SU}(2)^r$, and $\mathfrak{h}_2$ is a small positive multiple of the height function ζ_Q of Paper I, Section 9.2, a Morse function on Q with a unique maximum q_* . The perturbation $\mathfrak{h}_0$ is small, with P_0 transverse to P_1 and P_2 and disjoint from $P_1 \cap P_2$. The point q_* is the top generator of $CF(P_1, P_2)$, and the corresponding critical point of $Y_{12}^\#$ is again written q_* .

Write N_{ij} for the mixed map of the pair (P_i, P_j) , from the gauge complex of $Y_{ij}^\#$ to $C_{ij} = CF(P_i, P_j)$. Let F_r be the filling of $Y_{12}^\#$ by the boundary sum of r copies of $S^1 \times D^3$ and one copy of $T^2 \times D^2$. Its relative invariant θ_G is the signed count of index-zero solutions on F_r with cylindrical end $Y_{12}^\#$, projected to the framed complex; it is a cycle, equal to $\pm o_{q_*}$ (Paper I, Proposition 9.4, and Lemma 7.7). Put $\theta_S = N_{12}\theta_G$. The *torus-summed triangle* $X^\#$ is the union of $\Delta \times \Sigma_h$ and $\Delta \times T^2$, for a triangle Δ with edges e_0, e_1, e_2 , with the pieces $e_i \times C_i$ attached along $e_i \times (\Sigma_h \sqcup T^2)$. Its ends are $Y_{01}^\#, Y_{12}^\#$ and $Y_{02}^\#$, and $W_2^\# = X^\# \cup_{Y_{12}^\#} F_r$ (Paper I, Section 9.3). Its lifting group ϕ_X is the quotient of $(\mathbb{Z}/2)^3$, with one coordinate for each compression piece, by the diagonal; the class it restricts to on $Y_{ij}^\#$ is the sum of the i -th and j -th coordinates, written s_{ij} . Let $\mu_G: C_G(Y_{01}^\#) \otimes^r C_G(Y_{12}^\#) \rightarrow C_G(Y_{02}^\#)$ be the gauge map of $X^\#$ with lifting group ϕ_X . The symplectic map of the two-handle is

$$S_{W_2} = \mu_S(\cdot \otimes \theta_S),$$

where μ_S is the triangle product of (P_0, P_1, P_2) defined below, composed on both sides with the frame transports of Paper I, Theorem 9.13, which we suppress from the notation. Theorem 7.8 compares S_{W_2} with the gauge map of W_2 .

Order three Lagrangians L_0, L_1, L_2 by the boundary orientation of the triangle. Write $C_{ij} = CF(L_i, L_j)$. Use the end gradings induced by its chosen almost-complex gauge triangle. Holomorphic triangles carry the limiting triangle orientation of Definition 7.2: the mixed triangle at small positive collapse scale is oriented by Definition 5.5, this orientation is pulled back along collapse, and the lines of the corner pieces are then cancelled. The *corner pieces* of the zero-scale limit domain are the one-input mixed pieces $\mathcal{V}_{01}, \mathcal{V}_{12}$ attached at its incoming ends along ordinary strip necks (Section 7.3). Put all generator orientation lines before the determinant line of the triangle. The resulting product

$$\mu: C_{01} \otimes^r C_{12} \longrightarrow C_{02}$$

satisfies

$$(7.1) \quad d\mu(a, b) = (-1)^{|b|} \mu(da, b) + \mu(a, db).$$

Indeed the first strip break requires its translation line to pass the second input generator; the second break does not. This is the boundary calculation for the one-dimensional triangle space.

If $\theta \in C_{12}$ is a homogeneous cycle, then $R_\theta(a) = \mu(a, \theta)$ is a closed map of degree $|\theta|$. The product itself is even in this convention. For a change of reference almost-complex structure at any of the three ends of the triangle, the corresponding symplectic conversion is diagonal on the same generator orientation lines. Applying the conversions with (3.4) gives the product for the new reference structures. If the three conversions $\Theta_{01}, \Theta_{12}, \Theta_{02}$ have parities $|\Theta_{01}|, |\Theta_{12}|, |\Theta_{02}|$, the converted product has degree $m = |\Theta_{01}| + |\Theta_{12}| + |\Theta_{02}|$ and satisfies

$$d\mu(a, b) = (-1)^m ((-1)^{|b|} \mu(da, b) + \mu(a, db)).$$

Inserting a cycle θ then gives a closed map of degree $m + |\theta|$. All subsequent changes of reference structure use this converted formula, with the right tensor convention.

7.2. Collapse preserves the transported orientations. This subsection proves Theorem B. The decisive step is Proposition 7.5, which rests on Lemma 7.3: if a continuous family of Fredholm operators over a square is invertible on three edges, then transport along the fourth edge carries the canonical orientation to the canonical orientation. Proposition 7.4 supplies such a square at a fixed small collapse scale ϵ , for a regular index-zero limiting solution V . Its edge $s = 0$ is constant, equal to the linearized operator at the preglue $W_\epsilon(V)$, which is invertible, with inverse bounded by a fixed power of $1 + |\log \epsilon|$, by the attachment inverse of (I8). Its edge $t = 0$ is the derivative of the contraction equation of (I8) along the chart segment from $W_\epsilon(V)$ to the solution $\mathcal{G}_\epsilon(V)$; it is invertible because the inverse bound at the preglue dominates the quadratic bound on the contraction ball. Its edge $s = 1$ is the operator at $\mathcal{G}_\epsilon(V)$ with its slice centre moved from $W_\epsilon(V)$ to $\mathcal{G}_\epsilon(V)$; it is block triangular, with the invertible linearization on the quotient by gauge directions and an invertible scalar block. Its edge $t = 1$ is the path of linearized mixed operators along the segment, which carries the transported orientation. Lemma 7.3 therefore transports the canonical orientation at the preglue to the canonical orientation at the solution, and Lemma 7.1 identifies the former with the canonical orientation at V .

The bijection of solutions is Paper I's; what is new here is the comparison of determinant lines. The estimates below serve two purposes. Lemma 7.1 extends the linear analysis of Paper I from regular solutions to compact families of limiting configurations, whose operators may have kernel and cokernel; this is needed because orientations live on configuration spaces. Proposition 7.4 constructs the square; it is needed because the preglue is not a solution, so the block-triangular comparison of Lemma 5.2 does not apply along the segment.

The comparison must be defined on compact families of limiting configurations, where kernels and cokernels may occur. We first give that construction. Stationary mixed-end coordinates are included as finite-dimensional domain factors throughout.

Lemma 7.1 (Collapse of stabilized families). *Let $\{V_b\}_{b \in B}$ be a compact smooth family of limiting configurations with the specified autonomous end models and separated stationary coordinates. Assume that their remainders have a common exponential rate β at every required derivative order, and choose the comparison weights in common gaps strictly below that rate. Assume that the collapse datum is unperturbed, as in Paper I: the primary perturbation vanishes on the inserted compression body, and the dependence region and the output of every other perturbation avoid a fixed neighbourhood of the collapse. Primary data on the remaining pieces may vary in a compact regular family, and the length of the collapse arc may vary in*

a compact range. For sufficiently small $\epsilon > 0$, pregluing gives a continuous isomorphism of orientation lines

$$(7.2) \quad \det D_{V_b} \longrightarrow \det L_{W_\epsilon(V_b)}.$$

It is independent, up to positive homotopy, of the stabilization and cutoff choices, and it commutes with the lifting actions. It also commutes with ordinary gauge and strip gluing at an original end, in the following sense. There are constants C and T_0 , and a threshold ϵ_0 independent of the neck length $T \geq T_0$, such that the square formed by (7.2) and the gluing isomorphisms of Lemma 3.3 commutes in the cofinal neck-and-scale regime

$$T \geq C \log(1 + |\log \epsilon|) + T_0, \quad 0 < \epsilon < \epsilon_0,$$

and hence, by continuity, for all $T \geq T_0$ and $0 < \epsilon < \epsilon_0$. At a regular index-zero solution both operators are invertible and (7.2) carries the canonical element to the canonical element.

Proof. Stabilization. Choose a finite-dimensional real space F and smooth compactly supported target columns $j_b : F \rightarrow Y_b$ so that $\widehat{D}_b = (D_{V_b}, j_b) : X_b \oplus F \rightarrow Y_b$ is onto. Such columns exist by smoothness of the cokernel on the limit domain (Proposition 2.1(a),(b), which includes the zero-scale limit domains) and density of compactly supported tests; a finite cover of B and parameter cutoffs give one F . Write $K_b = \ker \widehat{D}_b$. These kernels form a vector bundle. Choose projections $\Gamma_b : X_b \oplus F \rightarrow K_b$ whose restriction to K_b is the identity. They can be formed from compactly supported pairings, together with the stationary coordinate projections at original ends and the coordinate tests on F : a nonzero kernel element cannot vanish in all three sets of tests. The F tests also detect redundant or inactive stabilization columns. A finite selection works over each compact parameter patch. A family of arc lengths in a compact range is identified with a family of metrics on one domain by diffeomorphisms supported along the arc; the constants of Paper I, Sections 7.1–7.7, are uniform on such families. Any free stationary coordinate at an artificial exterior end is also detected by compact pairings on K_b , after enlarging the compact set. Choose those pairings instead of an artificial-end limit projection: the artificial end will disappear on attachment. Then $(\widehat{D}_b, \Gamma_b)$ is an isomorphism, with uniformly bounded inverse over B .

Transport the columns to the preglued domain, away from the transition annuli, and denote them by $j_{\epsilon,b}$. Patch kernel fields by the harmonic lift on the collapsed compression and by the usual cutoffs elsewhere. Call this map $J_{\epsilon,b}$. Transport the compact pairings in Γ_b in the same manner, retaining their finite-dimensional factors in the original order. This defines

$$A_{\epsilon,b} = (L_{W_\epsilon(V_b)}, j_{\epsilon,b}, \Gamma_{\epsilon,b}) : X_{\epsilon,b} \oplus F \longrightarrow Y_{\epsilon,b} \oplus K_b.$$

Four estimates. Put $A = 1 + |\log \epsilon|$, and fix a weight $\delta_* < \min(\beta, \gamma/2)$, where γ is the transition gap of Paper I, Proposition 7.9. The arguments of Paper I, Sections 7.1–7.7, which are written for a regular index-zero solution, give the following four estimates for the stabilized family, with constants C_B , N and $\sigma > 0$ independent of ϵ and b . They are proved in Appendix A.5.

- (i) (Decay.) The stationary columns of the auxiliary exterior operator of Paper I, Section 7.7, built from the truncation of V_b at radius $\epsilon^{1/4}$, differ from those of D_{V_b} by $O(\epsilon^{1/4}) + O(\epsilon^{(\beta-\delta_*)/4})$, uniformly on B . The common rate β replaces the decay of solutions.
- (ii) (Exterior inverse.) The auxiliary exterior operator, augmented by j_b and Γ_b , is invertible, with inverse bounded uniformly in ϵ and b at the end weight δ_* , and by $C_B A^N$ at the end weight $\delta_L \asymp A^{-1}$ on the new ends. This replaces Paper I, Proposition 7.13.

(iii) (Attachment inverse.) $\|A_{\epsilon,b}^{-1}\| \leq C_B A^N$. This extends the attachment inverse of (I8) to the stabilized family.

(iv) (Kernel patching.) For $k \in K_b$, $\|A_{\epsilon,b} J_{\epsilon,b} k - (0, k)\| \leq C_B \epsilon^\sigma A^N \|k\|$. This carries the kernel estimate of Paper I, Proposition 7.8, from the model strip to the glued domain.

By (iii) and (iv), $\|A_{\epsilon,b}^{-1}(0, k) - J_{\epsilon,b} k\| \leq C_B^2 \epsilon^\sigma A^{2N} \|k\|$, which tends to zero. Hence $k \mapsto A_{\epsilon,b}^{-1}(0, k)$ identifies K_b with $\ker(L_{W_\epsilon(V_b)}, j_{\epsilon,b})$ and is asymptotic to the patched kernel map. Cancel $\det F$ in the two identities $\det D = \det \ker(D, j) \otimes (\det F)^*$ to obtain (7.2). Its independence of the stabilization and cutoff choices, and its compatibility with the lifting actions and with gauge and strip gluing, are proved in Appendix A.5.

At a regular index-zero solution the stabilization may be taken to be zero. Then $K_b = 0$, the realized operator is invertible, and the estimate above makes the preglued operator invertible, so (7.2) carries the canonical element to the canonical element; by stabilization independence the same holds for every stabilization. This proves all the assertions. $\square$

Definition 7.2 (Limiting orientation). Let V be a configuration on the zero-scale limit domain of the triangle, with gauge inputs a_{01}, a_{12} and output x_{02} . For $0 < \epsilon < \epsilon_0$, with ϵ_0 as in Lemma 7.1 for a compact family containing V , the *limiting orientation* of V is the map

$$L_G(a_{01}) \otimes L_G(a_{12}) \otimes \det D_V \longrightarrow L_S(x_{02})$$

obtained from the orientation of Definition 5.5 at the preglue $W_\epsilon(V)$ through (7.2). Let $\mathcal{V}_{01}, \mathcal{V}_{12}$ be the corner pieces of the zero-scale domain (Section 7.1). For a configuration u on the triangle domain with inputs x_{01}, x_{12} and output x_{02} , let $V = C_{x_{01}} \# C_{x_{12}} \# u$ be obtained by gluing the constant pairs on the corner pieces to the inputs of u . The *limiting triangle orientation* of u is the map

$$L_S(x_{01}) \otimes L_S(x_{12}) \otimes \det D_u \longrightarrow L_S(x_{02})$$

obtained from the limiting orientation of V by cancelling, in cyclic order, the canonically oriented lines of the two invertible constant pairs, with $L_G(a_{x_{ij}}) = L_S(x_{ij})$. The folded triangle is treated in the same way, with the pieces of the pairs (A_0, L_C) and (L_C, A_1) .

The orientation of Definition 5.5 at scale ϵ is formed after moving the metric, with the configuration fixed, along stages (A) and (B) of (I11), to a metric with a product collar at the cut. The configuration spaces do not depend on the metric (Appendix B.1), so this is transport along a path of data, continuous in ϵ . By part (V1) of Lemma 7.6 below, the strip actions of the corner pieces are those of C_{01} and C_{12} . A limiting triangle need not admit a shift to a constant triangle, because the triple intersection of its boundary labels can be empty; the orientation is therefore defined by collapse rather than by shifting.

The limiting orientations are independent of ϵ and of the auxiliary choices, and they are coherent at breaks, by the following arguments.

Continuity for $0 < \epsilon < \epsilon_0$ on a compact configuration family compares sufficiently small scales. On overlaps choose a common smaller scale. Stabilization independence and ordinary gluing compare the auxiliary columns, corner pieces and necks. For a symplectic strip at an incoming end, apply the one-input comparison $C_x \# u \sim A \# C_y$ in its corner piece, away from the central collapse neighbourhood, and absorb A into the ordinary cobordism factor. Thus the full end-gluing diagram commutes before the lines of the corner pieces are cancelled.

The same comparison applies with nonconstant configurations on the corner pieces. If c_{01}, c_{12} are such configurations, for instance solutions or the preglued configurations $C_x \# w$ used below, shift each to $A_{ij} \# C_{ij}'$ inside its domain defining the mixed map (Proposition 2.2(a); Appendix B covers the domains defining the mixed maps). This is a compact

family on a fixed limiting domain, of the kind allowed in Lemma 7.1. Its compatibility with gauge gluing at the original inputs, together with the coherence of the positive-scale orientation under input gauge gluing, absorbs A_{ij} . Hence at $c_{01} \# c_{12} \# u$ the pulled-back orientation is $\Gamma_{c_{01}} \otimes \Gamma_{c_{12}} \otimes o(u)$, in cyclic order. At an incoming break of the one-dimensional triangle space the internal neck grows with the breaking parameter. Regroup $C_x \#_T(w \#_{T'} u')$ as $(C_x \# w) \#_{T'} u'$ on a domain with a longer internal neck. This neck runs along the collapsing arc: the output strip of a corner piece is bounded by P_0 and P_1 , and in the limit domain the only P_1 boundary is the collapse arc (Paper I, proof of Theorem 9.10). Lemma 7.1 allows compact families of lengths of this arc, so the pulled-back orientation does not depend on the neck length over compact ranges. This proves coherence of the defined triangle orientation without an independent component sign.

For a nonsolution reference family supplied only in a strong weighted topology, use a comparison weight below its remainder rate. Solutions are used with their exponential decay, without truncation.

Now let V be a regular index-zero limiting solution. Its realized operator and the operator at its preglue are invertible. Let $\mathcal{G}_\epsilon(V)$ be the nearby solution obtained from the preglue by the contraction mapping argument of (I8). It remains to compare the operator at the preglue with the operator at $\mathcal{G}_\epsilon(V)$. The chart used for the contraction estimate need not be the chart used to define the determinant bundle, and the second derivative of a fixed mixed chart need not be bounded uniformly in the collapse norm. The comparison is therefore made around a square of operators.

Lemma 7.3 (Transport around a square). *Let $D_{s,t}$ be a continuous family of Fredholm operators on a square. Suppose the operators on three edges are invertible. Parallel transport of orientations on the fourth edge carries the canonical orientation at its first endpoint to the canonical orientation at its second endpoint.*

Proof. The orientation double cover of the determinant bundle is trivial over the square. Transport around its boundary is therefore the identity. On each of the three invertible edges the canonical orientation is a continuous section, so transport is positive. The remaining edge must have the same property at its endpoints. $\square$

Proposition 7.4 (Affine centres at positive collapse scale). *Fix the collapse datum of Lemma 7.1 and a sufficiently small positive scale ϵ . Let $c_s = \Phi(sz_*) = (A_s, u_s)$, $0 \leq s \leq 1$, be the segment in the exact matching chart from the preglue $c_0 = W_\epsilon(V)$ to the solution $c_1 = \mathcal{G}_\epsilon(V)$ obtained by the contraction mapping argument of (I8). At a common sufficiently small weight on the original ends there is a continuous Fredholm family $D_{s,t}$, $0 \leq s, t \leq 1$, whose $t = 0$ edge is the derivative of the contraction equation and whose $t = 1$ edge is the linearized mixed operator at c_s with slice centre A_s . It retains the finite stationary coordinates, both operator rows, and their weak-adjoint domains. If V is regular of index zero, the other three edges are invertible.*

The Fredholm construction also holds in compact smooth families of these small matching-chart neighbourhoods, including nonsolutions, with bounds depending on the positive scale.

Proof. The proof has three steps. Step 1 constructs the family $D_{s,t}$. Step 2 proves that it is a continuous family of Fredholm operators. Step 3 proves that the edges $t = 0$, $s = 0$ and $s = 1$ are invertible when V is regular of index zero.

Step 1: construction of $D_{s,t}$. Write $g = g_\epsilon$, $\lambda = \lambda_\epsilon$ and $\widehat{g} = \lambda^{-2}g$ for the metrics and rescaling in Paper I, Sections 7–8. Put

$$(7.3) \quad C_{s,t} = (1-t)A_0 + tA_s.$$

The centre need not solve the primary equation at a mixed end. The exact chart fixes the normal connection trace, so the covariant Neumann condition is common to all these centres. A single normal gauge for A_0 , applied to the whole family, makes this trace zero without changing the affine relation. For each fixed ϵ , the changed compression region is compact and its minimum scale is positive. Paper I's strong norms are then equivalent to ordinary weighted Sobolev norms, with the stationary coordinates adjoined; these equivalence constants need not be uniform as $\epsilon \rightarrow 0$.

At an original mixed end e , choose the primary representatives $B_e(p)$ in the spatial Neumann Coulomb slice of $B_e(0)$:

$$(7.4) \quad d_{B_e(0)}^*(B_e(p) - B_e(0)) = 0, \quad (B_e(p) - B_e(0))(n) = 0.$$

The spatial scalar inverse and implicit function theorem give this local section of the family of primary solutions. Extend it into the matching chart after subtracting its stationary value. This ensures that the affine scalar equation has weighted target. If an initial local primary section is used instead, its small spatial gauge gives this normalization; extend that gauge with a fixed longitudinal cutoff. The correction is quadratic in the stationary parameters when the initial tangent is Coulomb. It occurs on the fixed original-end geometry and preserves the initial collapse operator and the quadratic estimates of Paper I. The chart segment is smooth in all stationary-subtracted norms.

We first compare scalar operators without assuming compact support of $C_{s,t} - A_s$. With $d\mu = \lambda^2 dV_{\hat{g}} = \lambda^{-2} dV_g$, the operator $P_{C,A} = \lambda^2 d_C^{*,g} d_A$ has weak form

$$(7.5) \quad \langle \xi, P_{C,A} \eta \rangle_\mu = \int \langle d_C \xi, d_A \eta \rangle_{\hat{g}} d\mu.$$

The two factors on the right are generally different. Let C_P be the Poincaré constant for A , and let a bound the bracket with $C - A$ in the rescaled metric. Then

$$\operatorname{Re} \langle \xi, P_{C,A} \xi \rangle_\mu \geq (1 - aC_P) \|d_A \xi\|_\mu^2.$$

If $aC_P < 1$, the bounded coercive form gives a weak inverse. Conjugating by an original-end weight $w = e^\omega$ adds form terms controlled by $\|d\omega\|_\infty$ and its square. For a common sufficiently small weight they are absorbed. Neumann elliptic regularity gives the higher Sobolev inverses at this fixed scale. The identical spatial argument makes $d_C^* d_B$ invertible when B is the perturbed-flat limit of c_s at a mixed end and C is sufficiently close to it, whether or not C is perturbed flat.

We next specify all stationary columns. Let $\varphi_{s,0,e}(p)$ be the connection value of the contraction-chart column at the limit $B_e(s)$ of c_s at the end e , and let $C = C_{s,t,e}$ be the corresponding limit of the centre. Define

$$(7.6) \quad \psi_{s,t,e}(p) = (d_C^* d_{B_e(s)})^{-1} d_C^* \varphi_{s,0,e}(p), \quad \varphi_{s,t,e} = \varphi_{s,0,e} - d_{B_e(s)} \psi_{s,t,e}.$$

The new value is C -Coulomb and has the same primary boundary class. At $t = 0$ the correction vanishes by (7.4). Extend ψ by a fixed end cutoff to $\Xi_{s,t}(p)$ and put

$$H_{s,t}(p) = H_{s,0}(p) - G_s \Xi_{s,t}(p), \quad G_s \xi = (d_{A_s} \xi, 0), \quad J_{s,t}(w, p) = J_s^{\text{dec}} w + H_{s,t}(p).$$

Here J_s^{dec} is the decaying derivative of the exact chart. Reading the primary classes at infinity and subtracting $H_{s,t}$ proves that $J_{s,t}$ is a source isomorphism. It is the derivative of a matching chart: solve the nonlinear spatial C -Coulomb equation based at $B_e(s)$ and extend its small gauge with the same cutoff. This leaves the configuration c_s fixed and differentiates

to the displayed formula. The scalar stationary image vanishes, and all remaining scalar columns belong to the weighted target. Equivariance gives

$$(7.7) \quad \nabla \mathcal{F}_{c_s} G_s \Xi = \rho_*(\Xi) \mathcal{F}(c_s).$$

This weighted equation column is retained; it need not vanish at a nonsolution.

There is also a target derivative. On the section factor choose S_c by pointwise parallel transport from $u(z)$ to $u_0(z)$ along the short geodesic in the target. The gauge target is fixed. Shrinking the chart gives uniform normal neighbourhoods; $S_{c_0} = 1$ and $\nabla S|_{c_0} = 0$. Pointwise multiplication estimates preserve Paper I's quadratic bound and initial operator. Differentiating gives

$$(7.8) \quad d(S\mathcal{F})_c[v] = S_c \nabla_v \mathcal{F} + R_c v, \quad R_c v = (\nabla_v S)_c \mathcal{F}(c).$$

The coefficient of R_c tends to zero on the original ends, where the residual tends to zero. It is compact from the weighted decaying H^1 domain to L^2 , by compact truncation and Rellich, and its stationary columns are weighted finite rank. Since $\nabla S|_{c_0} = 0$, the formula (7.8) gives $R_{c_0} = 0$ even though the preglue is not a solution; also $R_{c_1} = 0$ since c_1 is a solution. The target map itself may have a nonidentity stationary limit.

The required square is

$$(7.9) \quad D_{s,t} = \left(S_{c_s} \nabla \mathcal{F}_{c_s} + (1-t) R_{c_s}, \lambda^2 d_{C_{s,t}}^{*,g} \pi_A \right) J_{s,t}.$$

At $t = 1$ the residual term is absent and the centre is A_s , so this edge is the path of linearized mixed operators at c_s .

Step 2: Fredholm property and continuity. The family (7.9) is a continuous family of Fredholm operators on the whole square. On the changed compact core, which is smooth at this fixed scale, this follows from the local estimates of [DFL21, DFL24]. On each mixed end the operator differs from the translation-invariant end operator at the stationary limit by a zeroth-order term bounded by a constant times the C^0 distance between the centre and that limit; when this distance is small, a Neumann series gives the end inverse and the inverse of its adjoint. The estimates are given in Appendix A.6. The smallness conditions used in Steps 1 and 2 follow from the contraction bounds of (I8). On the unchanged original ends the spatial gaps and cylinder constants are fixed. The stationary correction is bounded by $C(1 + |\log \epsilon|)^{m+1} e_\epsilon \rightarrow 0$, with e_ϵ the pregluing error of (I8). The strong-control estimate of Paper I, (8.1), bounds the variation sz_* on unit patches of the rescaled metric by its strong norm, so Sobolev embedding gives $\|A_s - A_0\|_{C^0(\widehat{g})} \rightarrow 0$. The temporal component in this norm is the rescaled component, with no additional ϵ^{-1} factor. The preglue's uniform scalar Poincaré estimate therefore persists, and both scalar and end smallness conditions hold for sufficiently small ϵ .

Step 3: the three invertible edges. The $t = 0$ edge of (7.9) is exactly the derivative of the contraction equation, including the affine scalar row and (7.8). By (I8) the operator at the preglue is invertible in the strong norms, with inverse bounded by a fixed power of $1 + |\log \epsilon|$. The derivative of the nonlinear part of the contraction equation along the segment is bounded by a constant times $\|z_*\|$, which is at most a fixed power of $1 + |\log \epsilon|$ times the pregluing error, a positive power of ϵ . Hence this entire edge is invertible. Strong invertibility implies first-order invertibility: kernels regularize into the strong domain, and a weak cokernel would annihilate every smooth target reached by the strong inverse. The first-order operator has closed range, being Fredholm by Step 2, and the smooth targets reached by the strong inverse are dense in its target, so such a cokernel element vanishes. At $s = 0$, the centres and source maps are constant and $R = 0$, giving the invertible preglued operator used in (7.2).

At $s = 1$, equivariance gives $\nabla \mathcal{F}_{c_1} G_1 = 0$. The scalar inverse splits the source into horizontal vectors and decaying gauge directions. The resulting blocks are the intrinsic quotient linearization and $P_{C_{1,t}, A_1}$. Changing (7.6) changes only the representative of that quotient tangent; (7.7) now vanishes. Regularity and index zero make the quotient block invertible, and scalar coercivity makes the second block invertible. This proves invertibility of the third edge.

Compact parameter families use finite local covers; the transition gauge transformations transport both rows and all stationary columns. Ordinary gluing at fixed positive scale commutes with these identifications by the bordered determinant construction, with corrections placed on the mixed factors and fixed ordinary end models. $\square$

Proposition 7.5. *The map (7.2), followed by transport from $W_\epsilon(V)$ to $\mathcal{G}_\epsilon(V)$, preserves the sign of the zero-dimensional solution.*

Proof. Use the family (7.9) of Proposition 7.4. Its contraction edge is invertible by the nonlinear estimate, its preglue edge is the initial invertible collapse operator, and its solution edge is invertible by the quotient and scalar block decomposition. Lemma 7.3 identifies the canonical orientations along the remaining edge, the path of linearized mixed operators at c_s . Composing this identification with (7.2) proves the claim.

That path, the edge $t = 1$, is an extended realization, not an permitted one in the sense of Definition 5.1: its target map S_{c_s} may have a nonidentity stationary limit, and its source is the chart at c_0 . It differs from a permitted realization L_{c_s} along the same configurations by bounded source and target isomorphisms depending continuously on s . At $s = 0$ these are the identity, since $S_{c_0} = 1$ and the chart is centred at c_0 . The naturality maps of determinant lines under these isomorphisms commute with transport, and at $s = 1$ both operators are invertible, so the naturality map carries the canonical element to the canonical element. Hence the identification along the edge $t = 1$ is the one in the determinant bundle of the permitted realizations along c_s , which carries the transported orientation. $\square$

If the conformal metric in the collapse estimate is changed within the unperturbed collar class, the zero set of the mixed equation is unchanged. At a solution Lemma 5.2 compares the gauge rows by invertible paths, so the sign is unchanged. This permits use of the metric of Paper I without introducing a new orientation convention.

The construction is compatible with compact parameter families because its stabilizations and operator bounds may be chosen locally uniformly. Associativity of ordinary gluing identifies the determinant maps at broken ends. This proves Theorem B. Together with Lemmas 7.6 and 7.7 below, it gives the two-handle square, Theorem 7.8.

7.3. Mixed maps of the degeneration pieces and the two-handle map. The degenerations of the mixed triangle contain one-input mixed pieces which are not the special quintuples on which the mixed maps are defined. At the collapsed end these are the two DFL pieces $\mathcal{V}_{01}, \mathcal{V}_{12}$ of the zero-scale domain, cut off along ordinary strip ends; they are the corner pieces of Section 7.1. Each carries an original mixed end and a new mixed end of Paper I, Proposition 6.9. At the end with the long cut collar the piece is the DFL piece $\mathcal{V}_{02}$ of the pair (P_0, P_2) beyond $Y_{02}^\#$. The folded triangle of Lemma 9.2 has the corresponding pieces for the pairs (A_0, L_C) , (L_C, A_1) and (A_0, A_1) . Each such piece $\mathcal{V}$ is a domain of DFL type in the class of (II) (Paper I, the paragraph following Lemma 9.7, and the proof of Theorem 11.8). By (II) its count is chain homotopic over $\mathbb{F}_2$ to the mixed map of the pair; Lemma 7.6 gives the integral statement.

Orient the configurations on $\mathcal{V}$ by the map $\Gamma_{\mathcal{V}}$ of Proposition 5.4; Proposition 2.2(a) applies to $\mathcal{V}$ by Appendix B.1. Let $\mathcal{V}^{\text{st}}$ be the special quintuple of the same pair, carrying the data of its mixed map N . Suppose that $\mathcal{V}$ has the same ordinary end data, including the almost-complex structure on the gauge end collar; otherwise compose with the continuations of Proposition 6.4 and the conversions of Section 4. Let $N_{\mathcal{V}}$ be the count of regular index-zero solutions on $\mathcal{V}$, signed by $\Gamma_{\mathcal{V}}$. For a regular strip u from x to y at the output, the action $R_{\mathcal{V}}(u) : L_S(x) \otimes \det D_u \rightarrow L_S(y)$ is defined as in Section 8.2: glue u to the constant pair C_x on $\mathcal{V}$ and evaluate $\Gamma_{\mathcal{V}}$ on the positive element of the invertible constant-pair operator.

Lemma 7.6 (Integral identification of the degeneration pieces). *Let $\mathcal{V}$ be a domain of DFL type in the class of (II), for a transverse pair whose mixed map N is defined on the special quintuple $\mathcal{V}^{\text{st}}$ of the pair, and joined to $\mathcal{V}^{\text{st}}$ by the path of domains and data in the proof of (II). Suppose that $\mathcal{V}$ and $\mathcal{V}^{\text{st}}$ have the same ordinary end data, including the almost-complex structure on the gauge end collar. Let $N_{\mathcal{V}}$ be the count of regular index-zero solutions on $\mathcal{V}$ signed by the orientation map $\Gamma_{\mathcal{V}}$ of Proposition 5.4, and let $R_{\mathcal{V}}(u)$ be the strip action defined above. The degeneration pieces of this section and the closed-cut piece M_{λ_0} of Proposition 8.14 are domains of this kind.*

- (V1) $R_{\mathcal{V}}(u) = R_{\mathcal{V}^{\text{st}}}(u)$ for every regular strip u at the output. Thus the strip orientations transported through $\mathcal{V}$ are those of the target complex of N .
- (V2) $N_{\mathcal{V}}$ is closed and even, and $N_{\mathcal{V}} - N = \delta H_{\mathcal{V}}$ over $\mathbb{Z}$. Here $H_{\mathcal{V}}$ is the signed count of the isolated parametrized solutions of vertical index -1 along the path of domains and data from $\mathcal{V}^{\text{st}}$ to $\mathcal{V}$ used in the proof of (II).

Proof. The proof of Paper I, Lemma 9.7, joins $\mathcal{V}^{\text{st}}$ to $\mathcal{V}$ by a path $\mathcal{V}_b$, $b \in [0, 1]$, of domains and data with fixed end types and fixed ordinary end data, which stays in the class of that lemma. It consists of compactly supported changes of the conformal structure of the section region, of the seam placement, of the metrics near the mixed ends and of the neck lengths, followed by a change of the translation-invariant data on the mixed ends. The configuration spaces along the path are identified by the decorated diffeomorphisms of Appendix B.1, and the data vary as permitted in Proposition 2.1. Hence $\Gamma_{\mathcal{V}_b}$ varies continuously with b , by Propositions 2.2 and 5.4, as for the parametric orientations of Section 6. With no secondary or inhomogeneous data the constant pair C_x is a solution on every $\mathcal{V}_b$, with invertible linearized operator (Lemma A.2); by openness, uniformly on the compact path, the operator remains invertible for the small data used here. Glue u at the output with one long neck along the whole path (Lemma 3.3). Then $R_{\mathcal{V}_b}(u)$ is a continuous family of isomorphisms between the fixed lines $L_S(x) \otimes \det D_u$ and $L_S(y)$, hence constant. This proves (V1); it is the argument of Corollary 8.7.

For (V2), choose small invariant secondary perturbations along the path, supported in the cores and regular as a parametrized family through index one (Paper I, Theorems 5.4 and 5.6(iii)). At $b = 0$ they are the regular data of N . At $b = 1$ they are the data of $\mathcal{V}$, after an arbitrarily small compactly supported change which makes index one regular; by compactness and local uniqueness this change leaves the index-zero count $N_{\mathcal{V}}$ unchanged, as in the proof of Proposition 6.4. On the compactly supported part of the path the parameter column is compactly supported; on the change of mixed-end data it is the weighted column of Paper I, Lemma 10.15, whose hypothesis of uniform energy decay is verified in the proof of Paper I, Lemma 9.7. This is the mixed-end family of Section 8.5, now on $\mathcal{V}$, with $\Omega = \Gamma_{\mathcal{V}_b}$ and $f = 0$. The proof of Lemma 8.12 uses only three properties of Ω : its compatibility with gauge paths at the input, its compatibility with strips at the output, and the constancy of the output differential along the family. For $\Gamma_{\mathcal{V}_b}$ the first two are Proposition 5.4 and the third is

(V1). That proof therefore gives $N_V - N_{V^{\text{st}}} = \delta H_V$, and its last assertion gives closedness at both ends. By construction $N_{V^{\text{st}}} = N$, and all these counts are even because their solutions have index zero. This is the construction which produces H_M in Proposition 8.14. $\square$

The piece counts are replaced by the mixed maps through the following identity. If F, F', G, G' are closed even maps and $F' - F = \delta H$, $G' - G = \delta K$ with H, K odd, then

$$F' \otimes^r G' - F \otimes^r G = \delta(H \otimes^r G' + F \otimes^r K),$$

since (3.3) and (3.4) give $\delta(H \otimes^r G') = \delta H \otimes^r G'$ and $\delta(F \otimes^r K) = F \otimes^r \delta K$. Composition with closed maps preserves such identities by (3.2).

Consider the cycles θ_G and $\theta_S = N_{12}\theta_G$ of Section 7.1. The unit diagonal of N_{12} and the filtration of C_{12} by the Morse index of ζ_Q identify θ_S with the top generator o_{q_*} , up to the chosen sign. The mixed triangle family has one endpoint equal to the gauge triangle followed by the count $N_{V_{02}}$ on the DFL piece beyond the cut, and the other endpoint is the symplectic triangle with the two counts $N_{V_{01}}, N_{V_{12}}$ of the zero-scale domain as inputs. Proposition 7.5 compares the counts at the collapse endpoint with the transported orientations, and the two endpoints of the family are related by its oriented boundary (5.3). By part (V1) of Lemma 7.6, the input strip orientations entering the limiting triangle orientation are those of C_{01} and C_{12} , and the output strip action at the long-collar end is that of C_{02} . Along the family the output strip action is defined through the terminal constant pairs on the truncated output regions, which are invertible (Lemma A.2); it is therefore constant, by the argument of Corollary 8.7, and the output differential is that of C_{02} throughout. The broken boundary terms give a homotopy $N_{V_{02}}\mu_G \simeq \pm\mu_S(N_{V_{01}} \otimes^r N_{V_{12}})$. Since μ_G and μ_S are closed and even, part (V2) of Lemma 7.6, the identity for tensors of homotopies above and (3.2) replace the three piece counts by N_{01}, N_{12} and N_{02} , hence

$$(7.10) \quad N_{02}\mu_G \simeq \pm\mu_S(N_{01} \otimes^r N_{12}).$$

The gauge two-handle map is itself a triangle map over the integers.

Lemma 7.7. *With the almost-complex orientations, the ordering of Lemma 3.2 and the conversions inserted as in Proposition 4.4,*

$$I_J^\#(W_2) \simeq \pm\mu_G(\cdot \otimes \theta_G),$$

where μ_G is the framed gauge triangle map of the torus-summed triangle (Paper I, Sections 9.3 and 9.4) and $\theta_G = \pm o_{q_*}$.

Proof. Paper I, Proposition 9.5, proves the identity modulo two, and its argument is integral once three points are checked. First, $W_2^\# = X^\# \circ W'$ is cut along $Y_{01}^\# \sqcup Y_{12}^\#$, and the connecting image in [KM11b, Proposition 5.4] vanishes by the H^0 computation in the proof of Paper I, Proposition 9.5. That proposition and the Künneth isomorphism of [KM11b, Section 5.2] hold over $\mathbb{Z}$; the latter is written in the left convention, and its transfer to the ordering of Lemma 3.2 is trivial on $a \otimes \tilde{\theta}_G$, because the invariant $\tilde{\theta}_G$ of the filling is even in its induced grading. Second, an element of the lifting group with $s_{12} = 1$ is realized by an automorphism supported on the torus region. The almost-complex orientations are invariant under all bundle automorphisms (Section 4), so it maps solutions to solutions preserving their signs, and the orbit count of Paper I gives $\mu^{\{s_{12}=0\}}(a \otimes b) = \mu^{\phi_X}(a \otimes \pi b)$ over $\mathbb{Z}$, with no factor two. The lifting-class sign of the conversion is trivial on these classes by Proposition 4.2, because J_W equals the fixed germ on the torus region. Third, the conversions to the reference structures of the objects are inserted and cancelled as in the proof of Proposition 4.4. Finally, the filling has exactly one index-zero solution, with limit q_* and invertible operator (Paper I,

Proposition 9.4), so $\theta_G = \pm o_{q_*}$. Paper I, Proposition 9.5, and [KM11b, Proposition 5.4] are stated on homology. The composition holds at chain level, up to an integral chain homotopy, by the stretching formula [KM11b, Section 3.9, (23)], with the conversions treated as in Proposition 4.4; so the identity holds up to chain homotopy. $\square$

Theorem 7.8 (Two-handle square). *Let $W_2: M_0 \rightarrow M_1$ be a two-handle cobordism with the data of Section 7.1, and let p_0 and p_1 be the presentations of M_0 and M_1 given by the pairs (P_0, P_1) and (P_0, P_2) , transported to the global frames as in Paper I, Section 9.8. Then*

$$(7.11) \quad N_{p_1} I_J^\#(W_2) \simeq \pm S_{W_2} N_{p_0} \quad \text{in } K(\mathbb{Z}).$$

Proof. The frame transports of Paper I, Theorem 9.13, are pushforwards under diffeomorphisms and bundle isomorphisms preserving the torus decoration; their mixed squares are those of Section 14. Evaluate the second input of (7.10) on the cycle θ_G , and apply Lemma 7.7. Since $N_{12}\theta_G = \theta_S$, this gives (7.11). Changing the reference structures conjugates this identity by the right-placed conversions and uses (3.2). Twisting the bundle on the two-handle changes the gauge triangle class but not the orientation argument; its relative bundle is retained throughout the mixed family. $\square$

For simultaneous gluing at two triangle ends, use one stabilization off both ends. The determinant comparison is independent of the order of gluing, up to the graded reordering already present in (3.4). Thus the triangle systems are coherent at corners, which is needed when the same product is later folded into a compression cup.

8. TORUS EXCISION OVER THE INTEGERS

This section proves the integral mixed excision theorem, Theorem 8.15, and with it Theorem C. The torus domains, the closed cut $Z = D_0 \cup (-D_1) \cong (-M) \# M^+ \# T^3$ and the reverse excision cobordism are those of Section 2.5. The gluing tori $T_a \subset D_0 = (-U_0) \cup_{T^2} V_0$ and $T_b \subset D_1$ are homologous odd tori in Z , the lifting group of Z is generated by $s_Z = \text{PD}[T_a] \pmod{2}$, and excising Z along T_a and T_b gives $Z' = (-M)^\# \sqcup (M^+)^\#$. Write $E: C_G(Z) \rightarrow C_G(Z')$ for the excision map of Kronheimer and Mrowka and $\bar{E}$ for the reverse excision map. The ordinary gauge theory is in Appendix D. There the almost-complex structures at the closed cut are chosen with (D.1), so that the conversion at Z is defined on the quotient complex. Proposition D.3 then shows that E and $\bar{E}$ are inverse over $\mathbb{Z}$ up to one sign, with the lifting groups retained: $\bar{E}E$ is chain homotopic to the identity up to one sign, and both maps are chain homotopy equivalences. Three points of that proof are not in the literature.

- (i) Regular cut data exist with secondary perturbations supported in fixed collars of the original ends and invariant under the lifting groups, while the new three-torus ends carry the unperturbed equation.
- (ii) For the nonseparating three-torus cut of the composite of excision and reverse excision, the raw determinant-one count is twice the graded partial trace over the full two-generator complex of the three-torus; the larger lifting-group quotient is taken afterwards.
- (iii) The parameter-first orientation of homotopies used here and the fibre-first orientation of [KM11b, Section 3.9] are related by $K = -H_0 \text{gr}$, where K and H_0 are the homotopies in the two conventions and gr is the grading operator.

The mixed content is in Sections 8.1 to 8.7: the orientation maps on the separate, completed and joined torus domains, their compatibility with torus completion, and the oriented comparison families, which together give the chain homotopy of Theorem 8.15. The position of

the conversion line at the closed cut is fixed by the following convention, which is also used in Sections 5 and 9.

Convention 8.1. The determinant line of the conversion at an interior cut is placed after the complete input generator orientation line and before the complete configuration determinant line. Its position is fixed relative to the entire configuration, independently of where the geometric cut is drawn.

Suppose the configuration is written as a sequence of glued pieces. Moving an odd conversion line with the cut through that sequence would move it past the determinant line of the intervening piece. The sign would then depend on its index. Convention 8.1 avoids this dependence: moving the line to a fixed reference position costs a sign determined by the entire configuration, which is unchanged by moving the cut. At index zero the alternative fixed placement after the complete configuration line agrees with this one.

8.1. The torus domains and their orientation lines. The remainder of this section proves the integral mixed excision theorem. The torus domains, namely the separate collared domains $W_{a,R}$ ($a = \pm, 0 \leq R < \infty$), their completions $W_{a,\infty}$ and the joined domains W_λ^J ($\lambda > 0$), with the bridge and the closed cut, are described in Section 2.5. The unsigned content is Paper I, Section 10: the index and energy formulas and the loss formula (Paper I, Theorem 10.13 and Proposition 10.14), the torus-completion bijections (I3), the collar and width families, and the closed-cut decomposition (I2). We use those results as they are proved there. The domain list of Section 2 includes the finite-width torus domains, their completed limits and the closed-cut mixed domain, so Proposition 2.1 and Proposition 2.2 apply to them and to the compact families below. What we add is an orientation for every count and every one-dimensional family, and the resulting chain homotopy over $\mathbb{Z}$.

At $R = 0$ the separate domain is the original special quintuple, and its count is the mixed map N_a . For every $\lambda > 0$ the closed cut of (I2) (Paper I, Lemma 10.18) decomposes the joined domain as

$$W_\lambda^J = \bar{\mathcal{E}}_\lambda \cup_Z M_\lambda.$$

Here, as in Section 2.5, $\bar{\mathcal{E}}_\lambda$ is the reverse excision cobordism of Kronheimer and Mrowka [KM11b, Section 5.3], with inputs $(-M)^\#$ and $(M^+)^\#$, and M_λ is a mixed domain of DFL type for the pair (A_0, A_1) , with compression bodies $D_{0,\lambda}, D_{1,\lambda}$ and single ordinary input Z . The lifting group is the fibre product $\phi_J = \phi_{\bar{\mathcal{E}}} \times_{\phi_Z} \phi_M$ over $\phi_Z = \langle s_Z \rangle$.

Labels and lines. An input label is an ordered pair $a = (a_-, a_+)$ of critical points on $(-M)^\#$ and $(M^+)^\#$, with generator orientation line $I_a = L_G(a_-) \otimes L_G(a_+)$. An output label is a point

$$x = (x_-, x_+) \in A_0 \cap A_1 = (\tilde{P}_0 \cap \tilde{P}_1) \times (R_0 \cap R_1).$$

Let $C_{S,-}$ and $C_{S,+}$ be the target complexes of N_- and N_+ , with their transported orientation systems, and let $L_\pm(x_\pm)$ be their generator orientation lines. Put

$$V_x = L_-(x_-) \otimes L_+(x_+), \quad C_I = C_G((-M)^\#) \otimes^r C_G((M^+)^\#), \quad C_V = C_{S,-} \otimes^r C_{S,+}.$$

In Section 9 these are $CF(\bar{p})$ and $CF(q)$.

Choose the invariant ordinary datum f_Z on Z to vanish on a neighbourhood of the entire critical set. The determinant-one critical classes on Z over x are exactly two, as in [DFL21, Lemmas 2.43 and 2.45]: the two genus matching components permit two central signs, and a constant central gauge on either compression piece changes both, which leaves $\{1, -1\}^2$ modulo its diagonal. A bundle automorphism of lifting class $s_Z = \text{PD}[T_a] \pmod{2}$, the projection of $\cos(\pi\rho) + \sin(\pi\rho)n$ on a genus collar, exchanges them, and it acts freely (Paper I,

Lemma 10.18(iv), for the lifting classes). Both are irreducible [DFL21, Lemma 2.43 and the discussion after (2.44)] and, the intersection being transverse, nondegenerate [DFL21, Lemma 2.47], so the quotient critical set is $A_0 \cap A_1$. Let β_x be either lift. The complex $C_Z = CF_{\mathbb{X}}(A_0, A_1)$ is the target of N_Z ; its generator orientation line over x is

$$Z_x = L_G^M(\beta_x),$$

the gauge line of β_x in the almost-complex structure of the mixed piece. The two lifts have lines identified by the action of s_Z , as in Proposition 5.4.

Finally, let Q_x be the conversion line at β_x from the structure induced by $\bar{\mathcal{E}}_\lambda$ to that of M_λ , and let $|\Theta|$ be its parity, the degree of the conversion Θ at Z . By Section 4, $|\Theta|$ is independent of x . Choose one orientation of Q at a reference connection and transport it along paths of connections on Z . Take the structure of M_λ with the product germ near the fibres T_a, T_b on its torus region, as in Section D.1, so that (D.1) holds for both structures at Z . By Proposition 4.2, the transport therefore gives coherent orientations ϑ_x of all Q_x , invariant under ϕ_Z . The composite $N_Z \bar{\mathcal{E}}$ always means the composite through the right-placed conversion Θ^r at Z , with its line in the position of Convention 8.1; it has degree $|\Theta|$.

Lemma 8.2. *Let $C_{M,x}$ be the constant pair on M_λ over x , with limit β_x on its ordinary input, for the primary perturbations alone, with no secondary or inhomogeneous data. Its linearized mixed operator, with stationary columns and scalar row, is invertible for every finite $\lambda > 0$. It is also invertible along the path of domains and data used to prove the deformation invariance (II) of the mixed map, which joins M_{λ_0} to the special quintuple of (D_0, D_1) by changing the finite torus-neck lengths and the compact metric. The inverse bounds are uniform on compact subfamilies.*

Proof. The path is that of the deformation invariance (II) of the mixed map, as applied in Paper I, Proposition 10.21. The lemma is part of Lemma A.2 in Appendix A: the kernel vanishes by the quadratic identity of [DFL21, Lemma 4.40] adapted to these domains, and the index is zero by [DFL21, Proposition 3.19] at the special quintuple and continuity along the family. The proof of that lemma uses the local theory of Section A.1, Lemma A.1, Paper I, Theorem 10.13, Lemma 8.4(iii) and Lemma 8.5. $\square$

The determinant line of an invertible operator has the canonical positive element 1, and conjugation by a gauge transformation preserves it.

8.2. Orientation maps on the torus domains. The proof of Theorem 8.15 compares two orientations of the same index-zero solutions on the joined domains. The closed-cut orientation G combines the gauge orientation of the reverse excision cobordism, the conversion at Z and the orientation of the closed-cut piece M_λ ; along a long closed neck its counts become the composite $N_M \bar{\mathcal{E}}$. The product orientation Π combines the orientations of the two separate domains across torus completion, and at large width its counts are the tensor product $N_{-, \infty} \otimes^r N_{+, \infty}$. Proposition 8.8 shows that G and Π are related by one isomorphism of lines $K_{0,x}$ for each output label x , read off at the joined constant pair. Proposition 8.9 shows that these isomorphisms form a closed isomorphism $\mathcal{K}^r$ of degree $|\Theta|$. The one-parameter families of Sections 8.5 and 8.6 supply the remaining homotopies.

All orientation maps below are defined up to positive scalars and take values in the lines of Section 8.1.

Separate maps. For a configuration c on $W_{a,R}$, $0 \leq R \leq \infty$, Proposition 5.4 gives

$$\Gamma_c^a : L_G(a_a) \otimes \det D_c \longrightarrow L_a(x_a).$$

For a configuration c on the closed-cut piece M_λ , with ordinary input β , the same proposition, with the positive orientation of Lemma 8.2 at the terminal constant pair, gives

$$\Gamma_c^M : L_G^M(\beta) \otimes \det D_c \longrightarrow Z_x.$$

Strip actions. For $b \in \{-, +, M\}$ and a regular strip u from x to y at the output of the corresponding domain, glue u to the constant pair at x and evaluate Γ^b on the positive constant element. This defines an even map

$$R_b(u) : L_b(x) \otimes \det D_u \longrightarrow L_b(y),$$

with $L_M(x) = Z_x$. Proposition 5.4 gives the coherence

$$(8.1) \quad \Gamma_{c\#u}^b(i \otimes (d\#e)) = R_b(u)(\Gamma_c^b(i \otimes d) \otimes e),$$

where $d\#e$ is the ordered determinant gluing of d and e . The differentials of $C_{S,\pm}$ and C_Z count index-one strips with these actions. For a split product strip $u = (u_-, u_+)$ put $n_- = \text{ind } D_{u_-}$ and

$$(8.2) \quad R_V(u)(v_- \otimes v_+ \otimes e_- \otimes e_+) = (-1)^{|v_+|n_-} R_-(u_-)(v_- \otimes e_-) \otimes R_+(u_+)(v_+ \otimes e_+).$$

The sign is the move of $\det D_{u_-}$ past $L_+(x_+)$. A regular index-one product strip for split data has exactly one nonconstant factor, of index one, since each nonconstant factor has a translation vector in its kernel and a constant factor at a transverse intersection is invertible. Thus (8.2) gives the right tensor differential (3.3) on C_V .

The decorated closed-cut map. A configuration on W_λ^J is terminal, in the sense of Definition 5.5, when it agrees with a constant pair $C_{M,x}^c$, for some representative c of the critical point over x , on the part of M_λ beyond a parallel copy of Z in the product collar of the cut, and on a two-sided collar of that copy. Its remaining factor e , the reverse excision cobordism enlarged by the collar, is an arbitrary configuration with the given inputs. Ordinary gluing at the copy of Z identifies its determinant line with $\det D_e \otimes \det D_{C_{M,x}^c}$, and the lines at c are identified with those at β_x as in Section 5. Let $\gamma_e : I_a \otimes \det D_e \rightarrow L_G^{\bar{e}}(\beta_x)$ be the almost-complex orientation of the gauge cobordism, with the ordered inputs, and let $\Theta_x : L_G^{\bar{e}}(\beta_x) \otimes Q_x \rightarrow Z_x$ be the conversion. Define

$$(8.3) \quad G_c(i \otimes \vartheta \otimes d_e \otimes 1) = (-1)^{|\Theta| \text{ind } D_e} \Theta_x(\gamma_e(i \otimes d_e) \otimes \vartheta).$$

The sign moves Q_x from its fixed position after the input line to the position adjacent to the cut. For a general configuration c on W_λ^J , the localized several-input shift of Section 5 (Definition 5.5) and Proposition 2.2 retains the ordinary cobordism and reaches a terminal configuration. Transport (8.3) back along it. This defines an even map

$$(8.4) \quad G_c : I_a \otimes Q_x \otimes \det D_c \longrightarrow Z_x.$$

Lemma 8.3. *The map (8.4) is independent of the shifting path, natural under ϕ_J , and continuous on compact families of finite widths. It has the following compatibilities.*

(i) *If A is an ordinary path at the input, then*

$$G_{A\#c}(i \otimes \vartheta \otimes (d_A\#d)) = (-1)^{|\Theta| \text{ind } D_A} G_c(\gamma_A(i \otimes d_A) \otimes \vartheta \otimes d).$$

(ii) *If u is a regular strip at the output, from x to y , then, with the coherent orientations ϑ_x, ϑ_y ,*

$$G_{c\#u}(i \otimes \vartheta_y \otimes (d\#e)) = R_M(u)(G_c(i \otimes \vartheta_x \otimes d) \otimes e).$$

- (iii) If $c = A \#_Z c'$ is an ordinary gluing at the closed cut of a configuration A on $\bar{\mathcal{E}}_\lambda$ from a to β and a configuration c' on M_λ from β to x , then, with the coherent orientations $\vartheta = \vartheta_x$ and ϑ_β ,

$$G_c(i \otimes \vartheta \otimes (d_A \# d')) = (-1)^{|\Theta| \text{ind } D_A} \Gamma_{c'}^M(\Theta_\beta(\gamma_A(i \otimes d_A) \otimes \vartheta_\beta) \otimes d').$$

Proof. Path independence is the square argument in the proof of Proposition 5.4, with the relative localized shift of Proposition 2.2(b). Apply it to a path whose endpoints are terminal. The endpoint tracks and the final edge remain terminal, where (8.3) is continuous. The orientation double cover over the parameter square then shows that the original path preserves (8.3). Applying this to the concatenation of two shifting paths from c proves independence. The bundle automorphisms of the lifting group ϕ_J preserve the almost-complex orientations, the lifting-class sign of the conversion on ϕ_Z and the positive constant element, and they map terminal configurations to terminal ones. The pullback of a shifting path by such an automorphism is a shifting path, so path independence shows that they preserve (8.4).

For (i), the path A is attached at an original input, away from the localized shift. At a terminal configuration, ordinary associativity gives $\gamma_{A \# e}(i \otimes (d_A \# d_e)) = \gamma_e(\gamma_A(i \otimes d_A) \otimes d_e)$, and the two sides of (i) carry the exponents $|\Theta|(\text{ind } D_A + \text{ind } D_e)$ and $|\Theta| \text{ind } D_A + |\Theta| \text{ind } D_e$. Continuity along the shift gives (i).

For (ii), glue u to the whole shifting family of c with one long neck. At the terminal end the configuration is $e \# C_{M,x}^c \# u$; shift $C_{M,x}^c \# u$ to $B \# C_{M,y}^{c'}$ with B an ordinary path on Z of index $n = \text{ind } D_u$, and insert a long product neck at Z so that B lies on the $\bar{\mathcal{E}}_\lambda$ side beyond a two-sided collar of Z ; then $e \# B \# C_{M,y}^{c'}$ is terminal. The ordinary action of B in the two structures satisfies the line form of the right conversion convention (4.2):

$$\gamma_B^M(\Theta_x(v \otimes \vartheta_x) \otimes d_B) = (-1)^{|\Theta|n} \Theta_y(\gamma_B^{\bar{\mathcal{E}}}(v \otimes d_B) \otimes \vartheta_y).$$

The right side of (ii) at the terminal configuration therefore carries the exponent $|\Theta| \text{ind } D_e + |\Theta|n$. The left side places Q before the determinant of $e \# B$, with exponent $|\Theta|(\text{ind } D_e + n)$. The other factors agree by associativity and the definition of $R_M(u)$. Continuity gives (ii).

For (iii), shift c' on M_λ to $B \# C_{M,x}^c$ and glue A to the whole family at Z with one long neck. Both sides of (iii) are continuous orientations of the same line along this family. At its end the configuration $A \# B \# C_{M,x}^c$ is terminal, and by (8.3) its left side carries $(-1)^{|\Theta|(\text{ind } D_A + \text{ind } D_B)}$. On the right, $\Gamma_{B \# C_{M,x}^c}^M$ is the action of B in the structure of M_λ , which by the display above contributes $(-1)^{|\Theta| \text{ind } D_B}$; the prefactor contributes $(-1)^{|\Theta| \text{ind } D_A}$. The two sides agree at the end, hence along the family. $\square$

The reference comparison. Let $C_{J,x}^{(\lambda)}$ be the joined constant pair on W_λ^J with output x . Its incoming labels are the critical points $a_{x\pm}$, so its input line is $L_G(a_{x-}) \otimes L_G(a_{x+}) = V_x$. Its operator is invertible for every λ (Lemma 8.4(iv) below). Evaluating (8.4) on the positive element gives an even isomorphism $G_{\text{ref}} : V_x \otimes Q_x \rightarrow Z_x$. Define the even map

$$(8.5) \quad K_{0,x} : Z_x \longrightarrow Q_x \otimes V_x, \quad K_{0,x} G_{\text{ref}}(v \otimes \vartheta) = (-1)^{|\Theta||v|} \vartheta \otimes v.$$

The lines are fixed and G_{ref} varies continuously with λ , so $K_{0,x}$ does not depend on λ . Replacing $C_{J,x}^{(\lambda)}$ by the joined constant pair of another representative conjugates the operator by a gauge transformation. It is assembled from the automorphisms relating the representatives on the two sides; these are central on T , and after multiplying one of them by -1 they agree there and glue across the bridge. The positive elements correspond, and so do the identified lines, by the remark on representatives in Appendix B; hence $K_{0,x}$ is unchanged.

Let $K_x^\vartheta : Z_x \rightarrow V_x$ be the map defined by $K_{0,x}(z) = \vartheta_x \otimes K_x^\vartheta(z)$, that is, $K_{0,x}$ followed by the isomorphism $Q_x \otimes V_x \rightarrow V_x$, $\vartheta_x \otimes v \mapsto v$. It has degree $|\Theta|$. Put

$$(8.6) \quad \mathcal{K} = K^\vartheta, \quad \mathcal{K}^r = K^\vartheta \text{gr}_{C_Z}^{|\Theta|}.$$

The isomorphism $\vartheta_x \otimes v \mapsto v$ involves no sign. The map $\vartheta_x^* \otimes^r 1$ of (3.4), with $\vartheta_x^*(\vartheta_x) = 1$, differs from it by $(-1)^{|\Theta||v|}$.

8.3. Product orientations across torus completion. On the completed pair $W_{-, \infty} \sqcup W_{+, \infty}$ the product of the separate maps is

$$(8.7) \quad \Pi_{(c_-, c_+)}(i_- \otimes i_+ \otimes d_- \otimes d_+) = (-1)^{|i_+| \text{ind } D_{c_-}} \Gamma_{c_-}^-(i_- \otimes d_-) \otimes \Gamma_{c_+}^+(i_+ \otimes d_+).$$

We extend it to large finite widths through the torus cutoff of Paper I. On the far torus tail of a completed side write a configuration as $\alpha + b$ in the torus-fibre slice of (I12). A compact family is *in torus normal form beyond r_0* if every member has this form for $r \geq r_0$, with b bounded in every weighted norm and with the weighted norms of b on $r \geq r_1$ tending to zero uniformly as $r_1 \rightarrow \infty$. For a pair $c = (c_-, c_+)$ in torus normal form beyond r_0 and $\lambda \geq 2r_0 + 2$, the preglue $P_\lambda(c)$ replaces b by $\chi_\lambda b$ on both tails, with $\chi_\lambda(r) = \chi(r - \lambda/2)$, and joins them by α on the bridge, as in Paper I, Lemma 10.25. On one side the truncation $P_R(c_a)$ on $W_{a,R}$ is defined in the same way; it equals α near the face and so satisfies the flat-fibre condition.

Lemma 8.4. (i) *Every compact family on the completed domains, in the topology of Section 2, is in torus normal form beyond some r_0 after a continuous family of small even gauge transformations which are the identity for $r \leq r_0 - 1$ and on the ordinary ends.*

(ii) *Let $\mathfrak{F}$ be a compact family of pairs in torus normal form beyond r_0 . The residual $\mathcal{F}(P_\lambda(c)) - \chi_\lambda \mathcal{F}(c)$, with χ_λ extended by 1 off the tails, is supported where $\chi'_\lambda \neq 0$ or $\chi_\lambda^2 \neq \chi_\lambda$, and tends to zero in every weighted norm, uniformly on $\mathfrak{F}$, as $\lambda \rightarrow \infty$. The same holds for P_R on one side.*

(iii) *The constant pairs $C_{a,x_a} = C^{c_a}$ whose representatives c_a are α on T , such as those produced in Appendix B, are equal to α on the whole added tail. Then $P_R(C_{a,x_a})$ is the constant pair on $W_{a,R}$, and $P_\lambda(C_{-,x_-}, C_{+,x_+}) = C_{J,x}^{(\lambda)}$ is the joined constant pair.*

(iv) *The operators at these constant pairs are invertible for every R and every λ .*

Proof. (i) Continuity and compactness make the set of tail differences $A - \alpha$ compact in H_δ^m . Their norms on $r \geq r_1$ therefore tend to zero uniformly as $r_1 \rightarrow \infty$. Sobolev embedding on unit boxes, whose geometry is fixed on the flat tail, makes them uniformly small in C^1 beyond some r_1 . By (I12) there is then a unique small even gauge transformation into the torus-fibre slice on $r \geq r_1$, continuous in the family. Cut off its logarithm on $[r_1, r_1 + 1]$, as in Paper I, Lemma 10.11, and take $r_0 = r_1 + 1$. The resulting gauges are connected to the identity by the path of logarithms, so determinant transport along them is canonical, and every orientation map below is natural under them.

(ii) Only the tail is changed. There the residual is $(d\chi_\lambda \wedge b)^+ + (\chi_\lambda^2 - \chi_\lambda)(b \wedge b)^+$, which is bounded by the norms of b on $r \geq \lambda/2 - 1$.

(iii) Choose the stationary representatives equal to α on the whole torus region of the slice, not only on the far collar. For a stationary representative $B = B_T + B_r dr$, flat on the torus band $[b, R]$ and equal to α near $r = R$ (Paper I, Proposition 10.7), solve $\partial_r g = -B_r g$, $g(R) = 1$, on a slightly larger interval $[c, R]$ in the determinant-one gauge bundle. Then B^g has no radial component there, and flatness gives $B^g = \alpha$ on $T \times [b, R]$. With a cutoff η equal to zero near c and one near $[b, \infty)$, the gauges $q_t(r) = g(R + t\eta(r)(r - R))$, $t \in [0, 1]$,

form a path from the identity, supported in a fixed collar, with $B^{q_1} = \alpha$ on $T \times [b, \infty)$. They preserve the classes of the perturbed-flat representatives and the genus germs, and depend smoothly on compact families of primary data. The constant pairs are pullbacks of these representatives, so they equal α on the tail and the truncation and the preglue have no cutoff error; the results are the constant pairs on $W_{a,R}$ and W_λ^J , which are solutions.

(iv) This is Lemma A.2. $\square$

The next lemma is the linear gluing across the torus bridge. The bridge is not an ordinary end, since its cross-section $T \times \mathbb{R}_\theta$ runs into the mixed ends, so Lemma 3.3 does not apply to it.

Lemma 8.5 (Determinant lines across torus completion). *Let $\mathfrak{F}$ be a compact family of pairs in torus normal form beyond r_0 . For $c \in \mathfrak{F}$ let $D_\lambda(c)$ be the linearized mixed operator at $P_\lambda(c)$, in the realization of Definition 5.1 with the end weights of Section 5.2, and let $D_\infty(c) = D_{c_-} \oplus D_{c_+}$.*

- (i) *There are a finite-dimensional space E , a linear map $\Lambda : E \rightarrow Y$ into smooth target sections supported in a compact subset K_Λ of the completed domains, chosen after $\mathfrak{F}$, and λ_1 with $K_\Lambda \subset \{r \leq \lambda_1/2\}$, such that $D_\lambda(c) \oplus \Lambda$ is onto, with a uniformly bounded right inverse, for all $c \in \mathfrak{F}$ and $\lambda \in [\lambda_1, \infty]$. If every $D_\infty(c)$ is onto, one may take $E = 0$.*
- (ii) *The kernels of $D_\lambda(c) \oplus \Lambda$ form a continuous vector bundle over $\mathfrak{F} \times [\lambda_1, \infty]$; in particular $\text{ind } D_\lambda(c) = \text{ind } D_\infty(c)$ for $\lambda \geq \lambda_1$. Near $\lambda = \infty$ its local frames are $\iota_\lambda(v) = \beta_\lambda v - Q_\lambda(D_\lambda \oplus \Lambda)\beta_\lambda v$, where β_λ is multiplication by χ_λ on both tails and Q_λ is the right inverse of (i).*
- (iii) *Consequently the lines $\det D_\lambda(c)$ form a continuous line bundle over $\mathfrak{F} \times [\lambda_1, \infty]$ whose fibre at $\lambda = \infty$ is $\det D_{c_-} \otimes \det D_{c_+}$. This structure does not depend on Λ . Where the operators are invertible, the canonical positive elements form a continuous section.*
- (iv) *Let an ordinary path or strip be glued, with a long neck, at an original gauge input or at the strip output, to every member of $\mathfrak{F}$. The identification of (iii) commutes with the ordinary gluing isomorphism of Lemma 3.3.*

The same statements hold for a compact family in torus normal form on one completed side and its truncations $P_R(c)$ on $W_{a,R}$, $R \rightarrow \infty$. They also hold for every family of operators with the properties listed at the beginning of its proof in Appendix A.9, in particular for the linearizations in the torus-completion chart (Paper I, Definition 10.23) and for the interpolating operators in the proof of Lemma 8.10.

The proof is given in Appendix A.9. Once a uniformly bounded right inverse is available, the construction is the standard gluing of determinant lines along a long neck, as in Lemma 3.3. Two points are specific to the torus bridge. First, the uniform estimates must be proved directly, because the cross-section of the bridge runs into the mixed ends; they are obtained by patching the flat odd-torus model on the bridge with the stationary models on the mixed ends, as for the uniform inverse of Paper I, Proposition 10.27. Second, the stationary columns at finite and at infinite width differ on the whole mixed ends, not only on the tails; their difference tends to zero exponentially in the width, because of the fibre gap μ_T on the torus neck.

Let $\mathfrak{F}$ be a compact family of pairs of configurations on the completed domains. By Lemma 8.4(i) we may assume it is in torus normal form. Examples are a pair of index-zero solutions and a pair of constant pairs with strips glued at their outputs. For $c \in \mathfrak{F}$ define

$$(8.8) \quad \Pi_{P_\lambda(c)} : I_a \otimes \det D_{P_\lambda(c)} \longrightarrow V_x, \quad \lambda \geq \lambda_1,$$

as the continuous extension of (8.7) along $\{P_\lambda(c)\}_{\lambda \in [\lambda_1, \infty]}$ given by Lemma 8.5. It is continuous on $\mathfrak{F} \times [\lambda_1, \infty]$ and natural under the joined lifting group $\phi_J \cong \phi_- \oplus \phi_+$.

The separate shifts of Proposition 2.2 give, for $c \in \mathfrak{F}$, a homotopy $\nu \mapsto H_\nu(c)$, $\nu \in [0, 1]$, continuous in the topology of Section 2, from c to $(A_- \# C_{-,x_-}^{c_-}, A_+ \# C_{+,x_+}^{c_+})$ with $A_\pm$ ordinary paths and representatives $c_\pm$ which are α on T . Its image is compact. By Lemma 8.4(i) it can be put in torus normal form beyond a common r_0 , and by Lemma 8.4(iii) its endpoints equal α on the tails. Applying (8.8) to this larger compact family gives Π on the rectangle $\{P_\lambda(H_\nu(c))\}$. At $\lambda = \infty$ it is (8.7), which along ν is the transport defining $\Gamma^\pm$ in Proposition 5.4. At $\nu = 1$ the configuration is $(A_- \sqcup A_+) \# C_{J,x}^{(\lambda)}$, and there Π is the ordinary orientation of $A_- \sqcup A_+$ followed by the positive element of the invertible constant factor, by Lemma 8.5(iii),(iv). The rectangle is contractible, so this description of Π at $P_\lambda(c)$ through $\nu = 1$ does not depend on the chosen shifting homotopy.

The same device compares the separate orientations on the finite and completed domains. By construction, Γ^a on the finite domain $W_{a,R}$ is a continuous orientation over its configuration space, defined by the one-input shift on that domain.

Lemma 8.6. *Let $\mathfrak{F}$ be a compact family on $W_{a,\infty}$ in torus normal form. For $R \geq R_1$, the orientation Γ^a of $W_{a,R}$ at $P_R(c)$ is the continuous extension of Γ_c^a along $\{P_R(c)\}_{R \in [R_1, \infty]}$ given by Lemma 8.5.*

Proof. Consider the rectangle $\{P_R(H_\nu(c))\}$, with H_ν the separate shift in torus normal form. The extension from $R = \infty$ is continuous on it by Lemma 8.5. For fixed finite R the orientation Γ^a of $W_{a,R}$ is continuous in ν . At $\nu = 1$ the configuration is $A \# P_R(C_{a,x})$, and $P_R(C_{a,x})$ is the constant pair of $W_{a,R}$ by Lemma 8.4(iii). There Γ^a is the ordinary orientation of A followed by the positive element of the constant factor. The extension from $R = \infty$ is the same: by Lemma 8.5(iv) it commutes with gluing A , and the constant factor is invertible for every R (Lemma A.2) with positive element continuous up to $R = \infty$. The two orientations therefore agree at $\nu = 1$, and hence on the whole rectangle. $\square$

Corollary 8.7. *The strip actions $R_\pm(u)$ computed on $W_{\pm,R}$ do not depend on $R \in [0, \infty]$, and $R_M(u)$ computed on M_λ does not depend on λ and agrees with the action on the special quintuple of (D_0, D_1) defining N_Z . Thus each of $C_{S,\pm}$ and C_Z carries one orientation system.*

Proof. Each action is an isomorphism between fixed lines, defined through an invertible constant pair with a strip glued at its output. For finite R it varies continuously with R , since the constant pairs are invertible for every R (Lemma A.2); at $R = \infty$ use Lemma 8.6 with $\mathfrak{F} = \{C_{a,x} \# u\}$, which is α on the tail, together with Lemma 8.5(iv) for the strip glued at the output. For R_M use the width interval and the family of Paper I, Lemma 9.7, along which $C_{M,x}$ stays invertible by Lemma 8.2. A continuous family of isomorphisms between fixed lines is constant. $\square$

Proposition 8.8. *On the family $P_\lambda(H_\nu(c))$, $c \in \mathfrak{F}$, for $\lambda \geq \lambda_1$ and $\nu \in [0, 1]$,*

$$(8.9) \quad K_{0,x} G_c = (1_{Q_x} \otimes \Pi_c)(\sigma_{I,Q} \otimes 1_{\det}),$$

where $\sigma_{I,Q}$ moves Q_x past the entire input line. Let $G_c^\vartheta(i \otimes d) = G_c(i \otimes \vartheta_x \otimes d)$ be the map obtained by evaluating G_c on ϑ_x in its middle factor, and let $m = \text{ind } D_c \pmod{2}$. Then

$$(8.10) \quad \mathcal{K} G_c^\vartheta = (-1)^{|\Theta| |I|} \Pi_c, \quad \mathcal{K}^r G_c^\vartheta = (-1)^{|\Theta| (m+1)} \Pi_c.$$

Proof. Both sides of (8.9) are continuous on the rectangle at finite λ , so it suffices to verify it at $\nu = 1$. There the configuration is $(A_- \sqcup A_+) \# C_{J,x}^{(\lambda)}$. Let k_A be the total parity of

$\det D_{A_-} \otimes \det D_{A_+}$ and let $v = \gamma_A(i \otimes d_A) \in V_x$, so $|v| = |i| + k_A$. By Lemma 8.3(i), $G_c(i \otimes \vartheta \otimes (d_A \# 1)) = (-1)^{|\Theta|k_A} G_{\text{ref}}(v \otimes \vartheta)$, and by (8.5)

$$K_{0,x} G_c(i \otimes \vartheta \otimes (d_A \# 1)) = (-1)^{|\Theta|(k_A+|v|)} \vartheta \otimes v = (-1)^{|\Theta||i|} \vartheta \otimes \Pi_c(i \otimes d).$$

This is the right side of (8.9). For the second assertion, the tensor of maps $1_Q \otimes \Pi_c$ with even Π_c introduces no sign, so removing the factor ϑ_x as in the definition of K_x^ϑ gives $\mathcal{K} G_c^\vartheta = (-1)^{|\Theta||I|} \Pi_c$. Since G_c is even, $|G_c^\vartheta(i \otimes d)| = |i| + |\Theta| + m$, and $\text{gr}^{|\Theta|}$ contributes $(-1)^{|\Theta|(|i|+|\Theta|+m)}$. The product is $(-1)^{|\Theta|(|\Theta|+m)} = (-1)^{|\Theta|(m+1)}$. $\square$

At index zero the coefficient in (8.10) is $(-1)^{|\Theta|}$; at the isolated points of a one-parameter family, where m is odd, it is $+1$.

Proposition 8.9. $\mathcal{K}^r : C_Z \rightarrow C_V$ is an isomorphism of graded groups, and

$$d_V \mathcal{K}^r = (-1)^{|\Theta|} \mathcal{K}^r d_Z.$$

Thus $\mathcal{K}^r$ is a closed isomorphism of degree $|\Theta|$.

Proof. $K_{0,x}$ is an isomorphism of lines, and the generators of C_Z and C_V are both indexed by $A_0 \cap A_1$; hence $\mathcal{K}^r$ is an isomorphism. Let $u = (u_-, u_+)$ be a regular index-one product strip from x to y for the split data. Apply Proposition 8.8 to the pair of separate configurations $(C_{-,x_-} \# u_-, C_{+,x_+} \# u_+)$, which is in torus normal form, and to its preglue $C_{J,x}^{(\lambda)} \# u$. By (8.1) on each side and (8.7), the product orientation at $\lambda = \infty$ is

$$\Pi_{C_J \# u}(v \otimes (1 \# e)) = R_V(u)(v \otimes e).$$

Indeed, $\Gamma_{C_\pm}^\pm(v_\pm \otimes 1) = v_\pm$, the determinant of $C_\pm \# u_\pm$ is $(1 \# e_\pm)$ with $\text{ind } D_{C_- \# u_-} = n_-$, and (8.1) gives $\Gamma^\pm(v_\pm \otimes (1 \# e_\pm)) = R_\pm(u_\pm)(v_\pm \otimes e_\pm)$. Formula (8.7) with $i_\pm = v_\pm$ then contributes $(-1)^{|v_+|n_-}$, which is the sign in (8.2); reordering the even constant elements costs nothing. By continuity in λ , and Lemma 8.5(iv) for the strip glued at the output, the identity holds at large finite widths. By Lemma 8.3(ii) with $c = C_{J,x}^{(\lambda)}$,

$$G_{C_J \# u}(v \otimes \vartheta_y \otimes (1 \# e)) = R_M(u)(G_{\text{ref}}(v \otimes \vartheta_x) \otimes e).$$

Apply $K_{0,y}$ and use (8.9) for $C_J \# u$ and (8.5). Both sides of the resulting identity, evaluated on $G_{\text{ref}}(v \otimes \vartheta_x) \otimes e$, equal $(-1)^{|\Theta||v|} \vartheta_y \otimes R_V(u)(v \otimes e)$. Let $J_u : Q_x \rightarrow Q_y$ be the conversion transport along the ordinary path obtained by shifting $C_{M,x} \# u$; by coherence $J_u \vartheta_x = \vartheta_y$. Hence

$$K_{0,y} R_M(u) = (J_u \otimes R_V(u))(K_{0,x} \otimes 1_{\det D_u}),$$

with Q to the left of both V and the strip determinant; since $R_V(u)$ is even, the tensor of maps inserts no sign. Removing the factors ϑ_x and ϑ_y as in the definition of K^ϑ gives $K_y^\vartheta R_M(u)(z \otimes e) = R_V(u)(K_x^\vartheta z \otimes e)$. Summing over index-one strips gives $d_V \mathcal{K} = \mathcal{K} d_Z$. Since $d_Z \text{gr}^{|\Theta|} = (-1)^{|\Theta|} \text{gr}^{|\Theta|} d_Z$, the assertion for $\mathcal{K}^r$ follows. $\square$

8.4. Signed counts at large torus length. Fix regular data in the following order, as in Paper I, Section 10.11: the separate mixed maps $N_\pm$, the closed-cut and joined data at a width $\lambda_0 \geq \lambda_*$, regular data on the completed domains, and lengths $R \geq R_0$ and $\lambda \geq \lambda_1$. The completed index-zero sets are finite and regular, and there are no completed solutions of negative index. By the torus completion (I3) (Paper I, Theorem 10.32(i) and (ii)), the gluing maps $\mathcal{G}_R$ and $\mathcal{G}_\lambda$ are bijections onto the regular index-zero solutions on $W_{a,R}$ and W_λ^J , and there are no solutions of negative index on these domains. Write $N_{a,\infty}$ for the completed separate counts, oriented by Γ^a , and $N_a^{(R)}$ for the counts on $W_{a,R}$. On W_λ^J let $F_\lambda^\Pi : C_I \rightarrow C_V$

count index-zero solutions oriented by Π , and $F_\lambda^Z : C_I \rightarrow C_Z$ the same solutions oriented by G^θ ; their degrees are 0 and $|\Theta|$. The orientation Π at these solutions is defined by the following lemma.

Lemma 8.10. *Let V be a completed index-zero solution, or a pair of them. For $L \geq L_1$ let $\mathcal{G}_L(V) = \Psi_L(z_L(V))$ be the glued solution of (I3), where Ψ_L is the gluing chart at the preglued configuration $P_L(V)$ and $z_L(V)$ is the unique small zero in that chart (Paper I, Proposition 10.30). Consider the rectangle of configurations*

$$\mathcal{R}_V = \{P_L(V) : L \in [L_1, \infty]\} \cup \{\Psi_L(sz_L(V)) : s \in [0, 1], L \in [L_1, \infty)\},$$

with $P_\infty(V) = V$. Here V is represented in the gauge in which its torus tail lies in the torus-fibre slice (the adapted representative of Paper I, Lemmas 10.11 and 10.25), so that the two edges meet, and the gauge transformation of Lemma 8.4(i) is the identity by the uniqueness of that slice (II2). The determinant lines of the linearized mixed operators form a continuous line bundle over $\mathcal{R}_V$, and determinant transport carries the canonical positive element at V to the canonical positive element at $\mathcal{G}_L(V)$.

Proof. Over the first edge the line bundle is that of Lemma 8.5 with $\mathfrak{F} = \{V\}$, and its extension clause covers the chart operators used below; the second edge is a continuous family on a fixed domain for each L .

Use also the chart operators T of Paper I, Definition 10.23. At the preglues they are invertible with a uniform bound by Paper I, Proposition 10.27; at $L = \infty$ the chart is centred at V and the operator is invertible, since it has index zero and is injective at the regular completed solution V , by the argument of Paper I, Proposition 10.27. Along the contraction path the chart operator is $D_L + DN_L(sz_L)$ with $\|DN_L(sz_L)\| \leq 2C_{\text{quad}}r \leq (2C_{\text{inv}})^{-1}$, by the quadratic estimate of Paper I, Lemma 10.29(c), and the choice of radius in Paper I, Proposition 10.30, so it is invertible. The chart operator at $\Psi_L(sz)$ is centred at the preglue. Moving its slice centre, trace chart and target projection to the configuration itself changes it by at most $C\|sz\|$ in operator norm, which is small, so the straight path between them consists of invertible operators. Here both operators are read in the fixed end trivializations, and $\|D\Psi_L(sz) - D\Psi_L(0)\| \leq C\|sz\|$ by the second-derivative bounds of the face chart (Paper I, Lemma 10.28(c)) and of the genus chart (Paper I, Lemma 10.29(c)). The resulting operator and the realization of Definition 5.1 at the same configuration have the same decaying block; their stationary columns and scalar rows differ by a bounded finite-rank map (Paper I, Remark 10.24). The straight path between these two is therefore a continuous Fredholm path, and these identifications are continuous over $\mathcal{R}_V$. The positive elements of the invertible chart operators thus give a continuous nonvanishing section of the determinant bundle of the linearized mixed operators.

At the two solutions V and $\mathcal{G}_L(V)$ this section is the canonical positive element. Indeed, there the change of stationary columns is unipotent: the chart lifts and the Coulomb lifts differ by infinitesimal gauge directions $d\psi_e$ and by decaying fields already in the domain; this holds because both lift the same deformation of the stationary limit with zero normal component, and the corrected lifts differ from the uncorrected ones by decaying fields (Paper I, (10.5)). A unipotent change of source coordinates preserves the positive element. The change of the scalar row is block triangular with invertible gauge block, as in the proof of Lemma 5.2. At V the gauge block is the Neumann map of Paper I, Lemma 10.26(ii), with $L = \infty$. At $\mathcal{G}_L(V)$ the C^m hypothesis of that lemma is not available, since Paper I, Proposition 10.30, bounds $z_L(V)$ only in X_L ; there the gauge block is invertible by Proposition 2.1(b), the connection

being irreducible on the odd surface and torus collars. At these points every operator of the interpolating paths is invertible. $\square$

Extend (8.8) from $P_\lambda(V)$ to $\mathcal{G}_\lambda(V)$ along the second edge of $\mathcal{R}_V$. Both sides of (8.9) are continuous along it, so the comparison persists at $\mathcal{G}_\lambda(V)$.

Proposition 8.11. *Fix the data as at the beginning of this subsection, and let $N_a^{(R)}$, $N_{a,\infty}$, F_λ^Π and F_λ^Z be the counts defined there, oriented by Γ^a , Π and G^ϑ . Then, for R and λ sufficiently large,*

$$N_a^{(R)} = N_{a,\infty} \ (a = \pm), \quad F_\lambda^\Pi = N_{-, \infty} \otimes^r N_{+, \infty}, \quad \mathcal{K}^r F_\lambda^Z = (-1)^{|\Theta|} N_{-, \infty} \otimes^r N_{+, \infty}.$$

Proof. Let V be a completed index-zero solution on a separate side. By Lemma 8.6 with $\mathfrak{F} = \{V\}$, the orientation Γ^a of $W_{a,R}$ at $P_R(V)$ is the continuous extension of Γ_V^a ; along the second edge of $\mathcal{R}_V$ it is continuous on the fixed domain. By Lemma 8.10 the canonical positive element is also transported continuously from V to $\mathcal{G}_R(V)$. Hence $\Gamma^a(i \otimes 1)$ has the same sign at $\mathcal{G}_R(V)$ as at V . The bijection $\mathcal{G}_R$ preserves labels and lifting orbits, so the counts agree entry by entry.

In the joined case Π at $\mathcal{G}_\lambda(V)$ is, by definition, the continuous extension of (8.7) over $\mathcal{R}_V$, and the same lemma transports the positive element. At index zero $\text{ind } D_{V_-} = 0$, so (8.7) evaluated on the positive element is the product of the two separate signs. The joined lifting group is $\phi_- \oplus \phi_+$, and $\mathcal{G}_\lambda$ is a bijection from pairs of classes with no further gluing parameter (13), so the orbit counts multiply. The generator factors in (8.7) are those of the right tensor of maps (3.4) of the even maps $N_{\pm, \infty}$. This proves the second identity. The third follows from the second and (8.10) with $m = 0$, applied solution by solution, since (8.9) holds at $\mathcal{G}_\lambda(V)$. $\square$

8.5. Boundary orientation of the comparison families. The comparison families below are those of Paper I: Propositions 10.16 and 10.17 for collar length and width, with the weighted parameter column of (19), Lemma 10.20 for the change of data at fixed width, the relative regularization of (17) (Paper I, Theorem 5.4) for compactly supported changes of secondary data, and the mixed-end family of the deformation invariance (II) of the mixed map, as in Lemma 7.6(V2), which produces H_M in Proposition 8.14. In each case the data are constant near $b = 0, 1$ and the endpoints are regular in index at most one. A relative invariant regularization then gives finitely many isolated parametrized solutions of vertical index -1 , each with invertible augmented operator $(v, \tau) \mapsto Dv + \tau C$. It also gives a compact one-manifold of parametrized index-zero solutions, whose boundary consists of the two endpoint fibres and one branch for each pair of an index -1 solution with an index-one trajectory at an ordinary end. Here C is the parameter column. On a varying mixed end it is the weighted column of (19): since the stationary references solve the equation for every parameter, the parameter derivative lies in the weighted target, although it is not compactly supported.

Let Ω denote the orientation map in use: Γ^a on a separate domain, Γ^M on the closed-cut piece, or G on a joined domain. Put $f = 0$ in the first two cases and $f = |\Theta|$ in the third. The endpoint counts F_0, F_1 are oriented by Ω on the positive element. At an isolated parametrized solution c , orient the parameter first: choose $d_c \in \det D_c$ such that $db \otimes d_c$ is the positive element of the invertible augmented operator, and count c by $\Omega(i \otimes d_c)$, or $\Omega(i \otimes \vartheta \otimes d_c)$ when $f = |\Theta|$. Let H be the resulting count. Its degree is $f + 1$.

Lemma 8.12. *Let one of the comparison families above be regularized relatively as described, and let Ω , f , the endpoint counts F_0, F_1 and the count H of its isolated parametrized solutions*

of vertical index -1 be as above. Then

$$(8.11) \quad F_1 - F_0 = d_{\text{out}} H + (-1)^f H d_{\text{in}} = \delta H.$$

For a fixed regular datum, the index-one boundary gives $d_{\text{out}} F = (-1)^f F d_{\text{in}}$, so F is closed of degree f .

Proof. Orient the parametrized index-zero space by $\det(T_b[0, 1]) \otimes \det D$ and Ω , as in Section 5; the boundary is oriented by the outward normal first. At $b = 1$ the outward normal is ∂_b , and at $b = 0$ it is $-\partial_b$; there the vertical operators are invertible, so these contribute $F_1 - F_0$.

At an outgoing strip break (c, u) let p_T be the preglued configuration of the augmented pair, with the strip translated to $u(s - T, \cdot)$. In the gluing of (I7) (Paper I, Theorem 5.1(iv)), for an index -1 piece of a family, the glued solutions are $U_T = P_{p_T}(z_T)$. Here P_{p_T} is the chart of Paper I, Lemma 5.3, and $z_T \in \text{im } Q_{p_T}$, a parameter increment together with a field, is the fixed point

$$z_T = -Q'_{p_T}(\widehat{\mathcal{G}}_{p_T}(0) + N_{p_T}(z_T)).$$

The operator Q_{p_T} is the patched right inverse of the augmented operator at p_T , whose image is complementary to the cut-off translation vector, Q'_{p_T} is the right inverse of $D\widehat{\mathcal{G}}_{p_T}(0)$ with the same image, and the map $z \mapsto -Q'_{p_T}(\widehat{\mathcal{G}}_{p_T}(0) + N_{p_T}(z))$ is a contraction with constant at most $\frac{1}{2}$. Identify the domains by maps which fix the mixed core and translate the strip core, and patch Q_{p_T} with cutoffs translated together with the strip. Then $\widehat{\mathcal{G}}_{p_T}$, N_{p_T} and Q'_{p_T} are differentiable in T with bounded derivatives, the T -derivative of $\widehat{\mathcal{G}}_{p_T}(0)$ is exponentially small in T , and the T -derivative of p_T is the cut-off translation vector $-\partial_s u$ plus an error which is exponentially small in T . Differentiate the fixed-point identity in T and absorb the term containing the derivative of z_T by the contraction constant. This bounds the T -derivative of z_T by a constant times $\|\widehat{\mathcal{G}}_{p_T}(0)\| + \|z_T\|$ plus the T -derivative of $\widehat{\mathcal{G}}_{p_T}(0)$, and all three tend to zero. For pseudoholomorphic curves this T -derivative estimate, with exponential decay, is [FOOO24, Theorem 6.4]; the same argument applies with the fixed mixed core and on the fixed ordinary gauge end. Hence $\partial_T U_T = -\partial_s u + o(1)$ on the strip core after projection to the slice. Linear gluing (Lemma 3.3) identifies $\det A_c \otimes \det D_u$ with the determinant along the branch, and the positive translation vector $\partial_s u$ orients $\det D_u$. The outward direction at $T = \infty$ is ∂_T , which is negative in this order. With the parameter already first in $\det A_c = \det(T_b[0, 1]) \otimes \det D_c$ and (8.1) or Lemma 8.3(ii), this branch contributes $-d_{\text{out}} H$.

At an incoming break (A, c) , with A an index-one trajectory, the trajectory is translated to $A(s + T, \cdot)$, and the same argument gives $\partial_T U_T = +\partial_s A + o(1)$, positive in $\det D_A \otimes \det A_c$. The outward direction is again ∂_T . Moving the one-dimensional parameter line to the front past the odd line $\det D_A$ contributes -1 . The input compatibility of Ω (Proposition 5.4 for Γ , Lemma 8.3(i) for G) contributes $(-1)^f$, because $\text{ind } D_A = 1$. This branch contributes $-(-1)^f H d_{\text{in}}$. For a joined domain with two inputs, the other-arm permutations are those of the right tensor differential, by Lemma 3.2. The total signed boundary vanishes, which is (8.11). Every branch is counted once with its matched lifting orbit. The relative regularization of (I7) gives one branch for each pair (Paper I, Theorem 5.4(ii)). The gluing of (I7) matches lifting orbits, since the lifting group of the glued problem is the fibre product of those of the pieces (Paper I, Theorem 5.1(ii),(iv), under hypothesis (H5) there).

For a fixed datum the same computation without the parameter line gives $-d_{\text{out}} F + (-1)^f F d_{\text{in}} = 0$. The sign determinations use only the direction of the translation vector, the ordinary linear gluing of Lemma 3.3 at the ordinary end where the break occurs, and the

invertibility of the bordered augmented operator along the gluing paths. By Corollary 8.7 the output differential is the same at every point of the family. $\square$

For a simple check of the incoming sign, take $D_A : \mathbb{R}v \rightarrow 0$, $D_c : 0 \rightarrow \mathbb{R}y$ and $C : \mathbb{R}\mu \rightarrow \mathbb{R}y$ the identity. In the order $\det D_A \otimes \det A_c$ the kernel is oriented by v . In the parameter-first order $\mu \otimes v \otimes y^*$ cancelling μ against y leaves $-v$.

8.6. Chain homotopies across collar length, width and the closed neck.

Proposition 8.13. *Let $N_{\pm}$ be the regular mixed maps at $R = 0$, and fix a regular joined datum at a width $\lambda_0 > 0$. Choose R and a width λ_2 as in Proposition 8.11, and at these lengths add small invariant secondary perturbations supported in the cores which make index one regular and preserve the index-zero counts. Then there are maps $P_{\pm}$ of degree one and B_J of degree $|\Theta| + 1$ with*

$$N_{a,\infty} - N_a = \delta P_a, \quad F_{\lambda_2}^Z - F_{\lambda_0}^Z = \delta B_J.$$

Proof. The added perturbations exist as in Paper I, Section 10.11: the index-zero sets are finite and regular, so openness and compactness preserve them, their signs and the absence of negative indices. For the separate side, join the data at $R = 0$ to those at the large length by one family which first changes the collar length and then makes a compactly supported change of secondary data. Its endpoints are regular through index one. Paper I, Proposition 10.16, with the weighted parameter column of (I9), and the relative regularization of (I7) give its relative regularization and compactification, applied to the concatenated family, whose two segments each satisfy the hypotheses of (I9) (Paper I, Lemma 10.15), and Lemma 8.12 with $\Omega = \Gamma^a$ gives $N_a^{(R)} - N_a = \delta P_a$. Proposition 8.11 replaces $N_a^{(R)}$ by $N_{a,\infty}$. For the joined side use one family consisting of the width change of Paper I, Proposition 10.17, from λ_0 to λ_2 , followed by the change of secondary data at width λ_2 , regularize it relatively as a single family, and apply Lemma 8.12 with $\Omega = G$. The joined families lie at finite widths throughout, so G is defined along them. $\square$

With the right tensor convention, $\delta(F \otimes^r G) = (-1)^{|G|} \delta F \otimes^r G + F \otimes^r \delta G$. Since $N_{\pm}$ and $N_{\pm,\infty}$ are even and closed,

$$(8.12) \quad P = P_- \otimes^r N_+ + N_{-,\infty} \otimes^r P_+, \quad \delta P = N_{-,\infty} \otimes^r N_{+,\infty} - N_- \otimes^r N_+.$$

By Proposition 8.9 and (3.2), $\delta(\mathcal{K}^r B_J) = (-1)^{|\Theta|} \mathcal{K}^r \delta B_J$. Together with Proposition 8.11 this gives

$$(8.13) \quad (-1)^{|\Theta|} \mathcal{K}^r F_{\lambda_0}^Z - N_- \otimes^r N_+ = \delta(P - \mathcal{K}^r B_J).$$

Proposition 8.14. *Let N_M be the count on the closed-cut piece M_{λ_0} of the decomposition (I2), oriented by Γ^M . Its data are the ordinary datum f_Z and the metrics on Z used for N_Z at width λ_0 , and the regular invariant secondary data chosen on this piece in the closed-cut data (Paper I, Section 10.9). There are maps B_{cut} of degree $|\Theta| + 1$ and H_M of degree one with*

$$N_M \bar{E} - F_{\lambda_0}^Z = \delta B_{\text{cut}}, \quad N_Z - N_M = \delta H_M.$$

Proof. Paper I, Proposition 10.21, inserts a closed neck $[-\ell, \ell] \times Z$ carrying the cut-off term f_Z . It shows that for large ℓ , after a small invariant perturbation making index one regular, the joined index-zero solutions are exactly the gluings $A \#_Z c'$ of index-zero solutions on $\bar{E}$ and on M_{λ_0} , one for each matched lifting orbit. The lifting group is the fibre product over ϕ_Z and all three pieces are connected, so by Lemma D.1 the gluing fibre is trivial and the

quotient count is the matrix composite. By Lemma 8.3(iii), at index zero the orientation G of $A \#_Z c'$ is $\Gamma_{c'}^M$ applied to the right-placed conversion of γ_A ; the prefactor $(-1)^{|\Theta| \text{ind } D_A}$ is one. Linear gluing along the long neck is continuous and the operators along the gluing paths are invertible, so the sign is that of the composite. Hence the count at the long-neck data is $N_M \bar{E}$. Paper I, Lemma 10.20, connects the data at width λ_0 to the long-neck data by a family with compactly supported parameter derivative, and Lemma 8.12 with $\Omega = G$ gives B_{cut} .

For the second identity, the path of domains and data of the deformation invariance (II) of the mixed map, as applied in Paper I, Proposition 10.21, changes the finite torus-neck lengths of D_{i,λ_0} to those of the special quintuple of (D_0, D_1) , together with a compact interpolation of the metric, with the ordinary datum on Z fixed. This is a noncollapsing family with the weighted parameter column of (I9). Concatenate it with the compactly supported change from the secondary data of M_{λ_0} to those of N_Z (relative regularization (I7)), and regularize the concatenated family relatively as a single family; this is the family used in the proof of Lemma 7.6(V2), which applies verbatim to M_{λ_0} , a domain of DFL type in the class of (II) with the ordinary end data of N_Z , whose constant pairs are invertible along it by Lemma A.2(c); it is traversed in the opposite direction, so that $H_M = -H_{M_{\lambda_0}}$. Lemma 8.12 with $\Omega = \Gamma^M$ gives H_M . The input complex $C_G(Z)$ is the same at both ends, since f_Z is fixed. $\square$

Since $\bar{E}$, with the conversion at Z , is closed, $\delta(H_M \bar{E}) = (N_Z - N_M) \bar{E}$ by (3.2). With $B = B_{\text{cut}} + H_M \bar{E}$,

$$(8.14) \quad N_Z \bar{E} - F_{\lambda_0}^Z = \delta B.$$

8.7. The mixed excision theorem.

Theorem 8.15 (Mixed excision). *The mixed map N_Z of the folded pair (A_0, A_1) , composed with the reverse excision map $\bar{E}$, is chain homotopic over $\mathbb{Z}$ to $(-1)^{|\Theta|}$ times the tensor product $N_- \otimes^r N_+$ of the two separate mixed maps, once the target of N_Z is identified with the tensor product of their targets by the closed isomorphism $\mathcal{K}^r$. The homotopy is an explicit sum of signed counts. Precisely, with the conversion at the closed cut placed as in Convention 8.1 and $\mathcal{K}^r$ as in (8.6),*

$$(8.15) \quad (-1)^{|\Theta|} \mathcal{K}^r N_Z \bar{E} - N_- \otimes^r N_+ = \delta H, \quad H = P - \mathcal{K}^r B_J + \mathcal{K}^r B.$$

Each term of H is a signed count of isolated parametrized solutions in a one-parameter family, or its composite or right tensor product with closed counts and with $\mathcal{K}^r$. Here $\bar{E}$ is the reverse excision map at the closed-cut data of Proposition 8.14; by Proposition 2.3(c) it is chain homotopic to the map at the data of that proposition. In particular

$$(8.16) \quad \mathcal{K}^r N_Z \bar{E} \simeq (-1)^{|\Theta|} (N_- \otimes^r N_+).$$

For $z \in C_G(Z)$ in the structure of $\bar{\mathcal{E}}$, write $N_Z z$ for $N_Z \Theta^r z$, as in the composite $N_Z \bar{E}$. Then for a homogeneous cycle z ,

$$(8.17) \quad [\mathcal{K}^r N_Z z] = \pm[(N_- \otimes^r N_+) E z].$$

The gauge complexes of $(-M)^\#$ and $(M^+)^\#$ carry the orientations defined by their reference almost-complex structures.

Proof. By (3.2) and Proposition 8.9, $\delta(\mathcal{K}^r B) = (-1)^{|\Theta|} \mathcal{K}^r \delta B$. Adding this, applied to (8.14), to (8.13) gives (8.15), and (8.16) is its restatement. For (8.17) apply (8.16) to $E z$. After the continuation homotopy of Proposition 2.3(c), the capped comparison gives $[\bar{E} E z] = \pm[z]$

for this homogeneous class, with one sign (Proposition D.3), and applying $\mathcal{K}^r N_Z \Theta^r$ gives the claim. $\square$

The sign $(-1)^{|\Theta|}$ in (8.16) refers to the fixed choices of ordinary orientations, of the orientation of Q chosen in Section 8.1 and of the reference comparison (8.5). Changing one of these choices changes the identity by one overall sign; no choice changes individual matrix entries.

Equality (8.17) is equality of homology classes of tensors. If the tensors differ by a boundary, applying the chain isomorphism (3.6) produces a boundary in the Hom complex. This is the mechanism by which the cup formula of Proposition 9.3 becomes a chain homotopy. No inference from equality of induced maps on homology is involved.

9. THE COMPRESSION CUP BEFORE REFLECTION

This section proves the cup formula, Proposition 9.3: the compression cup, before a reflection pairing is chosen, is chain homotopic, up to one sign, to the composite of the transposed reflected mixed map of Section 9.3, the gauge one-handle map and the mixed map of the upper presentation. The two inputs are the oriented folded triangle comparison, Lemma 9.2, and mixed excision, Theorem 8.15, applied to the gauge invariant of the one-handle filling.

9.1. The tensor supplied by the folded triangle. We use the one-handle geometry and notation of Section 2.5. A presentation p of M has torus compression bodies U_0, U_1 of genus $h \geq 2$, the compression body $C = C_{h,g}$ with $g = h + k$ adds k handles, $V_i = C \circ U_i$, and $M^+ = M \#^k (S^1 \times S^2)$. The correspondence L_C , the height functions ζ_i , the height Lagrangians R_i at height t , the lower Lagrangians $\tilde{P}_i$, the product $\mathbb{X}$ with $A_i = \tilde{P}_i^- \times R_i$, and the incidence conditions on (A_0, L_C, A_1) are as defined there.

The upper data may carry generic corrections γ from the span of the family F_* of (I5); these may move the generators, whose labels and gradings are continued (Paper I, Section 11.6). They are taken in a closed ball $|\gamma| \leq \delta(t)$ in that span. On this ball the corrected Lagrangians remain height Lagrangians at height t , and the restriction map of the family of perturbed flat connections parametrized by the primary perturbations remains a submersion [DFL21, Section 2.3]. Hence for $k = 0$ the incidence conditions hold for generic corrections, as in the proof of Paper I, Proposition 11.10; for $k \geq 1$ the upper data are fixed first, and the incidence conditions, for the pair and for its reflection, hold for generic lower data chosen afterwards (item (4) of Section 13.1; Paper I, Section 10.1). The resulting upper presentation is the upper height presentation of Definition 9.1.

Fix the height t . Choose an open C^1 -neighbourhood $\mathcal{N}_C(t)$ of the pair $(L(U_0), L(U_1))$, together with the corresponding neighbourhood $\mathcal{N}_{\tilde{C}}(t)$ of the reflected pair, for the reflected compression, with the following property: for every lower pair in $\mathcal{N}_C(t)$ and every correction in the ball, the two auxiliary height intersections remain transverse, and their critical points continue uniquely from the product Morse model. Such a neighbourhood exists by openness of transversality for these compact intersections and compactness of the ball of corrections. For $k = 0$ we shrink it so that it lies in the neighbourhood $\mathcal{N}(t)$ of (I5) at the same height. For $k \geq 1$ the upper Lagrangians R_i are fixed, and the lower primary data are in addition required to lie in the neighbourhood $\mathcal{N}_{C,R}$ of (I14)(i). The lower data are small presentations (Section 13.1) whose norm is small relative to t , so that their Lagrangians, and those along the segment joining them to zero in the small ball, lie in $\mathcal{N}_C(t)$. In both cases the grading, action and coefficient statements of (I13) and (I14) hold for the lower data: for $k \geq 1$ by (I14)(i),

and for $k = 0$ by (I14)(ii) applied along the path that first corrects the height datum and then follows that segment.

The two auxiliary complexes are $CF_{\mathbb{X}}(A_0, L_C)$ and $CF_{\mathbb{X}}(L_C, A_1)$. Their distinguished top cycles are q_0^* and q_1^* , respectively. Top degree refers to the auxiliary integer grading m of (I13), whose differences satisfy $m(a) - m(x) = \text{ind}(u) - 8\kappa_{\text{aux}}(u)$. Here κ_{aux} is the relative class of Paper I's auxiliary pair. Let $\mathcal{A}_i$ denote its critical action: the values at the intersection points of the local exact-graph primitive of the perturbed pair, with the sign fixed by the ordered pair. Writing $\omega_{\mathbb{X}}$ for the symplectic form of $\mathbb{X}$, the relation is

$$4\pi^2 \kappa_{\text{aux}}(u) = \int u^* \omega_{\mathbb{X}} - (\mathcal{A}_i(a) - \mathcal{A}_i(x)).$$

At the flat height data, $\mathcal{A}_i$ is the height function $t\zeta_i$ with the sign of the ordered pair. Under small changes of the primary data, $\mathcal{A}_i$ is the primitive of the changed pair, not the function $t\zeta_i$. Thus the symplectic area includes an endpoint correction; its quotient by $4\pi^2$ is the total topological energy. Attaching a positive generator sphere increases κ_{aux} by $1/2$ and the index by 4, leaving $m(a) - m(x)$ unchanged. The normalization of m agrees with the Morse grading. A top class is defined by m , not by the parity grading. Depending on the order of the pair, the corresponding critical point is, at uncorrected height data, the maximum or minimum of the height function, and in general its continuation.

Denote the upper height presentation by $q = C \circ p$; this notation includes its height perturbation. With a split almost-complex structure on $\mathbb{X}$, folding identifies the output complex with $CF(\bar{p}) \otimes^r CF(q)$. Let $\mathcal{K}^r$ be the closed identification (8.6) of Section 8, with its transported orientation lines. It uses the orientation system transported independently at $\bar{p}$; the reflection pairing with $CF(p)^\vee$ is constructed in Section 11. The triangle product for the ordered triple (A_0, L_C, A_1) , including the conversion Θ at its closed gauge cut Z in the position of Convention 8.1, is denoted by $\mu_{\mathbb{X}}^\Theta$. It gives the cycle

$$(9.1) \quad z_C = \mathcal{K}^r \mu_{\mathbb{X}}^\Theta(q_0^* \otimes^r q_1^*) \in CF(\bar{p}) \otimes^r CF(q).$$

Definition 9.1. Fix a height t , the ball of corrections $|\gamma| \leq \delta(t)$ and the neighbourhood $\mathcal{N}_C(t)$ above.

- (a) The *upper height presentation* at height t has the compression bodies V_0, V_1 , the Lagrangians R_0, R_1 , possibly corrected by some γ in the ball, and its height and auxiliary regular data (Paper I, Lemma 11.3).
- (b) Let p be lower data as above (small, with norm so small relative to t that the Lagrangians along the segment from zero lie in $\mathcal{N}_C(t)$, and, for $k \geq 1$, primary data in $\mathcal{N}_{C,R}$) which satisfies the incidence conditions of Section 2.5 with the upper height presentation $q = C \circ p$. The *unpaired compression cup* is

$$(9.2) \quad \check{U}_C = \text{bend}_{CF(\bar{p})}(z_C) : CF(\bar{p})^\vee \longrightarrow CF(q),$$

where bend is the bending map of Section 3 and z_C is the cycle (9.1).

- (c) Let $k = 0$, so that L_C is the diagonal, and write r for the upper height presentation. A *small fold pair* (p, r) at height t consists of a small lower presentation p whose norm is so small relative to t that the Lagrangians along the segment from zero lie in $\mathcal{N}_C(t)$, and an upper height presentation r at height t whose primary data lie in the ball $\mathfrak{B}$ of Section 10. The pair is subject to the incidence conditions and the auxiliary regularity conditions above (see also Section 13.1). For such a pair put $B(p, r) = \check{U}_C$.

Thus a fold pair is a particular pair of nearby presentations with a fixed diagonal triangle and its two height caps. Unfolded along the seam, the folded triangle is a quilted triangle

whose seam carries the correspondence L_C . Paper I, Remark 11.6, relates the mod-two cup to the quilted invariant of L_C ; that comparison is not used here.

Let $Z_0 = D_0 \cup (-C)$ and $Z_1 = C \cup (-D_1)$ be the gauge inputs of the auxiliary pairs (A_0, L_C) and (L_C, A_1) , each diffeomorphic to $V_i \cup (-V_i)$; let N_0, N_1 be their integral mixed maps, and let $\theta_i \in C_G(Z_i)$ be the relative invariant of the auxiliary filling $F_i = V_i \times [0, 1]$ of Z_i (Paper I, Lemma 11.1(i) and Lemma 11.2). The top cycles are normalized by $q_i^* = \pm N_i \theta_i$. The sign is one sign of each distinguished cycle. This is an integral normalization, for the following reason. The auxiliary gauge complexes $C_G(Z_i)$ are the framed complexes of $\#^r(S^1 \times S^2)$, with 2^r generators, and $I^\#(\#^r(S^1 \times S^2))$ is free of rank 2^r [Sca15, Section 7.8] (Paper I, proof of Proposition 9.4), so their differentials vanish. Along a fine subdivision of the path of end data from the height data, the gauge continuation maps between them are therefore isomorphisms, and they are triangular with respect to m by the energy bound in the proof of Paper I, Lemma 11.7(ii), which uses no signs; a triangular isomorphism has a triangular inverse, so each maps the generator of maximal m to $\pm$ itself (the argument of Paper I, Lemma 9.2, over $\mathbb{Z}$). At the height data the relative invariant of the filling is $\pm$ that generator, there being one regular solution (Paper I, Proposition 9.4; Lemma 7.7), and composition at the connected cut Z_i together with the data-change homotopy of Proposition 4.4, applied to the filling as in the proof of Lemma 7.7, transport it by these maps, the continuation cylinder being the product cobordism with varying Floer data; the differentials vanishing, the homotopies are equalities on chains. Hence θ_i is $\pm$ the generator of maximal m . The unit-diagonal calculation in Proposition 6.3, whose proof applies to the auxiliary pairs at their end data (the entry from the top generator to q_i^* has $\kappa_{\text{aux}} = 0$, so its solutions are the small deformations of the constant pair at the continued q_i^*), gives the coefficient ± 1 on q_i^* , and the other coefficients vanish by the grading argument of Paper I, Lemma 9.2 and Proposition 9.4, with (I13) and (I14) for the folded pairs, which use no signs.

9.2. The relative invariant of the one-handle filling. Let B_k be the gauge filling whose closed boundary is the three-manifold $Z = D_0 \cup (-D_1)$ of Section 2.5, obtained from the folded triangle and its two auxiliary fillings. Its relative invariant is a homogeneous cycle $\psi_{B_k} \in C_G(Z)$. Choose on this filling the almost-complex structure of Section D.1, which satisfies (D.1), and use its restrictions on the pieces. Let Θ' be the right-placed conversion at Z from the restriction of this structure to the restriction of $J_{\bar{\mathcal{E}}}$ (Section D.1), the structure in which (8.17) is stated, and write $\mathbf{E}\psi_{B_k}$ for $\mathbf{E}\Theta'\psi_{B_k}$. Since a composite of conversions agrees with the direct conversion up to one sign (Section 4), $N_Z\Theta^r\psi_{B_k}$, with the conversion of Lemma 9.2, agrees up to one sign with $N_Z\Theta^r\Theta'\psi_{B_k}$, with the conversion of (8.17). After excision the resulting relative invariant is a cycle on $C_G(-M) \otimes^r C_G(M^+)$. It is the relative invariant of the bent cobordism $W_k^\#$: by Paper I, Lemma 11.1(ii), $B_k \cup_Z \mathcal{E} \cong W_k^\#$ with the incoming end read as outgoing and with lifting group ϕ_{W_k} . Gluing at the connected cut Z has no extra gluing fibre (Lemma D.1), and the orbit count (D.2) gives the integral coefficients.

Reading the incoming end of the one-handle cobordism W_k as an outgoing end, that is, bending it, gives the identity

$$(9.3) \quad (\Phi_{G,M} \otimes^r 1)[\mathbf{E}\psi_{B_k}] = \pm[\text{ten}(I_J^\#(W_k))].$$

To check the sign convention, let ς be the relative invariant of the product cylinder with both ends bent outgoing. Gluing the doubly incoming cylinder of Section 4 to it along $-M$ gives the product cylinder. In ordered determinant gluing the line of $x \in C_G(M)$ passes the $-M$ factor of $\varsigma = \sum \varsigma_1 \otimes \varsigma_2$, so the glued map is $x \mapsto \sum (-1)^{|x||\varsigma_1|} \langle \varsigma_1, x \rangle_M \varsigma_2$, and it is ± 1 . By the definition of $\Phi_{G,M}$, $\langle y, x \rangle_M = \text{ev}'(\Phi_{G,M} y \otimes x)$, where $\text{ev}'(f \otimes x) = (-1)^{|f|} f(x)$

is the closed evaluation for (3.3). The same composite formed with ev' and $\text{ten}(1)$ is the identity, so $(\Phi_{G,M} \otimes 1)\varsigma = \pm \text{ten}(1)$, with one sign. Glue W_k to its second end and use the ordered determinant gluing. This gives the tensor of the right-placed map. An odd endpoint conversion is already part of that map through (4.3); moving it to the other side would change the tensor by a grading operator. The residual signs in (9.3) are therefore scalar. The same argument holds for $k = 0$, when the cobordism map is gauge continuation.

9.3. The cup formula. Define the transposed reflected mixed map

$$(9.4) \quad n_p = \text{ev}_{C_G(M)}^{-1} (N_{\bar{p}} \Phi_{G,M}^{-1})^\vee : CF(\bar{p})^\vee \longrightarrow C_G(M).$$

It is an even chain isomorphism when the reference almost-complex structures of M and $-M$ are paired and the gauge end data at $\bar{p}$ are the reflections of those at p , as for the bent product cylinder defining $\Phi_{G,M}$. The remaining regularizing data of the mixed domain defining $N_{\bar{p}}$ are chosen independently, as in Paper I, Section 11.2. Other gauge end data at $\bar{p}$ change $N_{\bar{p}}$ by the gauge continuation of Proposition 6.4(i), and n_p is then an even chain homotopy equivalence. The definition uses only gauge duality and the mixed map at the reflected presentation.

The integral version of the mod-two folded triangle comparison (I4) is the folded form of the triangle comparison (7.10).

Lemma 9.2 (Folded triangle comparison). *Let N_0, N_1 be the integral mixed maps of the product-target pairs (A_0, L_C) and (L_C, A_1) , with gauge inputs $Z_0 = D_0 \cup (-C)$ and $Z_1 = C \cup (-D_1)$, and let $\mu_{\mathscr{W}} : C_G(Z_0) \otimes^r C_G(Z_1) \rightarrow C_G(Z)$ be the integral gauge map of the cobordism formed by the pieces D_0, C, D_1 . With the conversion at Z placed as in Convention 8.1,*

$$\mu_{\mathbb{X}}^\Theta(N_0 \otimes^r N_1) \simeq \pm N_Z \Theta^r \mu_{\mathscr{W}}.$$

Proof. The unsigned argument is the proof of Paper I, Theorem 11.8: the family of data of Paper I, Lemma 9.12, collapses the middle compression body C , which is an odd compression body with Lagrangian L_C , stretches the two ordinary necks cutting off the domains defining the mixed maps of (A_0, L_C) and (L_C, A_1) , and stretches the closed cut Z . We orient every configuration in the family by the several-input construction of Section 5, with the conversion line at Z in its fixed position. At the collapsed end, Theorem B identifies the orientations of the index-zero solutions with the orientations of the symplectic triangle glued to the mixed domains of the two auxiliary pairs. It applies to this collapse: Section 7 rests on the collapse analysis of Paper I, Sections 6–8, and the proof of Paper I, Theorem 11.8, supplies its inputs for the collapse of $C = C_{h,g}$, with product target, both boundary surfaces matched and the lower sheet reversed, using the two-trace slice space $H^1(\Sigma_h; \text{ad}) \oplus H^1(\Sigma_g; \text{ad})$. The sign comparison of Proposition 7.5 uses only canonical orientations along paths of invertible operators, so the reversed sheet does not enter it. Ordinary gluing along the two stretched necks (Lemma 3.3) gives the count $\mu_{\mathbb{X}}^\Theta(N_{\mathcal{V}_0} \otimes^r N_{\mathcal{V}_1})$, with the input lines in cyclic order, where $\mathcal{V}_0, \mathcal{V}_1$ are the corner pieces of the zero-scale domain, the one-input mixed pieces for (A_0, L_C) and (L_C, A_1) attached at its incoming ends. Likewise the mixed map obtained below at the stretched cut is the count $N_{\mathcal{V}_Z}$ on the piece $\mathcal{V}_Z$ of DFL type for (A_0, A_1) beyond Z . By Lemma 7.6, the strip actions of these three pieces are those of the target complexes of N_0, N_1 and N_Z , and the three counts differ from N_0, N_1 and N_Z by boundaries δH with H odd. Along the family, the output strip action is defined through the invertible terminal constant pairs on the output regions (Lemma A.2). It is constant by the argument of Corollary 8.7; this is used at the breaks below. Since $\mu_{\mathbb{X}}^\Theta$ and $\Theta^r \mu_{\mathscr{W}}$ are closed, the identity for tensors of homotopies in Section 7.3 and (3.2) reduce the lemma to the same statement for $N_{\mathcal{V}_0}, N_{\mathcal{V}_1}$ and

$N_{\mathcal{V}_Z}$. In the rest of the proof N_0 , N_1 and N_Z denote these three counts. At the stretched cut, gluing at the connected closed cut Z has no extra gluing fibre (Lemma D.1). The stretched cut decomposes the domain into the triangle cobordism of D_0 , C , D_1 , glued at Z to the piece $\mathcal{V}_Z$. Lemma 8.3 is stated for W_λ^J , but its proof uses only a decomposition at a closed cut, the localized shift, which Proposition 2.2(b) supplies on folded triangles, and Convention 8.1. The same proof gives parts (i)–(iii) for this decomposition, and part (iii) identifies the orientation with that of the gauge cobordism followed by the conversion and the mixed map; this gives $N_Z \Theta^r \mu_{\mathcal{W}}$. Parts (i) and (ii) give the compatibility at the input and output breaks. The family between the two ends is a compact one-parameter family with compactly supported parameter derivative. Its boundary strata are the two endpoints and ordinary breaks at the two gauge inputs, ordered by Lemma 3.2 as in the proof of Lemma 8.12, and at the output strip; the collapse-scale stage is a family of the same kind. Lemma 8.12, with $f = |\Theta|$, turns its compactified index-zero space into the stated homotopy. The overall sign reflects the fixed orientation choices at the two ends. $\square$

Proposition 9.3 (Cup formula). *Let p be lower data and $q = C \circ p$ the upper height presentation at a height t , as in Definition 9.1(b), and let $W_k : M \rightarrow M^+$ be the cobordism of the k one-handles. With n_p as in (9.4),*

$$(9.5) \quad \check{U}_C \simeq \pm N_q I_J^\#(W_k) n_p.$$

In particular, for a small fold pair (p, r) (Definition 9.1),

$$(9.6) \quad B(p, r) \simeq \pm N_r \text{cont}_G(p, r) n_p.$$

Proof. Apply Lemma 9.2 to the two distinguished input cycles $q_i^* = \pm N_i \theta_i$. Since the gauge filling B_k is the triangle cobordism glued to the two auxiliary fillings, ψ_{B_k} is homologous to $\pm \mu_{\mathcal{W}}(\theta_0 \otimes \theta_1)$ by gauge composition. Hence the folded triangle output is, up to a boundary and one sign, $N_Z \Theta^r \psi_{B_k}$ with the conversion of Lemma 9.2, hence $N_Z \Theta^r \Theta' \psi_{B_k}$ with the conversion of (8.17). Apply (8.17) to the homogeneous cycle $\Theta' \psi_{B_k}$. With the convention for $E\psi_{B_k}$ stated above, the result is

$$[z_C] = \pm [(N_{\bar{p}} \otimes^r N_q) E\psi_{B_k}].$$

Equation (9.3) identifies this with the tensor of $N_q I_J^\#(W_k) n_p$. The identification follows directly from (3.6) and the graded double-dual map in (9.4). Since the difference of the tensors is a boundary, their images under $\text{bend}_{CF(\bar{p})}$ differ by δK . This proves (9.5). The same calculation with $k = 0$ gives (9.6). $\square$

This proves a formula for the cup tensor without selecting the reflection pairing needed for the source of Paper I’s map. The next sections construct that pairing and show that it gives the ordinary symplectic continuation when $k = 0$.

10. SPHERE CHARACTERS UNDER SMALL HOLONOMY PERTURBATIONS

The sphere character of a coherent strip orientation system is the sign by which the orientation changes when a sphere in M_h is attached to a strip (Definition 10.1). It decides whether coherent continuation orientations with prescribed ends exist (Lemma 12.1). Theorem 10.6 shows that it is constant on the ball $\mathfrak{B}$ of primary perturbations at an object with nonzero integral gauge homology. Nonvanishing is used only to find an intersection point that persists through each birth–death singularity of a generic path (Lemma 10.5); the character is read off at that point.

10.1. The character of an orientation system. For a transverse pair $P = (P_0, P_1)$ let $\mathcal{P}(P_0, P_1)$ be the path space of paths in $\mathbf{M}_h$ starting on P_0 and ending on P_1 . For $h \geq 2$, the compression-body Lagrangians are simply connected and have vanishing second homotopy group (Paper I, Section 3.2). The relative homotopy classes of strip configurations with fixed ends are consequently a torsor for $\pi_2(\mathbf{M}_h) \cong \mathbb{Z}$ in the lift convention of [DFL21, Section 2]. Fix its positive generator S_h ; attaching it changes the index by four. For $h = 1$ the moduli space is a point and the sphere-character assertions below are empty.

Definition 10.1. Let p be a presentation whose Lagrangians $P = (P_0, P_1)$ intersect transversely (its regularizing data need not be regular), and let Γ^p be the transported coherent orientation system on its strip configurations (Section 6). The determinant line of a sphere operator, together with the inverse determinant line of the tangent space of $\mathbf{M}_h$ at its node, carries its complex orientation. Attach a sphere to a strip configuration at an interior node. The orientation given by Γ^p on the glued configuration differs from the node-glued orientation, formed from Γ^p and this complex orientation, by a sign. The sign depends only on the class of the sphere, and it defines the *sphere character*

$$(10.1) \quad \chi^p : \pi_2(\mathbf{M}_h) \longrightarrow \{1, -1\}.$$

Associativity of node gluing makes χ^p a homomorphism to $\{\pm 1\}$ (a character). Multiplying generator orientation lines by signs does not change it: the two endpoint factors occur in both sides of the comparison. Multiplying them by their grading signs has the same property. Thus the character is information about the orientations of the determinant lines over all strip configurations, beyond the signs of the differential. The relation of χ^p to the orientation systems induced by relative spin structures in quilted Floer theory [WW15a] is not studied here; only the constancy of χ^p and its invariance under reflection are used.

At a transverse point x , a second description is available. Attach a fixed representative sphere to the constant mixed pair C_x and compare its transported mixed orientation with the product of the canonical orientation of C_x and the complex node orientation. Denote the ratio, the *pointed character*, by $\chi(\mathfrak{h}, x; d)$, where $\mathfrak{h}$ is the primary perturbation and d contains the regularizing and almost-complex data.

Lemma 10.2. *If the Lagrangians of the presentation p with data $(\mathfrak{h}, d)$ intersect transversely, then $\chi(\mathfrak{h}, x; d) = \chi^p(S_h)$ for every intersection point x . The ratio is constant along any compact smooth family of primary data carrying a persistent transverse intersection point. It is also independent of small regularizing data at fixed primary data and fixed reference almost-complex structure.*

Proof. Glue a strip to the constant pair and attach the sphere away from the ordinary gluing collar; moving the sphere from the constant pair to the strip keeps the node in the interior of a pure strip, where Proposition 2.1(c) applies. The two possible orders of these operations induce the same determinant-line isomorphism. The sphere and node factors are complex, hence even as real graded lines, so commuting them with the strip factor gives no sign. Cancelling the strip orientation gives equality of the pointed and strip definitions. A strip configuration exists between any two intersection points, in some relative homotopy class; it need not solve the holomorphic equation. Applying the same argument to that configuration shows independence of x .

For a family with a persistent point, the constant-pair operators and the node-gluing map vary continuously by Proposition 2.1, and the transported orientation of Proposition 5.4 by Proposition 2.2. The ratio is a continuous function to $\{1, -1\}$ and is therefore constant. To vary regularizing data, use the segment to zero data. The perturbations are supported off

the strip end, and the constant-pair operators stay invertible. Transporting their canonical orientations along the segment proves the last assertion. Compatible almost-complex structures can be varied through the same argument, with the reference almost-complex structure held fixed. $\square$

The persistent point need only remain transverse at that point. Other intersections of the two Lagrangians may become degenerate. The lemma is about determinant lines over configurations and therefore does not require a regular moduli space along the family.

10.2. A common ball of primary perturbations. Let $\mathcal{C}_i$ be the finite real span of the sign-sensitive holonomy functions on C_i° , modulo constants. Use the norm

$$(10.2) \quad \|\mathfrak{h}_i\|_i = \sup_B \|\nabla_B \mathfrak{h}_i\|_{L^\infty} + \sup_{B \neq B'} \frac{\|\nabla_B \mathfrak{h}_i - \nabla_{B'} \mathfrak{h}_i\|_{L^2}}{\|B - B'\|_{L^2}},$$

$$\|\mathfrak{h}\| = \max(\|\mathfrak{h}_0\|_0, \|\mathfrak{h}_1\|_1).$$

The suprema are taken over smooth connections in a fixed determinant bundle. Cylinder functions have finite norm by the holonomy estimates of [DFL21, Proposition 6.1]. If the norm vanishes, the gradient vanishes on the affine connection space, so the function is constant; thus this is a norm modulo constants.

Choose a finite trace family F_0 whose restrictions to the compact flat-connection moduli spaces give embedding coordinates. In the model $L(C_i) \cong \mathrm{SU}(2)^{h-1}$, quaternion coordinates can be obtained from sign-sensitive traces using the fixed holonomies on the odd torus. Such traces distinguish central sign changes; passing to unsigned $\mathrm{SO}(3)$ traces would lose that property. For a finite family F containing F_0 , put

$$\mathfrak{B}(F) = \{\mathfrak{h} \in \mathrm{span} F : \|\mathfrak{h}\| < r_*\}, \quad \mathfrak{B} = \bigcup_F \mathfrak{B}(F).$$

The radius is independent of the additional functions in F , for the fixed splitting and metric data. These sets are nested under inclusion of finite families and are convex. They differ from the sets $\mathcal{U}(F)$ of Paper I, Definition 11.12. The small ball of Section 13.1 is defined in the same way with a smaller radius r_χ and with F subject to the conditions of Paper I, Definition 3.3, chosen so that its elements are also admissible in the sense of Paper I.

Proposition 10.3. *Fix the splitting and the metric data. The radius $r_* > 0$ in the definition of $\mathfrak{B}$ can be chosen, independently of the finite family $F \supset F_0$, so that every primary perturbation $\mathfrak{h}$ with $\|\mathfrak{h}\| < r_*$ has the following properties.*

- (i) *The perturbed flat moduli space of each compression body is a smooth compact manifold identified with its unperturbed moduli space; its boundary restriction is an embedded Lagrangian close to the unperturbed one.*
- (ii) *The connection representatives and their first derivatives in the moduli parameters are close in L_1^2 and L^2 , respectively, uniformly as $\|\mathfrak{h}\| \rightarrow 0$.*
- (iii) *The functions from F_0 remain embedding coordinates on these perturbed moduli spaces.*
- (iv) *If the two boundary Lagrangians intersect transversely, the corresponding closed gauge critical points are nondegenerate and are exactly the lifts of those intersections.*

Proof. The norm gives $\|F_B\|_{L^\infty} \leq \|\mathfrak{h}_i\|_i$ for an $\mathfrak{h}_i$ -flat connection. If a sequence with norm tending to zero escaped every gauge neighbourhood of the flat moduli space, Uhlenbeck compactness on the compact compression body would give a subsequence weakly converging in $W^{1,4}$, after gauge, to a flat connection. The $W^{1,4}$ embedding in C^0 is compact in three

dimensions. A Coulomb–Neumann slice relative to the flat limit and the elliptic estimate for (d_B^*, d_B) then improve this to convergence in L_1^2 , a contradiction. The boundary slice and compactness statements are the versions for a manifold with boundary in [Weh04].

In a slice at a flat connection split into the finite-dimensional harmonic space and its complement. The linearized perturbed-flat operator on the complement has a bounded right inverse. The second term of (10.2) makes the holonomy part a small Lipschitz perturbation in L^2 , so the contraction gives a graph over the harmonic space. Differentiating its fixed-point equation in the finite-dimensional parameter gives the asserted first-derivative control. Hodge decomposition and boundary elliptic regularity identify these solutions with the smooth perturbed-flat solutions; the exact regularity inputs are recorded in Appendix A. Compactness of the flat moduli space makes the thresholds uniform over finitely many slice charts.

For a trace function τ , its bounded Lipschitz gradient extends the identity

$$\tau(B') - \tau(B) = \int_0^1 \langle \nabla_{B+s(B'-B)} \tau, B' - B \rangle ds$$

to L^2 connections. Differentiating on the graph family shows that the trace-coordinate map remains C^1 -close to its unperturbed embedding. Openness of embeddings of a compact manifold proves (iii). The matching exact sequence for the two compression-body boundary problems identifies the kernel of the closed Hessian with the tangent intersection of the Lagrangians, using the restriction and deformation sequences of [DFL21, Propositions 2.10 and 2.16], as in [DFL21, Lemmas 2.45 and 2.47]; hence transversality gives (iv). $\square$

The norm controls exactly the two estimates used in this proposition. On a fixed finite-dimensional perturbation family all higher smooth seminorms are finite and are handled separately in compact-family regularity statements.

10.3. Paths with only birth–death singularities. Enlarge F to include the monomials of degree at most five in the trace coordinates F_0 . These remain cylinder functions: a monomial is a smooth invariant function of the joint finite holonomy family.

Lemma 10.4. *Two regular primary data in $\mathfrak{B}$ can be joined within a common $\mathfrak{B}(F)$ by a smooth path with finitely many singular parameters. At each singular parameter there is one quadratic tangency, with versal normal form*

$$(10.3) \quad S_t(x, y) = x^3 + \sigma(t - t_*)x + Q(y), \quad \sigma \in \{1, -1\},$$

where Q is nondegenerate. All other intersections are transverse.

Proof. Start with the segment from $\mathfrak{h}_a$ to $\mathfrak{h}_b$ and perturb it to

$$\mathfrak{h}_t = (1 - t)\mathfrak{h}_a + t\mathfrak{h}_b + \beta(t)c,$$

where β vanishes near the endpoints and is positive on the interior where singularities may occur. The segment is compact in the open convex ball, so small c keeps the path in the ball.

Variation of the holonomy function induces an exact Lagrangian deformation; its Hamiltonian on the perturbed Lagrangian is the restriction of that function to the connection moduli space, by the Lagrangian cobordism of [DFL21, Proposition 2.18]. The trace coordinates remain embedding coordinates by Proposition 10.3. Polynomials of degree at most three in embedding coordinates realize arbitrary third jets at one point. Polynomials of degree at most five realize arbitrary second jets at two distinct points. For the latter assertion, choose an affine function ℓ with $\ell(a) = 0$ and $\ell(b) = 1$. Given a second jet at b , multiply ℓ^3 by the

quadratic Taylor polynomial at b of that jet divided by ℓ^3 . The product has degree at most five, realizes the prescribed jet at b , and vanishes to order three at a . Use $(1 - \ell)^3$ in the same way for the jet at a , then add the two polynomials. Restriction of ambient jets to an embedded submanifold is surjective.

In a Weinstein chart, varying one Lagrangian changes the relative generating function by this Hamiltonian to first order. The map on its jets is triangular with identity diagonal. Parametric transversality therefore excludes higher singularities, two simultaneous tangencies, and tangencies with zero crossing speed: these strata have codimension at least two. A generic one-parameter path meets only the codimension-one birth–death stratum, transversely. Its parameter-dependent splitting normal form is (10.3).

The singular-parameter set is closed by compactness of the Lagrangians. Each singular parameter is isolated by the normal form, and the endpoint neighbourhoods are regular. It is consequently finite. No uniform separation of the two points in the two-point jet argument is required; transversality is applied on the open configuration space of distinct pairs. $\square$

Lemma 10.5. *Suppose $H(C_G(M)) \neq 0$. The path in Lemma 10.4 has a finite subdivision with regular subdivision points such that each subinterval carries a continuous transverse intersection point throughout the subinterval.*

Proof. Every regular primary datum has at least one intersection. Otherwise its gauge complex would have no generators by Proposition 10.3(iv), contradicting invariance and nonvanishing of the gauge homology.

At a birth–death time the local pair in (10.3) is absent on one side. If it were the only intersection at the singular time, compactness would force the nearby regular intersection set on that side to be empty. Thus another intersection exists at that time and is transverse. The implicit function theorem continues it across the wall.

On a closed interval with no walls, the incidence set of intersections is a compact one-manifold projecting locally diffeomorphically to the interval. It is a finite covering, hence a disjoint union of sections. Extend the chosen transverse point from each wall to neighbouring regular subdivision times. This gives the required finite collection of persistent sections. $\square$

Theorem 10.6 (Constancy of the sphere character). *Fix a splitting, metric data and a reference almost-complex structure, and suppose $H(C_G(M)) \neq 0$. Then the sphere character $\chi^p(S_h)$ takes the same value at every regular presentation p with these data whose primary data lie in $\mathfrak{B}$.*

Proof. First remove the regularizing data at both endpoints using Lemma 10.2. Join the primary data by the subdivided path of Lemma 10.5. On each subinterval the pointed character is constant along its persistent transverse point. At a regular subdivision time the pointed character equals the global character for any choice of point, so one may switch to the persistent point for the next subinterval. Chaining these equalities gives constancy along the whole path. Restore the endpoint regularizing data. The same argument allows variation of the compatible almost-complex structure on M_h , with the reference almost-complex structure of the object fixed. $\square$

The argument uses nonvanishing only to find a point persisting through each birth–death. It does not assume that individual generators persist along an arbitrary path, or that the regular locus of the perturbation ball is connected.

11. REFLECTION OF THE TRANSPORTED ORIENTATIONS

This section constructs the reflection pairing $\Phi_p : CF(\bar{p}) \rightarrow CF(p)^\vee$ (Definition 11.4 and Proposition 11.5). Its existence rests on Corollary 11.3: the transported systems of p and $\bar{p}$ have the same sphere character. The proof compares the characters across the diagonal fold (Proposition 11.2), moving a sphere through the unfolded triangle at zero collapse scale.

11.1. Reflection of strip orientation systems. The reflected presentation $\bar{p}$ interchanges the two compression bodies. As in Paper I, Section 11.2, its ordinary strip data are the reflections of those of p , namely $J(-s, 1-t)$ and $-H(-s, 1-t)$ with the full one-form transformation, so that reflection is a bijection between the strip moduli spaces of p and $\bar{p}$. This convention is used in Proposition 11.5, which asserts a chain isomorphism. A strip is reflected by

$$\bar{u}(s, t) = u(-s, 1-t).$$

This map preserves the orientation of the strip domain and keeps the surface almost-complex structure. It reverses the translation direction. The target symplectic manifold retains the original orientation; the reversed order of the Lagrangians supplies the dual complex. The formula reflects strips only; the transported system of $\bar{p}$ is obtained through the mixed domain of $\bar{p}$.

For a strip from x to y , the reflection of its coherent orientation has differential sign

$$(11.1) \quad \text{sign}_{\text{dual}}(\bar{u}) = -(-1)^{|y|} \text{sign}(u).$$

The minus sign is the reversed translation vector; the other sign is the graded transpose. In general the *reflection transpose* from p to $\bar{p}$ assigns to the reflection of a strip of index n on p the sign $(-1)^{n|w|}$, relative to the identification of determinant lines induced by the reflection, where w is the target of the strip on p . Thus this reflected system is the one for which the paired generator bases give the graded dual differential. The reflection transpose from $\bar{p}$ to p is the same construction with p and $\bar{p}$ exchanged; its sign is read on the target of the strip on $\bar{p}$. The two reflection transposes are not inverse to each other: on a strip of index n their composite is $(-1)^n$, the graded evaluation sign in Proposition 11.5. The inverse of the reflection transpose from p to $\bar{p}$, whose sign is again read on the target of the strip on p , is called the *unfolding transpose*; it occurs in Lemma 11.1.

Attaching S_h changes the index by four. The determinant factor at a sphere node is complex and even. The additional transpose sign is $(-1)^{4|y|} = 1$. Consequently reflection of a coherent strip system preserves its sphere character. What remains is to compare this reflected system with the system transported through the reflected mixed domain.

11.2. Capped folded triangles: strip gluing and unfolding. We record the end-gluing comparison before using it to compare sphere characters. For a capped folded configuration c with outputs x on $\bar{p}$ and y on r , write F_c for its determinant line after contracting the two normalized auxiliary cap factors. The conversion line at Z is retained in its fixed position. Transport from the mixed configuration gives $\Gamma_c : F_c \rightarrow L_{\bar{p}}(x) \otimes^r L_r(y)$. Thus this is a tensor with two outgoing ends. Bending the first end gives a map from $L_{\bar{p}}(x)^\vee$, not from $L_{\bar{p}}(x)$.

Lemma 11.1. *Let u be any smooth outgoing strip from x to x' on $\bar{p}$, and w any smooth outgoing strip from y to y' on r . Capping commutes with their determinant gluing. For $k = 0$, unfolding along the diagonal commutes with the same gluing, with the first outgoing strip reflected into an incoming strip by the unfolding transpose. These statements hold for arbitrary strip indices and configuration classes.*

Proof. Put $d_u = \det D_u$ and $d_w = \det D_w$. The assertion about the end systems is the commutative diagram

$$\begin{array}{ccc} F_c \otimes d_u \otimes d_w & \xrightarrow{(\Gamma_u \otimes^r \Gamma_w) \Gamma_c} & L_{\bar{p}}(x') \otimes^r L_r(y') \\ \downarrow G & & \parallel \\ F_{c\#(u,w)} & \xrightarrow{\Gamma_{c\#(u,w)}} & L_{\bar{p}}(x') \otimes^r L_r(y'). \end{array}$$

In the upper arrow the strip determinant factors are inserted next to their corresponding generator factors using the right tensor convention. In particular moving d_u past $L_r(y)$ contributes $(-1)^{\text{ind}(u)|y|}$.

Here G is ordered determinant gluing, including the cap contractions. Before contraction take the order

$$\mathcal{J}, \quad \det \Theta, \quad \det D_c, \quad \det D_{a_0}, \quad \det D_{a_1}, \quad d_u, \quad d_w,$$

where $\mathcal{J}$ is the complete generator input line in cyclic triangle order and a_0, a_1 are the two caps. If a sphere is attached, append its complex node factor last. Both routes use this order. Associativity regroups factors without interchanging them; any permutation needed for an end convention is performed on both routes with its graded sign. In particular the conversion line is never carried through the moving configuration determinant. The unit cap maps contract their paired generator and determinant factors by even evaluation maps. No assumption that the individual cap generator has even parity is needed.

First consider the completed mixed triangle at positive scale. Its localized shift retains the ordinary two-input cobordism and ends at that cobordism followed by a constant mixed pair. There ordinary determinant associativity gives the end-gluing diagram, with the conversion line and the fixed unit cap maps in the order above. The relative family construction transports it over the positive-scale configuration families.

Pull back the entire ordered diagram through Lemma 7.1. Its ordinary-gluing compatibility applies to the strip and cap operators at their ends in the cofinal neck-and-scale regime defined in the statement of Lemma 7.1. Only after this diagram commutes do we cancel the lines of the corner pieces and contract the two cap factors. The fixed-width closed-cut line comparison retains the conversion factor; its completion is expressed through the separate output lines through the reference comparison $K_{0,x}$ of Proposition 8.8, not by completing the fixed closed- Z input itself. This proves the first diagram for the limiting configurations. The limiting triangle is reached by collapse, not by a shift to a constant triangle, which does not exist because the triple intersection of its boundary labels is empty.

For unfolding with $k = 0$, first choose doubling data retaining the Hodge almost-complex structure and zero inhomogeneous germ in the collapse neighbourhood, the compatible ordinary ends and the fixed cap models. Choose a folded configuration obtained from a smooth configuration on the unfolded strip. The two coefficient systems are then restrictions of one smooth system. At the $W^{1,p} \rightarrow L^p$ realization, $p > 2$, equal traces along the diagonal are exactly the condition for the two sections to join to a $W^{1,p}$ section after reflecting the lower side. Restriction is the inverse, and the same identification holds for L^p target sections. This is an operator identification. At higher Sobolev orders equal traces alone would not suffice. Instead identify the higher-order determinant lines through their common smooth kernels and cokernels, using smooth stabilization columns. Elliptic regularity makes these identifications commute with gluing and parameter variation.

The comparison is first made on this doubling locus. Any finite set of smooth configuration classes has representatives there: smooth the joined maps in a small collar of the diagonal and keep the end data fixed. The smoothing is a relative homotopy and changes no sphere class.

Interpolate the auxiliary almost-complex and perturbation data to doubling data within the admitted data space, preserving these germs and end values; Proposition 2.1(a) gives the continuous Fredholm families. Determinant transport extends the resulting comparison back to the original finite configuration families. The first gluing diagram and path independence in Proposition 5.4 make it independent of these homotopies. With stabilizations off the joining regions the operator identification on the doubling locus gives

$$\begin{array}{ccc} F_c \otimes d_u \otimes d_w & \xrightarrow{G} & F_{c\#(u,w)} \\ \downarrow U_{\text{ord}} & & \downarrow U \\ d_{Ru} \otimes \det D_{Uc} \otimes d_w & \xrightarrow{G} & \det D_{Ru\#Uc\#w}. \end{array}$$

Here Ru runs from x' to x on p , so it attaches to the incoming end of Uc , which runs from x to y . The left arrow U_{ord} applies unfolding and reflection and then interchanges F_c and d_u , with sign $(-1)^{|F_c|\text{ind}(u)}$. At the operator level the maps are restriction, reflection and the same ordinary-end cutoff patching, whose supports are disjoint from the diagonal. Their two composites agree. This is a diagram on the zero-scale section problem: unfolding is applied only to zero-scale operators, after collapse. The displayed graded interchange expresses that agreement in the chosen tensor order. Cancelling the ordered stabilizations proves the displayed statement. The diagonal normal factor has its complex orientation and is even. Reflection of the strip u of index n moves its index line across the paired endpoint line $L_{\bar{p}}(x)$ and contributes the sign $(-1)^{n|x|}$. This is the unfolding transpose defined after (11.1): the inverse of the reflection transpose from p to $\bar{p}$, applied to the strip Ru , which runs from x' to x on p , so that the sign is read on the target x of Ru on p . It is not the reflection transpose from $\bar{p}$ to p , whose sign $(-1)^{n|x'|}$ is read on the target x' of u on $\bar{p}$ and differs from $(-1)^{n|x|}$ by $(-1)^n$. At index four both signs equal 1, as for $(-1)^{4|y|}$ above. For a strip counted modulo translation the reversed translation vector adds the minus sign in (11.1). To check the total sign, put $a = |x|$, $b = |y|$ and $f = |F_c| = a + b$. The outgoing gluing sign is $(-1)^{nb}$. On bending, the reflected input sign $(-1)^{na}$ and the swap $(-1)^{nf}$ have product $(-1)^{n(a+f)} = (-1)^{nb}$. Thus the diagrams give the same sign at every index. $\square$

11.3. Sphere characters across the diagonal fold. Fix one small fold pair with lower primary datum $\tilde{p}$ and upper height datum r , both in the small ball. Choose the height after the radius of that ball, as in Section 13. Double the folded triangle and attach its auxiliary caps. Before bending, the resulting disk has two outgoing strip ends modelled on $\tilde{p}$ and r (Figure 2). By Lemma 11.1, its transported orientation is coherent with the independently transported strip system on $\tilde{p}$ and the transported outgoing system on r . The comparison with the dual system on p is made only after the sphere characters have been identified.

Proposition 11.2. *The sphere character of the reflected lower presentation equals that of the upper height presentation:*

$$\chi^{\tilde{p}}(S_h) = \chi^r(S_h).$$

Proof. Choose a smooth reference configuration c_0 on the doubled disk. Insert a representative sphere on the first outgoing strip end at a fixed small positive plumbing parameter. Move this smooth inserted configuration through the unfolded disk to the second outgoing end, then refold. The two resulting configurations lie in the same relative class. Under diagonal unfolding the two sphere classes are identified with the same generator of the relative sphere group. In the product description this moves the class between the two factors modulo the diagonal class. A reversal of its representative replaces the generator by its negative and does not change the sphere character, whose values are ± 1 .




FIGURE 2. The comparison of Proposition 11.2, for $k = 0$, where L_C is the diagonal. (a) A folded triangle in $\mathbb{X} = M_h^- \times M_h$ with boundary on (A_0, L_C, A_1) , $A_i = \tilde{P}_i^- \times R_i$, with the two height caps at its auxiliary corners and its output end at $A_0 \cap A_1$. (b) Doubling, capping and unfolding along the diagonal give a disc with two outgoing strip ends, modelled on $\tilde{p}$ and on r ; the image of the diagonal is the interior dashed line joining the two caps. At zero collapse scale a sphere representing the generator of $\pi_2(M_h)$, attached at an interior node near the first end, is moved through the disc to the second end. Node gluing and the transported orientation are continuous along the move, so the two sphere characters agree.

There is a continuous orientation on the determinant bundle along this path. Its construction on the folded domain uses the collapse transport of Section 7, the closed-cut conversion, and the reference comparison of Proposition 8.8. For the endpoint comparisons place the sphere in an ordinary strip factor attached away from the collapse neighbourhood. At the fixed positive plumbing parameter, Lemma 7.1 commutes with gluing that entire strip. Its interior-node determinant comparison is the usual punctured Cauchy–Riemann bordering and complex evaluation map; the gauge rows are absent on this pure strip. Ordinary-gluing associativity thus gives the endpoint comparison with the complex even node factor appended last. The collapse scale and ordinary neck can be chosen after this plumbing parameter, so the node and the collapse never degenerate together. At the two ends, the first diagram of Lemma 11.1 identifies the transported orientation with both specified strip-node gluing orientations. No Koszul sign is introduced by moving the complex node factor.

It remains to carry the node-glued orientation along the path, on which the node crosses the image of the diagonal. Carry out the move at zero collapse scale, on the doubling locus, with the doubling data of Lemma 11.1: the Hodge almost-complex structure and zero inhomogeneous germ in the collapse neighbourhood. There the folded operator is identified with the unfolded pure real Cauchy–Riemann operator, by the $W^{1,p} \rightarrow L^p$ identification in the proof of that lemma, and in the unfolded picture the image of the diagonal is an interior line, not a matching or boundary collar. Along the whole path the node lies in a disc separated from the matching and boundary collars, so Proposition 2.1(c) makes the node gluing a continuous isomorphism of line bundles along the path. Hence transport carries the node-glued orientation at the start to the node-glued orientation at the end. The transported orientation is continuous along the same zero-scale path by Lemma 7.1: the path is a compact smooth family of limiting configurations, and the collapse datum is unperturbed near the diagonal. The ratios at the two ends are therefore equal, giving the assertion.

This comparison uses arbitrary smooth configurations, not only holomorphic solutions. To justify coherence where the fold is attached to prescribed strip configurations, select one reference configuration for each of the finitely many generator pairs and their required sphere classes. With the height already fixed, choose the ordinary neck lengths and then the collapse scale so that ordinary gluing applies to this finite set, in the order of choices of Section 13. The determinant-line construction then extends over each configuration component by path independence. Proposition 2.1 and Proposition 2.2 are used on all the domains in this finite construction, not only at the constant solution. $\square$

Changing the reference almost-complex structure on the closed cut changes the output generator lines by a fixed diagonal conversion. That conversion cancels in a sphere ratio. Thus the character comparison is unchanged by the conversion required in Section 8, although that conversion remains essential for the cup identity as a chain-level formula.

Corollary 11.3 (Reflection invariance). *Let p be a small presentation of an object with nonzero integral gauge homology. Then p and its reflection have the same sphere character:*

$$\chi^{\bar{p}}(S_h) = \chi^p(S_h).$$

Proof. Apply Theorem 10.6 on $-M$ to compare $\bar{p}$ with $\tilde{\bar{p}}$. Apply Proposition 11.2 to compare the latter with r . Apply character constancy on M to compare r with p . The reflected ball is the same ball under the chosen reflection identification. Gauge duality makes nonvanishing for M and $-M$ equivalent. $\square$

11.4. The reflection pairing.

Definition 11.4. Let p be a small presentation of an object with nonzero integral gauge homology. On the strip configurations of $\bar{p}$ compare two orientation systems: the transported system of $\bar{p}$, and the dual system of (11.1), which is the reflection transpose from p to $\bar{p}$ of the transported system of p . Fix a reference intersection point x_0 . For an intersection point x , let $\kappa_p(x) \in \{1, -1\}$ be the ratio of the two systems on a strip configuration of $\bar{p}$ from x_0 to x . The *reflection pairing* is

$$(11.2) \quad \Phi_p : CF(\bar{p}) \longrightarrow CF(p)^\vee, \quad o_x^{\bar{p}} \longmapsto \kappa_p(x)(o_x^p)^\vee.$$

Thus Φ_p is the integral form, for transported orientations, of the standard duality $CF(L_1, L_0) \cong CF(L_0, L_1)^\vee$ given by reversing strips, here with $(L_0, L_1) = (P_0, P_1)$.

The ratio of the two systems is constant on smooth configuration components and multiplicative under concatenation. Changing the configuration from x_0 to x changes its class by a sphere; Corollary 11.3 says that the ratio on this sphere is one. Hence $\kappa_p(x)$ is well defined, and the ratio on a strip from x to y is $\kappa_p(x)\kappa_p(y)$, independent of the strip class.

Proposition 11.5. *The map Φ_p of (11.2) is an even chain isomorphism. The comparison of the orientations of the determinant lines over all strip configurations determines it up to one sign. The pairings may be chosen so that*

$$(11.3) \quad \Phi_{\bar{p}} = \Phi_p^\vee \text{ ev}.$$

Proof. Equation (11.1) and the ratio just computed give exactly the dual differential under (11.2). The paired reference almost-complex structures make the sum of the two endpoint parities zero, so the map is even. Any two comparisons have constant ratio: configurations exist between every pair of intersection points, and the ratio on all those configurations fixes the relative generator signs. Only the initial reference sign is free.

Apply to a strip of index n with endpoints x, y the reflection transpose from p to $\bar{p}$ and then the one from $\bar{p}$ to p , which defines $\kappa_{\bar{p}}$. Their signs are read on the two different endpoints and multiply to $(-1)^{n(|x|+|y|)} = (-1)^n$. The corresponding change of the paired basis is the sign in the graded evaluation map. Thus the pairing obtained by the second reflection is exactly $\Phi_p^\vee \text{ev}$, after the common reference sign is fixed. This proves (11.3) for the orientations over all strip configurations. $\square$

A chain-map condition alone would not determine κ_p up to one sign. If the differential splits into two subcomplexes, one can change the sign on one summand and still obtain a chain map. The preceding construction uses all strip configurations, including those which are not solutions, to exclude that extra freedom.

Put

$$(11.4) \quad \widehat{\Phi}_p = \Phi_p^\vee \text{ev} : CF(p) \longrightarrow CF(\bar{p})^\vee.$$

This is the reflection pairing which will be composed with the cup map (9.2).

12. CONTINUATION AND THE DIAGONAL COMPRESSION CUP

12.1. Coherent orientations with prescribed ends. A continuation datum from p to q has a strip-like surface with Lagrangian boundary conditions varying on a compact interval and equal to those of p, q at its ends. The almost-complex and Hamiltonian data are likewise constant outside that interval. Throughout this section the primary paths lie in the common convex ball of Proposition 10.3; after Hamiltonian identification, their exact-graph primitives lie in a fixed convex neighbourhood. Homotopies have fixed endpoints and remain in this neighbourhood. Almost-complex and Hamiltonian continuation data are varied with these fixed boundary identifications. All statements of independence below concern this class of data. For generators x, y , let $\mathcal{B}_{p,q}(x, y)$ be its smooth configuration space, with the prescribed exponential limits. A coherent orientation system assigns relative orientations to the determinant lines over these spaces and commutes with attaching strips at either end.

The space of configurations in a fixed relative homotopy class is path connected. One may first choose a continuous relative homotopy, smooth it relative to the fixed ends and labels, and then make it constant beyond common end cutoffs. Local charts are given by an exponential map for a metric making the boundary labels totally geodesic. The topology in Section 2 makes these paths continuous. The relative classes are a torsor for the same sphere group as fixed-label strips, and the index changes by four under its generator. In particular there is at most one index-zero component for fixed endpoints.

Lemma 12.1 (Rigidity of coherent systems). *Fix a continuation datum from p to q in the class described at the start of this subsection. Two coherent continuation orientation systems with the same end systems differ by one sign. Such a system exists if and only if the sphere characters at the ends agree.*

Proof. The ratio of two systems is constant on each configuration component. If a strip is attached at either end, the common end orientation cancels, so the ratio does not change. Choose one reference configuration from x_0 to y_0 . For any x, y , choose strips from x to x_0 and from y_0 to y . Attaching these strips transports the ratio to a configuration with ends x, y . Any difference in relative class can be supplied by a sphere class on either strip. Since the end systems are the same, this also leaves the ratio unchanged. Thus there is only one sign for all configurations.

For existence choose a reference orientation and extend it by the same gluing procedure. The determinant bundle on each configuration component is orientable: the boundary Lagrangians are the parallelizable compression- body Lagrangians [WW15a, Theorem 5.1.6], and the remaining sphere factors have complex orientation. The only compatibility left when the gluing paths are changed is moving a sphere class from the incoming end to the outgoing end. The two extensions agree precisely when the two sphere characters agree. This also proves necessity by moving the node of one sphere through a continuation configuration. The construction is continuous because the determinant line is a line bundle over the smooth configuration space. $\square$

In an index-zero comparison a finite form of the lemma suffices. Choose one strip for each input generator and one for each output generator, in the classes needed to reach the index-zero component. Coherence on that finite collection fixes the ratios of all zero-dimensional counts. This form allows the neck lengths and collapse scales to be chosen after a finite set of configurations has been fixed, so that estimates are needed only on that finite set.

At an object with nonzero integral gauge homology, Theorem 10.6 supplies existence whenever p and q have the same splitting, metric data and reference almost-complex structure, and primary data in the ball $\mathfrak{B}$ of Section 10. This includes the auxiliary upper data of Section 13.1, which lie in $\mathfrak{B}$ but in general not in the small ball. For a lower datum and an auxiliary upper height datum, the folded configuration itself also supplies a coherent system and hence the required character equality; moving the node through the unfolded configuration crosses the image of the diagonal, and this is justified as in the proof of Proposition 11.2.

12.2. Parity and continuation maps. For a continuation configuration c from x to y , the *continuation parity*

$$\text{par}(p, q) = |y| - |x| - \text{ind } D_c \pmod{2}$$

is independent of the configuration and its end generators. It is additive under composition of continuation data. The configuration description by sphere classes, together with a homotopy in the specified small family, shows its independence of the chosen datum in that family. For the paired reference almost-complex structures of this paper, $\text{par}(p, q) = 0$ at the data used in the comparison, at an object with nonzero integral gauge homology. The proof uses the integral cup formula, so that only the nonvanishing of the integral homology enters.

Indeed, at a small fold pair (p, r) the map $B(p, r)$ is homogeneous and (9.6) identifies it up to homotopy with the even homotopy equivalence $N_r \text{cont}_G(p, r) n_p$. On nonzero integral homology this map cannot be both even and odd. Thus its degree is zero. The doubled configuration has the same index as the continuation problem, and the gauge reflection pairing is even; hence $\text{par}(p, r) = 0$. A common upper height datum and additivity give $\text{par}(p, q) = 0$. The integral cup formula used here is proved in Section 9 without symplectic continuation maps, which are introduced only below.

With a coherent system define

$$\text{cont}_S(p, q) : CF(p) \longrightarrow CF(q)$$

by the signed zero-dimensional continuation count. It is even and commutes with the differential. The compactification of a one-dimensional continuation moduli space proves this exactly as in Proposition 6.1.

Proposition 12.2. *Let p, q, r be regular presentations of an object with nonzero integral gauge homology, with the same splitting and metric data and with primary data in the ball $\mathfrak{B}$*

of Section 10, and define their continuation maps as above. These maps satisfy, up to sign and chain homotopy,

$$(12.1) \quad \text{cont}_S(q, r) \text{cont}_S(p, q) \simeq \pm \text{cont}_S(p, r),$$

$$(12.2) \quad \text{cont}_S(p, p) \simeq \pm 1.$$

Their homotopy classes are independent of the continuation datum. They are homotopy equivalences.

Proof. For a long strip neck, ordinary gluing identifies the zero-dimensional solutions with pairs of continuation solutions. The determinant gluing is coherent, so the signed matrix is the composite matrix. The glued orientation system has the same end systems as the chosen direct continuation system. Lemma 12.1 compares them by one sign. Varying the neck length gives the chain homotopy in (12.1).

A homotopy of continuation data, constant near the ends, gives the corresponding homotopy by its oriented one-dimensional boundary. The pure symplectic compactness and gluing used here are Proposition 2.1(c). For constant data the zero-dimensional solutions are constant strips, with canonical positive orientation. This normalizes the unit. Apply composition to a datum and its reverse to obtain a homotopy inverse. $\square$

12.3. Continuation maps under reflection. There is no translation quotient for a continuation configuration. Reflecting it therefore has the reflection transpose sign of Section 11 but no additional translation sign. Using the pairings of Proposition 11.5 gives

$$(12.3) \quad \text{cont}_S(\bar{p}, \bar{q}) \simeq \pm \Phi_q^{-1} \text{cont}_S(q, p)^\vee \Phi_p.$$

To verify coherence of the reflected system, glue an input strip, a continuation configuration, and an output strip, and reflect the three pieces in reverse order. If their indices are n_u, n_c, n_w , the reordering sign is $(-1)^{n_u n_c + n_u n_w + n_c n_w}$. Combining it with the three transpose signs leaves $(-1)^{n_u \text{par}(p, q)} = 1$. The generator corrections κ_p and κ_q identify the reflected end systems with the transported systems of $\bar{p}$ and $\bar{q}$. Rigidity then gives the one sign in (12.3). This argument also shows why parity zero is needed here.

12.4. The diagonal compression cup equals continuation. For $k = 0$ the compression correspondence is the diagonal. Double the folded triangle, glue the two distinguished caps, and unfold along that diagonal. The resulting surface is a continuation strip. This is the geometric identification used in Paper I. To obtain an integral statement we compare its coherent systems before counting.

Proposition 12.3 (Diagonal cup). *At every small fold pair (p, r) ,*

$$(12.4) \quad B(p, r) \hat{\Phi}_p \simeq \pm \text{cont}_S(p, r).$$

Proof. Use both diagrams of Lemma 11.1. After unfolding, the incoming system is the image of the transported system of $\bar{p}$ under the unfolding transpose, the inverse of the reflection transpose from p to $\bar{p}$ (see the paragraph after (11.1)). Its comparison with the system on p is the ratio of orientation systems over all strip configurations that defines κ_p (Definition 11.4). Indeed, by the definition of κ_p , on every strip of $\bar{p}$ between x and y the transported system of $\bar{p}$ is $\kappa_p(x) \kappa_p(y)$ times the reflection transpose of the transported system of p . Applying the inverse transpose, on every strip of p from x to y the incoming system is $\kappa_p(x) \kappa_p(y)$ times the transported system of p . These are precisely the two endpoint corrections in $\hat{\Phi}_p$. The reflection transpose from $\bar{p}$ to p , which enters the double reflection in the proof of Proposition 11.5, would instead give the ratio $(-1)^{|x|+|y|} \kappa_p(x) \kappa_p(y)$, and the comparison below would

leave the grading operator $\text{gr}_{CF(p)}$, which anticommutes with the differential. On a homogeneous generator the coefficient of $\widehat{\Phi}_p$ is $(-1)^{|x|}\kappa_p(x)$, since the graded evaluation map enters (11.4). Bending by $\text{bend}_{CF(\bar{p})}$ supplies the second factor $(-1)^{|x|}$; the two evaluation signs cancel. Thus precomposition by $\widehat{\Phi}_p$ changes the unfolded input system into the prescribed system on p . The outgoing system is already the prescribed one on r . This holds for every configuration class and every strip index, as rigidity requires.

The pulled-back capped system and the chosen continuation system therefore have the same two end systems. Their relative sign is constant under the two gluing operations in the diagrams. Rigidity (Lemma 12.1) applies and leaves one scalar sign. A nonconstant generator-sign automorphism would change one of those gluing maps of determinant lines over strip configurations; a grading operator would change an odd-index gluing map. Both have been fixed before invoking rigidity. This proves the required comparison of coherent systems.

For sufficiently long cap necks, the gluing theorem gives a bijection of zero-dimensional solutions and preserves these orientations. A cap of index zero has one regular solution with unit coefficient, so no extra multiplicity occurs. The counts are therefore the same up to that sign. Varying the cap length and the triangle data to the ones defining $B(p, r)$ gives a chain homotopy through the one-dimensional moduli spaces. The preglue-to-solution path is within the pure section chart, and determinant transport on that path preserves the canonical orientation of a regular solution. This is the pure symplectic gluing input of Proposition 2.1(c). $\square$

This completes the proof of Theorem D. The cup in that theorem is the geometric compression cup of Paper I, and it is compared directly with the geometric continuation map.

12.5. Objects with vanishing homology. No object with vanishing integral gauge homology is known (see the remark after Theorem A); the convention below makes the statements unconditional. Suppose $H(C_G(M)) = 0$. Lemma 3.1 and Proposition 6.3 make both $C_G(M, p)$ and $CF(M, p)$ contractible for every presentation. Every closed map from or to a contractible complex is a boundary: compose it with a contracting homotopy and use (3.2). Hence there is exactly one morphism between it and any other object in the homotopy category.

At such an object we take the continuation and reflection maps needed for the comparison to be zero representatives of those unique morphisms. In particular a reflection map here is a homotopy equivalence between zero objects of $\mathbf{K}(\mathbb{Z})$. Character constancy and the reflection pairing of Definition 11.4 are used only at objects with nonzero integral gauge homology. Every identity involving an acyclic object is interpreted in the homotopy category.

This convention also handles a word passing through an acyclic intermediate object: its class is zero, as is the corresponding gauge composite. Reduction modulo two preserves a contracting homotopy. Thus Paper I's maps incident to that object are null-homotopic as well.

Nonzero mod-two gauge homology implies nonzero integral gauge homology, since an acyclic finite free integral complex would be contractible. Integral framed calculations place rational homology spheres and $\#^k(S^1 \times S^2)$ in the nonvanishing case: for a rational homology sphere Y the Euler characteristic of $I^\#(Y)$ is $|H_1(Y; \mathbb{Z})|$ [Sca15, Corollary 1.4], and $I^\#(\#^k(S^1 \times S^2))$ is computed in [Sca15, Section 7.8]. Only this implication is used. The converse fails in general: integral homology that is finite of odd order gives zero homology modulo two.

13. SMALL DATA, NATURALITY OF THE MIXED MAPS AND HANDLE SQUARES

13.1. Order of the auxiliary choices and small data. This subsection constructs the radius r_χ of the small ball of Definition 2.4, and fixes the order in which the remaining auxiliary data are chosen. Items (1)–(3) below are made once for each Heegaard splitting with its metric data, together with its reflection, before any comparison is considered. They are transported by isomorphisms of decorated data: for such an isomorphism ϕ , the choices for $\phi(\Sigma)$, with the transported metric data and reference almost-complex structure, are the transports of those for Σ . Since the norm (10.2) is invariant under bundle automorphisms, frame changes preserve them. Items (4)–(8) concern the finitely many presentations of one comparison. The order avoids requiring a perturbation to be simultaneously small relative to a parameter chosen after it.

- (1) Fix the trace-coordinate family F_0 of Section 10, and let the primary regularity threshold r_H be the radius r_* of Proposition 10.3.
- (2) Choose the auxiliary height t_* below the height thresholds of Section 9 and of the admissible neighbourhoods (15) of Paper I. Choose it so small that its height perturbations, together with every generic correction in the ball $|\gamma| \leq \delta(t_*)$ of Section 9, have norm less than r_H . These data are used as upper auxiliary presentations. This t_* is taken as the height t_* fixed in Paper I, Definition 11.12, together with its companion height $t'_* \leq t_*$ chosen there.
- (3) Choose the radius $r_\chi < r_H$ with the following property, for the splitting and for its reflection. If a primary datum lies in the span of a finite family containing F_0 and has norm less than r_χ , then its Lagrangians lie in the lower-data neighbourhood $\mathcal{N}_C(t_*)$ of Section 9 for the product compression body, and in the neighbourhoods $\mathcal{N}(t_*)$ and $\mathcal{N}(t'_*)$ that define admissible primary data in Paper I, Definition 11.12 (see (15)). Such a radius exists because these neighbourhoods are C^1 -neighbourhoods of the unperturbed pair of Lagrangians, and by Proposition 10.3(ii),(iii) and its proof the perturbed Lagrangians converge in C^1 , uniformly in the finite family, as the norm tends to zero. A small presentation is then admissible in the sense of Paper I: the segment from zero to its primary datum lies in the convex small ball, so along it the Lagrangians, and those of the reflection, stay in these neighbourhoods, and the datum lies in the path component $\mathcal{U}(F)$ of (15).
- (4) Choose smaller heights t'' for the proof of Proposition 13.2, t_f for the reflection fold and t_C for the compression cup and, for a Heegaard stabilization or destabilization, the two-handle heights t_2, t'_2 and the generic perturbations R'_0, R''_0 of R_0 used for stabilization in Section 14 (Paper I, proof of Proposition 12.4). Choose them so that the resulting upper and intermediate presentations are small in the sense defined below. Choose the lower datum of a fold pair or of a cup after its height t'', t_f or t_C , with norm so small relative to that height that the Lagrangians along the segment from zero to it lie in $\mathcal{N}_C$ at that height.
- (5) Fix the ordinary gauge trajectory data and the symplectic strip-end data for every required end, including the intermediate closed cuts. Make the simultaneous regular end choices before selecting their weights. For the auxiliary pairs, for every k , choose the end perturbations and the secondary data with $4 \sup \|f\|_\infty + 2 \sup \|\eta\|_{L^2}^2 < 2\pi^2$, as (I14) requires; the secondary data chosen in items (6) and (8) below respect this bound. Choose a common sufficiently small weight below the relevant primal, adjoint and scalar thresholds. For configuration reference paths it must also lie below their

- remainder rates; relative end preparation then supplies faster reference paths when needed. The completed-domain estimates at these choices are Proposition 2.1(b).
- (6) Select the finite compact reference families needed for coherence, and regularize, with these fixed ends, the limiting problems on the mixed domains defining the maps N_p , on the triangles and on the completed cores. The transport of their primary data, their finite stabilizations and their compact secondary supports are fixed before taking a geometric limit. In particular an auxiliary upper datum is a full presentation with its own thresholds for the mixed map. Keep the collapsing compression unperturbed, including the dependence regions of the holonomy terms.
 - (7) Choose the ordinary necks after these limiting data have been fixed. Then choose the collapse scales from the bounds for the compact reference families and the finite limiting index-zero solution sets, with their decay rates and the already fixed weights. For simultaneous determinant comparisons use the cofinal neck-and-scale regime in Lemma 7.1. Each threshold therefore depends only on parameters already chosen.
 - (8) At the resulting positive geometric parameters choose any further compactly supported secondary regularizers within the openness thresholds for the endpoint counts, preserving every end equation fixed in item (5). Extend them relatively over the comparison families. If an end datum or a limiting regularization is changed, reselect the dependent necks and collapse scales and use the corresponding data homotopies. Independence over the full permitted weight range is applied only after this common small-weight construction, as in Proposition 5.3.

Every choice has a positive threshold once the preceding finite data have been fixed.

With the radius r_χ of item (3), the small ball and the small presentations are those of Definition 2.4. The slices of the small ball by the finite families are convex, and two of its elements lie in a common slice. It lies in the ball $\mathfrak{B}$ of Section 10, and its elements are admissible primary data in the sense of Paper I (see (I5)). Small fold pairs are those of Definition 9.1(c): the lower presentation is small, with norm so small relative to t that the Lagrangians along the segment from zero lie in $\mathcal{N}_C(t)$, and the primary data of the upper height presentation lie in $\mathfrak{B}$. For $k \geq 1$, *compatible handle data* consist of a small lower presentation p and the upper height presentation $q = C \circ p$ at the height t_C , subject to the same conditions. The Lagrangians of p , and those along the segment from zero, lie in $\mathcal{N}_C(t_C)$, and its primary data, and those of its reflection, lie in the neighbourhoods $\mathcal{N}_{C,R}$ and $\mathcal{N}_{\bar{C},\bar{R}}$ of (I14)(i), on which the grading and coefficient statements of (I13) and (I14) hold for the pair and for its reflection, as the three-handle cap requires. By item (4), q is small. A *small decomposition* of a morphism, from (M_0, p_0) to (M_1, p_1) , is a word of elementary moves as in (I15): changes of presentation, frame changes, one- and three-handle attachments and two-handle cobordisms. All presentations in it are small, its maps S_e are those of Section 14, and its auxiliary data are chosen as in items (4)–(8).

The auxiliary upper data of item (2) lie in $\mathfrak{B}$ but in general not in the small ball. At an object with nonzero integral gauge homology, Theorem 10.6 applies to them, so their coherent continuation systems exist by Lemma 12.1. The upper data at the heights of item (4), and in particular the handle targets at which reflection or mixed duality (Proposition 13.2) is used, are small.

13.2. Naturality under small changes of holonomy perturbation. The mixed maps commute, up to sign and homotopy, with the gauge and symplectic continuation maps for changes of primary data within a small ball. All maps in this subsection and the next are even.

Proposition 13.1 (Primary naturality). *For small presentations p, q with the same splitting and metric data,*

$$(13.1) \quad N_q \operatorname{cont}_G(p, q) \simeq \pm \operatorname{cont}_S(p, q) N_p.$$

Proof. Choose one upper height presentation a of $-M$ at the height t_* of item (2) of Section 13.1, with its correction in the ball $|\gamma| \leq \delta(t_*)$ generic for $\bar{p}$ and $\bar{q}$. By item (3) there, $(\bar{p}, a)$ and $(\bar{q}, a)$ are small fold pairs. Apply the cup formula and Proposition 12.3 to both fold pairs. They give

$$\begin{aligned} \operatorname{cont}_S(\bar{p}, a) &\simeq \pm N_a \operatorname{cont}_G(\bar{p}, a) n_{\bar{p}} \widehat{\Phi}_{\bar{p}}, \\ \operatorname{cont}_S(\bar{q}, a) &\simeq \pm N_a \operatorname{cont}_G(\bar{q}, a) n_{\bar{q}} \widehat{\Phi}_{\bar{q}}. \end{aligned}$$

All factors being eliminated are homotopy equivalences. Cancel the common upper datum using continuation composition. Thus

$$\begin{aligned} \operatorname{cont}_S(\bar{p}, \bar{q}) &\simeq \pm (n_{\bar{q}} \widehat{\Phi}_{\bar{q}})^{-1} \operatorname{cont}_G(\bar{p}, \bar{q}) \\ &\quad \cdot n_{\bar{p}} \widehat{\Phi}_{\bar{p}}. \end{aligned}$$

Use the reflection formula (12.3) on the symplectic side and its gauge analogue, Lemma 4.5 for gauge continuation, together with (4.5). The symmetric conventions give

$$\widehat{\Phi}_{\bar{p}} = \Phi_p, \quad \Phi_{G, M} n_{\bar{p}} = N_p^\vee,$$

and the same equalities at q . Cancelling these pairings transforms the displayed identity into

$$\operatorname{cont}_S(q, p)^\vee \simeq \pm (N_p \operatorname{cont}_G(q, p) N_q^{-1})^\vee.$$

Transpose and rename p, q to obtain (13.1). Every double-dual map occurs together with its inverse. Because the maps here are even, the cancellations introduce no grading-dependent sign. $\square$

This proof uses the same pairing in the diagonal-cup and reflection identities. Choosing an unrelated chain-level duality isomorphism would leave a diagonal automorphism after the cancellation.

13.3. The reflected mixed map inverts N_p . Following Paper I, write $N_p^\dagger$ for the transpose of the mixed map of the reflected presentation $\bar{p}$, read through the gauge and symplectic reflection pairings:

$$N_p^\dagger = n_p \widehat{\Phi}_p : CF(p) \longrightarrow C_G(M),$$

with n_p as in (9.4) and $\widehat{\Phi}_p$ as in (11.4). Modulo two it is the map $N_p^\vee$ of Paper I, Proposition 11.9. The next proposition shows that it is a homotopy inverse of N_p , up to sign.

Proposition 13.2 (Mixed duality). *For every small presentation p ,*

$$(13.2) \quad N_p N_p^\dagger \simeq \pm 1.$$

Proof. First choose p so that (p, r) is a small fold pair at the height t'' ; by item (4) of Section 13.1, r is small. The cup formula (9.6) and Proposition 12.3 give

$$\operatorname{cont}_S(p, r) \simeq \pm N_r \operatorname{cont}_G(p, r) n_p \widehat{\Phi}_p.$$

By primary naturality the right side is homotopic, up to sign, to $\operatorname{cont}_S(p, r) N_p N_p^\dagger$. Cancel the continuation equivalence.

For a general small presentation q with the same splitting and metric data, transpose primary naturality for the reflected presentations. The symmetric pairings give

$$N_q^\dagger \text{cont}_S(p, q) \simeq \pm \text{cont}_G(p, q) N_p^\dagger.$$

Combine this with primary naturality and (13.2) at p :

$$N_q N_q^\dagger \text{cont}_S(p, q) \simeq \pm \text{cont}_S(p, q).$$

Cancel the continuation once more. The contractible case has the interpretation of Section 12.5. $\square$

13.4. One- and three-handle squares. At compatible handle data (Section 13.1) define the integral compression cup by

$$(13.3) \quad U_C = \check{U}_C \hat{\Phi}_p : CF(p) \longrightarrow CF(q), \quad q = C \circ p.$$

Define the cap by reflected transposition. If $\bar{C}$ is the reflected compression and $\check{U}_{\bar{C}} : CF(p)^\vee \rightarrow CF(\bar{q})$ is its unpaired cup, put

$$(13.4) \quad \check{V}_C = \text{ev}_{CF(p)}^{-1}(\check{U}_{\bar{C}})^\vee, \quad V_C = \check{V}_C \hat{\Phi}_q.$$

Then $V_C : CF(q) \rightarrow CF(p)$. Equivalently it is the transpose of the paired reflected cup with the two pairings inserted. These are the orientations of the compression cup and cap counts of Paper I. At other presentations they are conjugated by the geometric continuation maps.

Theorem 13.3 (One- and three-handle squares). *Let $C = C_{h,g}$ be the compression body adding k handles, $W_k : M \rightarrow M^+$ the corresponding one-handle cobordism, and $W_3 : M^+ \rightarrow M$ the three-handle cobordism whose reversal $\bar{W}_3 : -M \rightarrow -M^+$ is the one-handle cobordism of the reflected compression $\bar{C}$. Let (p, q) be compatible handle data, with U_C and V_C as in (13.3) and (13.4). Then*

$$(13.5) \quad U_C N_p \simeq \pm N_q I_J^\#(W_k),$$

$$(13.6) \quad V_C N_q \simeq \pm N_p I_J^\#(W_3).$$

The identities hold at compatible handle data and, with the cup and cap conjugated by continuations, at all small endpoint presentations with the splittings and metric data of the handle data; other endpoints are reached through Section 14.

Proof. The cup formula (Proposition 9.3) and mixed duality (Proposition 13.2) give

$$U_C N_p \simeq \pm N_q I_J^\#(W_k) n_p \hat{\Phi}_p N_p \simeq \pm N_q I_J^\#(W_k),$$

which proves (13.5). For the cap, transpose the cup formula at the reflected compression. The graded evaluation maps and $\Phi_{G,-M} = \Phi_{G,M}^\vee \text{ev}$ give the gauge transpose of the reversed handle map (Lemma 4.5 applied to $\bar{W}_3$). Mixed duality at q cancels the reflected mixed map. In detail, in the homotopy category,

$$\begin{aligned} V_C N_q &= \check{V}_C \hat{\Phi}_q N_q \\ &\simeq \pm \check{V}_C n_q^{-1} \\ &\simeq \pm N_p \Phi_{G,-M}^{-1} I_J^\#(\bar{W}_3)^\vee \Phi_{G,-M^+} \\ &\simeq \pm N_p I_J^\#(W_3). \end{aligned}$$

The only potentially odd map in the transpose comparison is the handle map; all pairings and mixed maps are even. The last step is Lemma 4.5 for $\overline{W}_3 : -M \rightarrow -M^+$, whose reversal is W_3 ; it holds with one sign although the handle map may be odd.

Finally conjugate both handle maps by continuations from arbitrary small endpoints to the compatible handle data. Primary naturality moves the mixed maps past these continuations, and gauge functoriality identifies the gauge composite with the same cobordism. Homotopies compose by (3.2). $\square$

14. DESCENT TO COBORDISMS AND PROOF OF THE MAIN THEOREM

14.1. Squares for the remaining elementary moves. The geometric presentation moves are the ones in Paper I. We record how the integral orientations enter them. A change of meridian disk systems does not change the Lagrangians; its map is the identity. A change of regularizing data uses the family defining Proposition 6.3. A change of strip-end almost-complex structure uses Proposition 6.4(ii), with the outgoing continuation map of that proposition and the fixed transported end systems. For a change of the reference almost-complex structure, the gauge and symplectic generator orientation lines are converted by the same right-placed conversion maps. The constant pair used as reference has no further line, so its orientation comparison commutes with this conversion.

An isomorphism of decorated data pulls back the equation, its linearization and the shifting construction; the pullback of a shift is a shift, and the orientation maps do not depend on the shifting path (Proposition 5.4). Choose pulled-back stabilizations in the determinant construction. The kernel gluing maps then commute with the isomorphism, and the mixed square commutes exactly on chains at the pulled-back data. This applies to frame changes and diffeomorphisms preserving the torus decoration. The transported presentation is small, since the choices of Section 13.1 are natural under these isomorphisms. For a frame change it has the same splitting and metric data, and a subsequent continuation (Proposition 13.1) gives the square at the originally chosen data, up to sign and homotopy. For a diffeomorphism the transported reference almost-complex structure is compared with the given one by the conversion of the preceding paragraph, and the subsequent change of presentation is a composite of the moves treated in this section.

A Heegaard stabilization is realized by the cancelling one- and two-handle pair of Paper I. Let r be the stabilized presentation. The one-handle square ends at the presentation $C \circ p$, which carries height perturbations on both bodies, while the two-handle square holds from $(C \circ p)'$ to r' . Here $(C \circ p)'$ is $C \circ p$ with the Lagrangian R_0 of its first body replaced by a generic small perturbation R'_0 from the span of F_* , which may move its generators, and with zero primary perturbation on its second body; r' is r with R'_0 on its first body and the two-handle height perturbation at height t_2 on its second body (Paper I, proof of Proposition 12.4). As in Paper I, insert the continuations $\text{cont}_S(C \circ p, (C \circ p)')$ and $\text{cont}_S(r', r)$; their squares are Proposition 13.1, since $(C \circ p)'$ and r' are small by item (4) of Section 13.1. Compose (13.5), these two squares and (7.11). The decorated gauge cobordism is the identity, so its map is continuation up to sign and homotopy. The reverse stabilization is treated by the cancelling two- and three-handle pair. Passing through a common stabilization treats changes of metric data at a fixed splitting. An isotopy of the splitting fixed near the ball is treated as in Paper I, Proposition 12.4: transport by the isotopy, as in the preceding paragraph, followed by a change of metric data and a change of primary data. The order of choices in Section 13.1 allows the new presentations to satisfy the integral height and ball conditions. Every orientation comparison in this argument is one of the signed squares above.

Thus every elementary move $e : p \rightarrow q$ has a symplectic map S_e , constructed from the curves of Paper I, and a square

$$(14.1) \quad N_q G_e \simeq \pm S_e N_p.$$

Here G_e is the gauge map of the decorated cobordism of e ; for a change of presentation it is a gauge continuation map. Changes of presentation give even maps, and handle attachments may give odd ones.

14.2. Independence of the decomposition.

Lemma 14.1 (Projective descent). *Suppose that the presentations of each object are connected by elementary moves, and that every cobordism, with any small endpoint presentations, is represented by a word of elementary moves. Suppose that the gauge maps G_e compose, up to sign and homotopy, to the gauge map of the represented cobordism. Suppose N_p are isomorphisms in a homotopy category and the elementary maps satisfy (14.1). Then S_d depends only on the represented cobordism and the endpoint presentations, up to sign and homotopy.*

Proof. For a word $d = e_m \cdots e_1$, compose the squares to obtain

$$S_d N_{p_0} \simeq \pm N_{p_m} G_{e_m} \cdots G_{e_1}.$$

Every composition of homotopies is valid by (3.2). If d' represents the same cobordism, gauge functoriality identifies the two gauge words up to sign and homotopy. Cancel N_{p_0} to identify S_d and $S_{d'}$.

For words representing the identity cobordism this also gives the transitive system of comparison maps between presentations. The identity word is the identity map, and concatenation of words gives composition up to sign. Hence these systems descend to a projective functor and the classes of N_p give a projective natural isomorphism. $\square$

Proof of Theorem A. Proposition 6.3 gives the integral mixed isomorphisms. Equation (7.11), Theorem 13.3, Proposition 13.1 for changes of primary data within a small ball, and the presentation comparisons above supply all elementary squares. Generation and connectivity are (II5); the presentations in these words may be taken small, because the slices of the small balls are convex and contain regular data, two primary data in one small ball lie in a common slice, and the balls are transported by isomorphisms (Section 13.1). Proposition 4.4 and Lemma 14.1 therefore give (1.1) and the projective functoriality.

The mixed maps preserve the chosen object gradings. On homology, the comparison with the homogeneous gauge map therefore gives the stated parity of the symplectic map. The parity is a property of the homology class, whatever chain representative of a long word is chosen.

Modulo two all determinant orientations give the same counts. The conversions become the identity on the generator basis, the signs in the tensor and dual formulas disappear, and the pairing (11.2) becomes the paired-basis reflection pairing. Thus the geometric maps reduce to those of Paper I at the same data. At an acyclic object the zero representatives reduce to maps homotopic to those of Paper I by Section 12.5. $\square$

14.3. Further questions. The scope of Theorem A, including its relation to homology-oriented instanton maps, to the integral map of [DFL21] and to relative-spin orientations, is described in the Scope subsection of Section 1. One further point concerns the descent itself. The symplectic maps form a projective functor, and exact composition signs would require more than projectivizing the maps; the regrading discussed after (4.6) indicates the additional grading information involved.

APPENDIX A. ANALYTIC FOUNDATIONS

This appendix proves Proposition 2.1, the four estimates in the proof of Lemma 7.1 (Section A.5), Lemma 8.5 (Section A.9) and Proposition 5.3 (Section A.10). All references to Paper I mean its version of 19 October 2026 [Wue26a]. The unsigned collapse, torus excision and generation results of that manuscript are used as they are stated there. The exceptions are the proofs here which repeat an argument of Paper I for a wider class of operators, such as stabilized families of nonsolutions, and state which steps change; the principal case is Lemma 7.1. The arguments below combine the local theory of [DFL21, DFL24] with explicit end models. They are the local-plus-ends arguments of [DFL21, Lemma 4.34], carried out on the domains of this paper. Theorem, proposition and remark numbers in [DFL24] are those of arXiv:2109.07038v1.

A.1. Local theory and end models. *Local estimates.* On a compact region meeting the matching seam, with a product metric and the fixed almost-complex structure near the seam, the local estimates for the mixed operator are [DFL21, Theorem 4.17]; they are uniform along C^∞ -convergent sequences of configurations by [DFL21, Remark 4.19]. For domain-dependent almost-complex structures, and for the formal adjoint, the local estimates and weak regularity are [DFL24, Theorem 5], with [DFL21, Remark 4.33]; their uniformity on neighbourhoods of the configuration and of the data is [DFL24, Remarks 5.73 and 5.76] for the operator; applied to the adjoint $d/d\theta + D_{(B_\theta, S_\theta^*)}$, which has the same form [DFL24, Section 5, before Theorem 5.9], they give it for the adjoint. Away from the seam these are the interior and boundary estimates for the anti-self-duality operator with the normal boundary condition and for real Cauchy–Riemann operators with totally real boundary conditions. Holonomy perturbations contribute zeroth-order terms with the bounds of [DFL21, Proposition 6.1]. The conformal factor of Definition 5.1 does not change the anti-self-duality equation; it rescales the scalar row by f^{-2} and adds zeroth-order terms to it. A compact family in the topology of Section 2 has coefficients converging in every C^k norm on compact sets, so these estimates are uniform on it.

Charts. At a smooth configuration the matching-boundary chart is built from [DFL21, Lemmas 4.1, 4.2 and 4.4] and the construction following Lemma 4.4, with the slice of [DFL21, Proposition 4.6] and the L^q radius of Paper I, Remark 5.2, in place of the L^2 radius stated in [DFL21, Lemma 4.1]. The construction is local, so on a compact family the charts vary continuously with their centres. Charts with a radius uniform as a torus length tends to infinity are used only at the preglued references of Paper I, Definition 10.23 and Lemmas 10.28–10.29.

End models. (O) At an ordinary gauge or strip end the operator is translation invariant with invertible asymptotic operator, the extended Hessian at a nondegenerate irreducible critical point or $J\partial_t$ with transverse boundary conditions. Weighted estimates hold for weights below the gap, by [Don02, Section 3.3.1] for gauge ends and by the same Fourier argument for strip ends. (E) At a mixed end the stationary operator is self-adjoint with kernel the tangent space of the manifold of stationary solutions at that end, which is the perturbed Lagrangian of its compression body, and it has a spectral gap [DFL21, Section 4.2.1 and Proposition 4.21]. On the torus-extended slices Y_R , $R \in [0, \infty]$, the gap is uniform, and on the joined slices $D_{i,\lambda}$ it is uniform for $\lambda \geq \lambda_*$ (Paper I, Proposition 10.7). For every $\lambda > 0$ the slice $D_{i,\lambda}$ is compact and its stationary operator is self-adjoint, by the argument of that proposition, so its spectrum is discrete; its kernel is the tangent space of $P_i \times Q_i$ by the same kernel computation, which uses only flatness on the bridge and $H^*(T; \text{ad } \alpha) = 0$ and does not involve λ_* . Hence each $D_{i,\lambda}$ has a gap, uniform on compact ranges of λ by continuity, as in the

proof of Paper I, Proposition 10.17. The moving boundary domain is transported to a fixed one by [DFL24, Proposition 5.27 and Lemma 5.67], and the weighted estimate for weights $\delta \in (0, \gamma)$ follows by Fourier transform in the end coordinate: the stationary operator $\mathcal{A}$ is self-adjoint and its spectrum meets $(-\gamma, \gamma)$ only in 0, so $\|(\partial_\theta + \mathcal{A} - \delta)^{-1}\| \leq \min(\delta, \gamma - \delta)^{-1}$ (Paper I, Proposition 10.27). (T) On a flat torus region the operator and its adjoint are the invertible operators of Paper I, Proposition 10.3, with the stated weighted bounds; at a flat face the normal derivatives are recovered from the equation as in the proof of Paper I, Theorem 10.5(b).

A.2. Proof of Proposition 2.1(a).

Lemma A.1. *Let $\mathfrak{F}$ be a compact family as in Proposition 2.1. There are a compact subset K of the cores and a constant C such that, for every member and every x in the domain of its realized operator D and every η in the domain of its weak adjoint,*

$$\|x\| \leq C(\|Dx\| + \|x\|_{L^2(K)} + |p(x)|), \quad \|\eta\| \leq C(\|D^*\eta\| + \|\eta\|_{L^2(K)}),$$

where $p(x)$ denotes the stationary coordinates.

Proof. Choose cutoffs subordinate to a fixed compact core, the ordinary ends, the mixed ends and, on torus domains, the torus region beyond a fixed radius. Their gradients can be made small on the noncompact overlaps: there the cutoffs vary over a length ℓ , so their gradients are $O(1/\ell)$, and ℓ is chosen after the constants of the model estimates (Paper I, Proposition 10.27). On each end apply the corresponding model of Section A.1; the coefficients differ from the model by terms which are small beyond a common truncation, uniformly on $\mathfrak{F}$. On the core apply the local estimates. Summing and absorbing the small commutator and coefficient terms gives the first estimate for the decaying part; the finite stationary columns and the scalar row on them are bounded, which accounts for $|p(x)|$. The adjoint estimate is the same argument with the adjoint models, whose invertibility is recorded in Section A.1.

Varying geometry is handled as follows. The collar identifications ψ_R of Paper I, Definition 10.1(a), are the identity outside the region N_X , which is compact in the normal direction and extends along both mixed ends; they change the mixed-end slices from Y to Y_R . The width identifications in the proof of Paper I, Proposition 10.17, are the identity outside the bridge. Pulled back by them, the metrics form a compact family of uniformly elliptic metrics, and on the mixed ends model (E) is applied on the slices Y_R and $D_{i,\lambda}$ for lengths in the given compact range. At positive collapse scale, the domains X_ϵ , $\epsilon \in [\epsilon_1, \epsilon_2]$, of Paper I, Definition 6.2, are built in fixed coordinates, with section complex structure j_ϵ and seam embedding F_ϵ depending smoothly on ϵ . An isotopy of the collar embeddings F_ϵ , extended to a diffeomorphism of the section surface, identifies them with X_{ϵ_1} as decorated manifolds. At each fixed $\epsilon > 0$ the collapsing region is compact with a smooth metric (proof of Paper I, Proposition 6.9). The zero-scale limit domains are obtained, in the proofs of Paper I, Theorems 9.10 and 11.8, by joining DFL-type domains of compression pairs, with the new mixed ends of Paper I, Proposition 6.9, and a polygon domain along ordinary strip ends. The joined domain has a compact core, and the argument above applies to it directly, with model (E) at the new ends. On the completed torus domains a compact family is in torus normal form after the gauge transformations of Lemma 8.4(i), and model (T) applies beyond a fixed radius. $\square$

Proof of Proposition 2.1(a). Charts are recorded in Section A.1. By Lemma A.1 and Rellich's theorem on K , each realized operator has finite-dimensional kernel and closed range, and so does its adjoint; hence it is Fredholm with cokernel the kernel of the weak adjoint.

Local regularity for the operator and its weak adjoint makes these elements smooth. After transport to fixed spaces by the boundary maps of Section A.1 and the domain identifications above, the coefficients vary continuously in operator norm on the family, so the operators form continuous Fredholm families and their determinant lines form continuous line bundles. For the $W^{1,p} \rightarrow L^p$ realization, $p > 2$, of the Cauchy–Riemann summand used in unfolding in Section 11, the standard L^p theory of real Cauchy–Riemann operators with totally real boundary conditions and nondegenerate ends gives the Fredholm property, and elliptic regularity identifies its kernel and cokernel with those of the Hilbert realization. The comparison of realizations is Lemma 5.2; its paths are compact families of this kind, and its proof uses only the Fredholm and continuity statements just proved. The affine-centre comparison is Proposition 7.4. $\square$

A.3. Proof of Proposition 2.1(b). The end critical sets are finite and nondegenerate. The manifolds of stationary solutions at the mixed ends, the perturbed Lagrangians, are compact, and the stationary operators there have kernel exactly their tangent spaces. Hence there is a common gap for the finitely many end models of a compact family, and a common small weight below all gaps and scalar thresholds; at mixed ends the stationary coordinates are split off before the weight is imposed [DFL21, Definition 3.9(ii) and Section 4.2.1].

The scalar operator. The Neumann operator $d_A^* d_A$ is Fredholm by the argument of Lemma A.1: its end models have positive gaps by irreducibility of the end connections, and on the torus region the fibre operator has mass μ_T because $H^0(T; \text{ad } \alpha) = 0$. Integration by parts identifies its kernel at weight zero with parallel sections, which vanish on an irreducible collar; it is self-adjoint of index zero there, and sufficiently small weights preserve the index and add no kernel. This is the argument in the proof of Paper I, Theorem 5.6(ii). Near the preglued references the uniform version is the proof of Paper I, Lemma 10.26, which is stated for completed index-zero solutions and applies without change to the constant pairs and to the references along the shifting rectangles.

A regular zero-dimensional solution has an invertible realized operator by definition, and this does not depend on the realization by Lemma 5.2. Uniformity for the preglued families as the torus lengths tend to infinity is Lemma 8.5. At positive collapse scale the estimates are those of Lemma A.1 on compact scale intervals; as the scale tends to zero, the uniform statement used is Lemma 7.1. The constant reference pairs are treated in the next lemma. For small secondary or inhomogeneous data the operators remain invertible by openness, with the bounds of Lemma A.1.

Lemma A.2. *The constant reference pairs, with no secondary or inhomogeneous data, have invertible linearized operators on each of the following domains.*

- (a) *Every domain of DFL type in the class of (I1), including the domain defining the mixed map and the truncated output regions of the triangles.*
- (b) *The separate domains $W_{a,R}$, $R \in [0, \infty]$, and the joined domains W_λ^J , $\lambda > 0$.*
- (c) *The closed-cut piece M_λ of (I2), and every domain of the deformation from the proof of (I1) which joins it to the special quintuple of (D_0, D_1) (Paper I, Proposition 10.21).*

The inverse bounds are uniform on compact families of lengths, and the positive elements extend continuously through torus completion.

Proof. The quadratic identity. Write $d_A^{\mathfrak{h}}$ for the perturbed exterior derivative of [DFL21, (4.42)] for the primary perturbation $\mathfrak{h}$, which adds the primary Hessian term on the compression cylinders and equals d_A on the flat genus and torus regions. For every variation

(ζ, ν) of connection and section satisfying the linearized matching and arc conditions, with the stationary behaviour of the realization,

$$(A.1) \quad \int_X \operatorname{tr}(d_A^h \zeta \wedge d_A^h \zeta) + \int_U \Omega(d\nu, d\nu) = 0.$$

On the original pieces this is [DFL21, Lemma 4.40]: the interior terms are total derivatives by the symmetry of the primary Hessian and the primary equation, the matching terms cancel because the surface trace is closed with class the section value, and the arc terms vanish on the Lagrangian tangent spaces. A primary tangent vector satisfies the linearized primary equation, so its pure stationary contribution at a truncated mixed end vanishes; the terms involving the decaying remainder tend to zero with the truncation, and the ordinary ends contribute zero. The flat torus region contributes only Stokes terms. At a finite outer torus face the conditions are $\zeta(\partial_r) = 0$ and $d_\alpha(\zeta|_T) = 0$ (Paper I, Definition 10.23). Write the trace as $a + \psi d\theta$ with $d_\alpha a = 0$; since $H^1(T; \operatorname{ad} \alpha) = 0$ we have $a = d_\alpha f$, and $\int_T \operatorname{tr}(a \wedge \partial_\theta a) = \int_T \operatorname{tr}(d_\alpha f \wedge d_\alpha \partial_\theta f) = 0$ by Stokes on T . The cross term $\int_T \operatorname{tr}(a \wedge d_T \psi)$ vanishes by Stokes on T because $d_\alpha a = 0$, as in [DFL21, (4.52)]. On a completed domain use spatial cutoffs equal to one on all primary supports, whose gradients are of size C/R on the flat tail; the resulting errors tend to zero because the tail field is in H^1 . On a joined domain the bridge has no boundary, and its interfaces cancel those of the two cores. On the closed-cut piece the ordinary input Z contributes nothing.

The kernel. Let x be in the kernel and let (ζ, ν) be the corresponding variation of connection and section. The mixed equations give $(d_A^h \zeta)^+ = 0$ and the linear Cauchy–Riemann equation for ν , so both integrands in (A.1) are nonnegative with the trace convention of [DFL21]. Hence $d_A^h \zeta = 0$ and $d\nu = 0$. The section variation decays at its strip output, so $\nu = 0$ on each component. The section limit of the stationary coordinates is their image in the tangent space of the primary Lagrangian, and the restriction map is an embedding, so all stationary coordinates vanish. What remains is a decaying ζ with $d_A^h \zeta = 0$, $d_A^* \zeta = 0$ and the normal and face conditions.

A primitive. Construct η with $d_A \eta = \zeta$ as in [DFL21, Proposition 4.54]. On each compression cylinder take the temporal primitive from the gauge input; it converges because the kernel field solves the full gauge-fixed equation on the input cylinder, whose asymptotic operator is nondegenerate, so the field decays exponentially there. The primary equation and Hessian symmetry make it a primitive of the spatial part. On a flat genus product the surface class of ζ is $\nu = 0$, and the surface scalar inverse gives the primitive; the base components agree because a parallel adjoint scalar on an odd surface vanishes. On the flat torus region put $\eta_T = \Delta_\alpha^{-1} d_\alpha^* \zeta_T$ fibrewise. Since $H^1(T; \operatorname{ad} \alpha) = 0$ this is a primitive of the fibre part, and since $H^0(T; \operatorname{ad} \alpha) = 0$ the base components agree. The primitives agree on overlaps for the same reason, also across a joined bridge. The fibre estimate $\|\eta_T\| \leq \mu_T^{-1} \|\zeta_T\|$, the scalar spatial estimates on the mixed ends and the irreducible gauge input put η in the weighted H^1 space. Integration by parts with exhausting cutoffs gives $\|\zeta\|^2 = \langle d_A \eta, \zeta \rangle = \langle \eta, d_A^* \zeta \rangle = 0$; the boundary term vanishes by the normal and face conditions. Thus the kernel is zero.

The index. On the special quintuple [DFL21, Proposition 3.19] gives index zero for the realization of [DFL21]. That realization and the one of Definition 5.1 use the same decaying space and the same stationary parameter space; their stationary columns, including the scalar ones, differ by a bounded finite-rank map, so the indices agree (as in Paper I, Remark 10.24); a change of weight inside the gap does not change the index (Lemma 5.2). If the almost-complex structure near the matching line is not the one of [DFL21], a homotopy of compatible structures localized in the seam collar keeps the constant pair a solution and the operators

Fredholm, by Lemma A.1. For a DFL-type domain in the class of Paper I, Lemma 9.7, and for the closed-cut piece along the family in the proof of that lemma, the constant pair remains a solution. On each mixed end the metric is conformal to a product, and the perturbed anti-self-duality condition on two-forms is conformally invariant. The flat references do not depend on the end metric, and the set of stationary solutions at each mixed end, the perturbed-flat connections on its compression body, does not depend on the metric or the almost-complex structure (proof of Paper I, Lemma 9.7). The section is constant. The index is therefore constant along the family. For $W_{a,R}$ with finite R the index is constant along the collar-length family (Paper I, Theorem 10.13(i)). By Lemma 8.4(iii), which does not use the present lemma, $P_R(C_{a,x})$ is the constant pair on $W_{a,R}$ and $P_\lambda(C_{-,x-}, C_{+,x+}) = C_{J,x}^{(\lambda)}$. Hence Lemma 8.5(ii) gives $\text{ind } D_R = \text{ind } D_\infty = 0$ for large R , and the index at $C_{J,x}^{(\lambda)}$ is the value 0 at $\lambda = \infty$ for large λ ; it is constant in $\lambda \in (0, \infty)$ by Proposition 2.1(a). In all cases the kernel is zero and the index is zero, so the operator is invertible.

Uniformity. On a compact family of finite lengths the inverses vary continuously. Through torus completion the positive elements are continuous by Lemma 8.5(iii), and uniform inverse bounds for large lengths follow from its part (i), with $E = 0$ since the limiting operators are onto. $\square$

A.4. Proof of Proposition 2.1(c). *Compactness and gluing of pure moduli spaces.* The compression-body Lagrangians are simply connected with $\pi_2 = 0$, and $\pi_2(\mathbf{M}_h) \cong \mathbb{Z}$ is generated by a class of Maslov index four; in the product target the correspondence L_C is simply connected (Paper I, Section 11.3) and has π_2 generated by spheres of zero area and zero Chern number (Section 10). The pairs and correspondences are therefore monotone with minimal Maslov number at least three: for the compression-body Lagrangians this is [DFL21, Proposition 2.27], and for the correspondence L_C it is Paper I, Section 11.3. Sphere and disc bubbling has codimension at least two in the dimensions used. The compactification and boundary structure of the one-dimensional pure continuation, cap and strip moduli spaces are then as in [WW15b, Theorem 3.9(c)] for the continuation and cap spaces and as in [WW10, Theorem 4.2.3(c)] for the Floer strips. Transversality there is obtained from Hamiltonian perturbations. In the perturbation class used here, with zero Hamiltonian and domain-dependent almost-complex structures, it is Paper I, Lemma 11.3, for the strips and triangles from which the cup and cap maps are built, and Proposition 6.4(ii) for continuation strips. The ordinary Heegaard strips are regular by the definition of the symplectic complex (Paper I, Section 3.4). Strip gluing is [Sim14, Theorems 5.4.1, 5.4.3, 5.4.5 and 5.4.8], with the compactness theorem [MWW18, Theorem 3.3] and the boundary description of [MWW18, Theorem 3.7], whose gluing step is [Mau09], for the quilted disc families, and exponential decay on half-strips is [Sim14, Corollary 2.3.2]. For a finite collection of regular configurations with nondegenerate strip ends, every nearby zero-dimensional solution at a long neck is the unique correction of a preglued configuration, the correction tends to zero with the neck length, and the estimates are uniform over compact families. Along the path from the preglue to the solution the operators are invertible, by the uniform right inverse at the preglue and the smallness of the correction, as in the proof of Lemma 8.10; hence determinant transport preserves the canonical orientation. The unsigned low-dimensional counts agree with those of Paper I, Section 11.

Interior nodes. Linear gluing of Cauchy–Riemann determinant lines along strip ends is Lemma 3.3. For a family of real Cauchy–Riemann operators on surfaces degenerating at

an interior node, the determinant lines form a topological line bundle [WW15a, Proposition 2.4.12], and the line isomorphism is given by the exact sequence (18) of [WW15a, Theorem 2.4.5] and [WW15a, Corollary 2.4.7]; it is associative [WW15a, Proposition 2.4.8]. The sequence is proved there for a boundary node, and the interior case is stated there to be similar. In (18) the fibre at the node is a summand of the third term; at an interior node it is the fibre E_z of the bundle, a complex vector space (compare [WW15a, Remark 2.4.6]), which is even and canonically oriented, so the node contributes no additional sign. Numbering is as in arXiv:1503.07803v6. These statements localize, by the following stabilization argument. Choose stabilizations of a mixed or pure operator supported away from a disc containing the node and separated from the matching and boundary collars. Then the stabilized kernel near the node is obtained by the local gluing map of the cited proposition, and the rest of the operator is unchanged. Hence the line bundle structure and the gluing isomorphism hold for the real Cauchy–Riemann part of any such operator. Associativity follows from the cited associativity, with disjoint supports. Commutation with ordinary strip gluing follows because the two stabilizations and cutoffs have disjoint supports, as in the proof of Lemma 8.5(iv). In the collapse application the positive plumbing parameter is fixed first, and the endpoint insertion lies in an external strip.

A.5. Proof of the estimates of Lemma 7.1. We use the notation of the proof of Lemma 7.1. Paper I, Sections 7.1–7.7, proves its linear estimates for the linearization at a regular index-zero limiting solution, whose operator is invertible. Here $\{V_b\}$ is an arbitrary compact family of limiting configurations, and D_{V_b} may have kernel and cokernel. After the stabilization of that proof, the arguments of Paper I extend to the stabilized operators through the following four modifications. The exterior inverse of Paper I, Proposition 7.13, is proved for the augmented operator (D_{V_b}, j_b, Γ_b) . The common exponential rate β of the family replaces the decay of solutions. The attachment inverse of (18) (Paper I, Proposition 7.15) is proved with the exterior operator augmented by F and K_b . The kernel-patching estimate, proved in Paper I for the model strip (Proposition 7.8), is transplanted to the glued domain. The remaining steps use only bounds on the coefficients of the linearized operator, not the equation, and apply unchanged.

Proof of (i) and (ii). The linear estimates of Paper I, Section 7, apply here before assuming that V_b solves an equation. On the collapsed strip their inputs are a smooth path of flat compression connections and a smooth section with the prescribed boundary values. The Hodge, trace and energy identities use these conditions and bounded derivatives of their coefficients. The zero-order term of the section linearization is retained. On the unchanged pieces the operator is the full D_{V_b} . The compactness of B bounds all these coefficients. The end and transition estimates use only the fixed end operators, the separated stationary coordinates, and the common gaps. Choose the auxiliary exterior weight $\delta_* < \min(\beta, \gamma/2)$, where γ is the transition gap of Paper I, Proposition 7.9, as in Paper I, Proposition 7.13. Paper I’s exterior truncation gives the stationary-column error

$$O(\epsilon^{1/4}) + O(\epsilon^{(\beta-\delta_*)/4}),$$

uniformly over B . Its exponentially convergent coefficient estimates and their weak-adjoint counterparts therefore apply to the original configurations. Exact constant tails are not required. Put $A = 1 + |\log \epsilon|$. First work on the auxiliary exterior of Paper I, Section 7.7, with its free stationary limits and fixed positive end weight. On this exterior the localized estimate holds with a compact remainder on the unchanged pieces and the limiting section. It also holds for the adjoint. If the augmented exterior estimate failed, normalized fields with

image tending to zero would have a nonzero limiting field in $\ker(D_{V_b}, j_b, \Gamma_b)$ after passing to a subsequence of b . The compression Hodge estimate eliminates the fast complementary component, and harmonic matching identifies the remaining component with the limiting section. Compact convergence and the finite stationary and F tests detect the limiting field. The adjoint argument gives a field in $\ker(D_{V_b}, j_b, \Gamma_b)^*$. Both are impossible. Thus the augmented exterior inverse is uniformly bounded. Its finite columns and projections are supported off the new necks. Before attachment change the exterior weight from the fixed gap weight to $\delta_L \asymp A^{-1}$, using the small-weight exterior argument of Paper I. The stationary part is integrated explicitly; the normal part retains its fixed spectral gap. The finite compact columns and tests are unaffected. This gives the same polynomial loss at the small weight and avoids exponential weight costs for sources whose supports move out along an overlap.

Proof of (iii). The preglued domain itself requires an additional attachment argument. Remove the sources in the middle of each overlap using the stationary line inverse of Paper I, the transition inverse, with its small weight of order A^{-1} . This costs $O(A)$; cutoff derivatives of order A^{-1} leave bounded sources. Solve those sources on the augmented exterior and the transitions, setting each transition's stationary value equal to the corresponding free exterior value. The two stationary values are independent. Patch in the source-free bands. The normal components decay exponentially across a band of length comparable to $|\log \epsilon|$, so the patching error is bounded by $C_B A^N e^{-c|\log \epsilon|}$ for a fixed finite N . A Neumann series gives a right inverse. Cutting a homogeneous field in the same bands and applying the component inverses proves injectivity. The stabilization columns and projections occur only in the exterior equation and are unchanged by these operations. Consequently

$$\|A_{\epsilon,b}^{-1}\| \leq C_B A^N.$$

This is the source-removal proof of the attachment inverse in Paper I, with the exterior operator augmented by F and K_b . A bound independent of ϵ on the whole domain would be false: stationary fields supported on a long new neck have norm ratio of order $|\log \epsilon|$. The polynomial loss suffices below.

Proof of (iv). For $k \in K_b$, cutoff differentiation and the horizontal velocity of the harmonic lift give

$$\|A_{\epsilon,b} J_{\epsilon,b} k - (0, k)\| \leq C_B \epsilon^\sigma (1 + |\log \epsilon|)^N \|k\|,$$

for some $\sigma > 0$ and finite N . These are the endpoint and patching estimates of Paper I, Sections 7.4–7.7, namely the kernel estimate of Proposition 7.8 and the spatial estimate of Proposition 7.9. The fast lift has weighted size $O(\sqrt{\mathfrak{s}})$, where $\mathfrak{s}$ is the local compression scale ($\mathfrak{s} = \epsilon/r$ on the middle compression in the conformal representative of Paper I, Section 7.7, and $\mathfrak{s} = \epsilon e^s$ in the endpoint neck coordinate s). The kernel field is stationary on the truncation region, so the lift contributes only where $\mathfrak{s} \leq \epsilon^{3/4}$. The stationary factors are unchanged, and the fixed compact family has uniformly bounded coefficients. Take σ no larger than the positive exponents in these endpoint errors. The truncation estimates are uniform on the compact exponentially convergent family.

Choices, lifting actions and gluing. Two choices of stabilization are compared by their direct sum. The exact kernel sequence for adding the extra columns identifies the two cancellation maps. Cutoffs and compact pairings can be interpolated; the same estimates on the compact interpolation give a homotopy of the resulting maps. Lifting actions preserve the operators and the harmonic lift, so applying a lift to this construction gives the same determinant map after stabilization independence.

For ordinary end gluing choose the stabilizations off both transition regions. If T is the ordinary gluing length, the constants of the cofinal regime in Lemma 7.1 are chosen so that $e^{-cT} \mathbf{A}^{N'} \rightarrow 0$ for each of the finitely many polynomial losses occurring in the construction; $T \geq C \log \mathbf{A} + T_0$ with C, T_0 large suffices. Correcting the common patched kernel at once, or first at one neck and then at the other, now gives maps uniformly close to that patched map. Their straight interpolation on the finite-dimensional kernels is invertible. The threshold for ϵ may be taken uniform in $T \geq T_0$. Indeed the stabilizations avoid the neck. In the contradiction argument for the augmented exterior estimate a normalized sequence with $T_i \rightarrow \infty$ retains mass on the compact core of one of the two pieces, by the compact-remainder estimate and the gap on the neck; there it converges to a nonzero field annihilated by the augmented operator of that piece, which is an isomorphism. Thus the cofinal regime meets the range of the construction. Any two sufficiently large gluing choices connect through this cofinal regime. The ordinary gluing determinant maps are unchanged along that connection by continuity. We obtain the asserted compatibility, including its graded permutations.

A.6. Fredholm property of the family of Proposition 7.4. We use the notation of the proof of Proposition 7.4 and prove the estimates behind its Step 2.

For Fredholmness, first work on the weighted spaces of decaying fields. The local primal estimates are [DFL21, Theorem 4.17 and Remark 4.19]; the weak-adjoint estimates are [DFL21, Remark 4.33] and [DFL24, Theorem 5]. The matching and Lagrangian boundary conditions depend on the configuration; the maps of [DFL24, Proposition 5.27 and Lemma 5.67] transport them to fixed boundary conditions, and their differences from the identity are bounded by the distance between the configurations. The changed compact core is smooth at this fixed scale. At a mixed stationary limit B , the translation-invariant end operator with slice centre B has the form $\partial_\theta - \mathcal{A}_B$, with self-adjoint spatial operator and a positive normal gap γ . Its weighted whole-cylinder inverse satisfies

$$M_\delta \leq \frac{1}{\min(\delta, \gamma - \delta)}, \quad 0 < \delta < \gamma,$$

as well as the spatial first-order estimate. Stationary coordinates are adjoined separately. Changing the scalar centre to C adds a zeroth-order row of norm at most $c\|C - B\|_{C^0}$. If $cM_\delta\|C - B\|_{C^0} < 1/2$, a Neumann series supplies the affine cylinder inverse and the inverse of its adjoint, without a self-adjoint spectral decomposition of the affine operator. The true end coefficients and matching domains converge to these models, so their errors are absorbed after a common far cutoff. Ordinary end models are unchanged. Compact-core centre changes are relatively compact by Rellich. Hence, after weight conjugation,

$$\|v\|_{H^1} \leq C(\|D_{s,t}^{\text{dec}} v\|_2 + \|v\|_{L^2(K)}), \quad \|y\|_{H^1} \leq C(\|(D_{s,t}^{\text{dec}})^* y\|_2 + \|y\|_{L^2(K)}).$$

Adding the finite stationary columns preserves Fredholmness. Higher estimates and smooth finite stabilizations identify the resulting determinant lines with the strong realizations.

For clarity, the adjoint here is taken after weight conjugation. A zeroth-order change of slice centre leaves the principal Green form and its matching annihilator unchanged. Write $K_{s,t}(p) = D_{s,t}(0, p)$ for the full target column corresponding to the source lift $H_{s,t}$. A cokernel element satisfies both $(D_{s,t}^{\text{dec}})^* y = 0$ and $K_{s,t}^* y = 0$. A target map S pulls the weak domain back by S^* ; it does not leave the old adjoint boundary label fixed. The source chart, which includes a boundary trace extension, need not extend to a bounded map on bulk L^2 . We use it as a bounded isomorphism of Sobolev domains, which preserves the cokernel annihilator. Equivalently, complete its decaying source in the pulled-back norm $\|J_s^{\text{dec}} w\|_{L^2}$ before forming a coordinate adjoint. Keep the finite stationary space as a separate summand.

A.7. Boundary Hodge theory and holonomy regularity. For a flat orthogonal bundle on a compact compression body, the input needed in Proposition 10.3 is the Hodge decomposition on $W^{1,2}$ with normal boundary condition. Its harmonic fields are smooth and finite dimensional, and its exact part is the covariant differential of $W^{2,2}$ sections. Since $H^0(C; \text{ad}) = 0$, the Neumann Green operator G_N for $d_B^* d_B$ exists and preserves smoothness, so, on forms b with $*b|_{\partial C} = 0$, the map $1 - d_B G_N d_B^*$ is a smoothness-preserving projection onto the coclosed forms with $*b|_{\partial C} = 0$; the homogeneous Neumann condition leaves the normal component unchanged. The decomposition is the normal form of the flat covariant de Rham case of [KL25, Section 1.6 and Theorems 3.10–3.11], and the Neumann Green operator follows from [KL25, Theorem 3.10] in degree zero together with $H^0(C; \text{ad}) = 0$.

The boundary elliptic estimate has the form

$$\|b\|_{W^{1,p}} \leq C(\|d_B b\|_{L^p} + \|d_B^* b\|_{L^p} + \|b\|_{L^p})$$

for $*b|_{\partial C} = 0$. For smooth fixed B , adjoint bundle coefficients add only bounded zeroth-order terms in local trivializations. This estimate is Theorem D(a) of the introductory statements in [Weh04], and the compactness and Coulomb-slice inputs are Theorems A and F there.

Finite holonomy functions have bounded gradients and Lipschitz gradient differences as used in (10.2) [DFL21, Proposition 6.1]. Their higher derivatives obey the local tame estimates obtained by differentiating parallel transport. Combined with the boundary elliptic estimate, these give the bootstrap from the slice fixed point to the smooth perturbed-flat solution. For holonomy families with moving basepoints, restriction and composition are performed in the smooth topology.

A.8. Shifting and self-gluing. Relative preparation and shifting are proved in Appendix B. Naturality under isomorphisms of the decorated data is automatic: the pullback of a shift is a shift, and the orientation maps are independent of the shifting path (Proposition 5.4).

For ordinary gauge self-gluing, Proposition 2.3 applies the low-dimensional instanton compactification, verifies matching on the prescribed bundle and identifies the lifting fibre. Its orientation comparison first traces over the full determinant-one torus complex and then takes the larger quotient. Together with the finite group calculation in Section D.2, this gives integral unit multiplicity.

A.9. Proof of Lemma 8.5.

Proof of Lemma 8.5. Once a uniformly bounded right inverse is available, the construction is the standard gluing of determinant lines along a long neck, as in the proof of Lemma 3.3 and in [Don02, proof of Proposition 3.9 and Proposition 5.11]. Two points are specific to the torus bridge. The uniform estimates must be proved directly, because the cross-section of the bridge runs into the mixed ends. The stationary columns at finite and at infinite width differ on the whole mixed ends, not only on the tails.

Uniform estimates. The uniform inverse of Paper I, Proposition 10.27, is obtained by patching model estimates. On the torus region beyond a fixed S the operator is the flat odd-torus operator D_α of Paper I, Proposition 10.3, plus a small zeroth-order term, and the normal derivatives at a flat face are recovered as in the proof of Paper I, Theorem 10.5(b). On the mixed ends the operator is transported to the stationary models by the maps of [DFL24, Proposition 5.27], with errors tending to zero, and the slice estimates of Paper I, Proposition 10.7 and Corollary 10.4, hold. The ordinary asymptotic operators are invertible, and local elliptic estimates hold on a fixed compact region. The patching also uses that the coefficients are bounded in C^m on unit patches of finitely many isometry types, and that the

weight satisfies $\delta < \gamma/2$ and $\delta < \mu_T$, with the weight function equal to $\tau_0(\theta)$ on the torus region (Paper I, Definition 10.1). For the operators $D_\lambda(c)$ these inputs hold uniformly on $\mathfrak{F} \times [2r_0 + 2, \infty]$. The smallness beyond S follows from the uniform tails in the definition of torus normal form and Sobolev embedding, and the other inputs follow from the local estimates and end models of Section A.1 and the compactness of $\mathfrak{F}$. The finite stationary columns and the scalar row of the realization are bounded, and they act through the fixed cores and the mixed ends. The same patching therefore gives

$$\|x\| \leq C(\|D_\lambda(c)x\| + \|x\|_{L^2(K_1)} + |p(x)|),$$

where K_1 is a fixed compact subset of the two cores and $p(x)$ the stationary coordinates. The adjoint of D_α is invertible with the same bound, the stationary models are self-adjoint, and the local estimates for the weak adjoint are recorded in Section A.1. Hence the same patching gives

$$\|\eta\| \leq C(\|D_\lambda(c)^*\eta\| + \|\eta\|_{L^2(K_1)})$$

on the dual weighted spaces.

(i) At $\lambda = \infty$ each $D_\infty(c)$ is Fredholm by Proposition 2.1(a). Smooth compactly supported target sections are dense, so finitely many of them span a complement of its range, and by continuity also near c . A finite cover of $\mathfrak{F}$ gives Λ , with sections supported in a compact set K_Λ ; for $\lambda \geq \lambda_1$ this set lies in the part of W_λ^J identified with the completed pair, so Λ is defined there. If every $D_\infty(c)$ is onto, no sections are needed and $E = 0$. Suppose there were $\lambda_j \rightarrow \infty$, $c_j \in \mathfrak{F}$ and η_j with $\|\eta_j\| = 1$ and $\|(D_{\lambda_j}(c_j) \oplus \Lambda)^*\eta_j\| \rightarrow 0$. By the adjoint estimate, $\|\eta_j\|_{L^2(K_1)} \geq 1/C$. Pass to $c_j \rightarrow c$. The parts of $W_{\lambda_j}^J$ within distance $\lambda_j/2$ of each core are isometric to the corresponding parts of $W_{\pm, \infty}$. Local elliptic bounds give a subsequence converging weakly on compact sets and strongly in $L^2(K_1)$ to some η_∞ on the completed pair. The stationary columns converge by the estimate for $(C_\lambda - C_\infty)p$ in (ii) below, so $C_{\lambda_j}^*\eta_j \rightarrow C_\infty^*\eta_\infty$. Then $(D_\infty(c) \oplus \Lambda)^*\eta_\infty = 0$ and $\eta_\infty \neq 0$, which contradicts the choice of Λ . Hence $\|\eta\| \leq C\|(D_\lambda(c) \oplus \Lambda)^*\eta\|$ for $\lambda \geq \lambda_1$. Then $Q_\lambda = T^*(TT^*)^{-1}$, with $T = D_\lambda(c) \oplus \Lambda$ and T^* its Hilbert adjoint for the weighted inner products, is a uniformly bounded right inverse. The estimate for the formal adjoint on the dual weighted space is equivalent to one for T^* through the Riesz isometry; it is proved first at the H^1 level and then upgraded by the local elliptic estimates on unit patches. The map Q_λ is continuous in (c, λ) for finite λ .

(ii) At $\lambda = \infty$ the kernels form a vector bundle over $\mathfrak{F}$, since the operators are onto. The frames ι_λ are the transplanted kernels of [Don02, Step 3 of the proof of Proposition 3.9], with Q_λ in place of the patched inverse there. They are injective for large λ and converge to the identity once $\|(D_\lambda \oplus \Lambda)\beta_\lambda v\| \rightarrow 0$ uniformly for unit v , where β_λ acts on the decaying part and leaves the stationary and E -components unchanged. The unit sphere bundle of the limiting kernels is compact, so the tails of its elements on $r \geq r_1$ are uniformly small as $r_1 \rightarrow \infty$. The error consists of $\sigma(d\chi_\lambda)v$ and of the coefficient difference on the tail, both supported where $r \geq \lambda/2 - 1$, and of the difference $(C_\lambda - C_\infty)p$ of the stationary columns, read through the identification of (i). The last term is not localized on the tail, but it tends to zero uniformly for $|p| \leq 1$. The spatially Coulomb stationary representatives on the joined slices $D_{i, \lambda}$ and on the separate slices $Y_{\pm, \infty}$ differ by the gradients of their scalar Neumann corrections, whose sources coincide and are supported away from the torus neck. On the neck the scalar Laplacian is bounded below by μ_T^2 , the square of the fibre gap, so the two corrections differ by $O(e^{-\mu\lambda})$ in every norm, for some $\mu > 0$, uniformly in the stationary limits, as in the proof of Paper I, Proposition 10.7. The frames are also onto: otherwise unit kernel elements orthogonal to their image would, by the estimate for D_λ and the argument in

(i), converge to a nonzero kernel element at $\lambda = \infty$ orthogonal to the whole limiting kernel. For finite λ the kernels of the surjective family form a bundle as usual.

(iii) This is the stabilization construction of [Don02, Step 4 of the proof of Proposition 3.9, and Proposition 5.11], as in the proof of Lemma 3.3. The exact sequence $0 \rightarrow \ker D \rightarrow \ker(D \oplus \Lambda) \rightarrow E \rightarrow \operatorname{coker} D \rightarrow 0$ identifies $\det D$ with $\det \ker(D \oplus \Lambda) \otimes (\det E)^*$, and (ii) makes these lines continuous. At $\lambda = \infty$ the ordered direct-sum isomorphism gives $\det D_{c_-} \otimes \det D_{c_+}$. Two choices Λ, Λ' are compared through $\Lambda \oplus \Lambda'$ by the standard stabilization maps, which are continuous in (c, λ) because they are built from the same kernel bundles. On an invertible operator the stabilization identification carries the canonical element to the canonical element, so the positive elements are continuous where the operators are invertible.

(iv) The sections of Λ are supported in K_Λ , which lies at a fixed distance from the ordinary gluing neck and from the torus bridge once both are long. The glued operators then form a family over the product of the two closed neck-length intervals, including their limits. The kernel transplants of (ii) and of Lemma 3.3 act by cutoffs and projections in disjoint regions, and the estimates of (i) hold uniformly on this product. The product of the two intervals is contractible, so the two orders of gluing give the same identification.

For one side use the half-plane and strip models of Paper I, Proposition 10.3, with the flat-fibre face, in place of the bridge. $\square$

A.10. Proof of Proposition 5.3. We use the notation of Section 5.2: the autonomous end operators $\partial_s - \mathcal{A}_e$ in outward coordinates, the gaps $\gamma_e, \gamma_e^{\text{opp}}$ and the weights $r_e < \gamma_e$, and P for the finite summand of mixed-end stationary coordinates.

Proof of Proposition 5.3. All comparisons use closed intervals strictly inside the gaps. At spatially completed torus ends no extra spatial weight is introduced. The longitudinal weights are constant on surface fibres, have zero normal derivative at the boundary faces of the domain and agree on matching traces. Write each longitudinal exponent as $r_e W_e$. Choose $|dW_e| \leq 1$, with cutoff derivatives supported in disjoint longitudinal transitions. The complete stationary mixed operator is self-adjoint on the domain given by its boundary conditions, by Green's formula. On a completed compression, its compact-remainder estimate and the acyclic torus model give a normal gap, without compactness of the resolvent. The following argument uses its spectral projections.

Solutions decay exponentially at some positive rate along every end, by the mixed-end estimates of Paper I and [DFL21, Section 5.3]. To raise the rate, choose a far end and put the solution in the small Coulomb slice of the stationary connection, with a Dirichlet condition on an artificial initial slice and the Neumann condition of the problem at the boundary faces. The small-weight scalar inverse provides this gauge and preserves the chosen stationary limit of the gauge transformation. Subtract the exact stationary primary limit and use its translation-invariant matching chart. The equation becomes

$$(\partial_s - \mathcal{A}_e)v = f(v), \quad f(0) = df(0) = 0, \quad v(s) \rightarrow 0,$$

and the chart and holonomy estimates bound $f(v)$ quadratically in v on nested unit windows, with a fixed finite loss of derivatives. For such an equation the standard bootstrap, by variation of constants in the spectral decomposition of $\mathcal{A}_e$ and local elliptic estimates on nested windows, raises the decay rate to every rate below the spectral gap; compare [Don02, Section 4.1] for nondegenerate gauge ends. It gives every rate $b < \gamma_e$ with all derivatives, and uniform bounds on compact families with common tail-entry windows and spectral margins. Two points are specific to mixed ends. The component in the stationary kernel is integrated from infinity and uses the prescribed limiting primary coordinate. The component in the

positive spectral subspace is also integrated from infinity, so the rate is not restricted by the opposite gap γ_e^{opp} . A rate $b > r_e$ gives membership in the desired weighted space.

Gauge transformations between these representatives have their prescribed stationary limit g_∞ . The temporal comparison equation $\partial_s g = g a_s - a'_s g$, integrated from infinity, gives the same rates for $g - g_\infty$. The spatial comparison equation and differentiation give all higher derivatives, after subtracting the stationary intertwining relation. The lifting classes and central transition data are unchanged. This concerns gauge transformations between the representatives controlled above, not arbitrary gauge transformations applied to a stationary solution.

Parameter derivatives require differentiating the stationary subtraction. In a family the identity $\mathcal{F}_b(0, q) = 0$ cancels the nondecaying parameter column. Its remainder and every finite stationary column then have all the above decay rates. The linearized equation improves the first parameter derivative; induction improves higher ones. Polynomial factors caused by parameter differentiation are absorbed in a strict rate margin. A common data parameter contributes only its one original column.

Write the full realized operator as

$$D_r : X_r^{\text{dec}} \oplus P \longrightarrow Y_r, \quad D_r(w, p) = D^{\text{dec}} w + C p.$$

On an end, weight conjugation gives $\partial_s - \mathcal{A}_e - r_e$. Its inverse and adjoint inverse are uniformly bounded on a closed weight interval by the reciprocal spectral distance. Small coefficient and domain-transport errors are absorbed beyond a common cutoff. On an infinite spatial torus region the unweighted flat operator has mass $\mu_T > 0$. The conjugation term has norm at most the largest relevant mixed-end exponent r_T ; its inverse bound is at most $(\mu_T - r_T)^{-1}$. The permitted completed spatial gap is bounded by this essential-spectrum threshold. The original closed incoming weights do not vary on that added torus region. Compact-core estimates for the operator and for its weak adjoint give Fredholmness over the entire weight rectangle. Smooth stationary and parameter columns remain in every target. If a scalar inverse at a larger weight is needed, scalar Fredholmness throughout its indicial gap, index zero from the small weight and the positive kernel energy identity give it. Scalar gauge modes occur in the full end operator, so the full permissible gap is no larger than that scalar threshold. This avoids extending the sufficient small-weight coercivity bound beyond its range.

For $r_- \leq r_+$ the inclusions of the weighted spaces give a commuting square of source and target inclusions, with identity on P . A kernel element is a stationary column plus a remainder. The coefficient errors and the images of stationary columns have every permitted rate. The same spectral formulas improve the remainder to every weight. This applies also to kernels with compactly supported stabilization columns and to kernels of parametrized operators.

For the cokernel use distributions paired in the unweighted volume form of the domain. A weak cokernel element at r_+ lies at dual weight $-r_+$. Window estimates initially bound its growth by e^{r+s} . Choose β with $r_+ < \beta < \gamma_e$; multiplying by the coefficient error gives a decaying forcing of rate $\beta - r_+$. The adjoint spectral formula excludes growing normal modes, whose rates are at least γ_e . Its zero mode has a finite limit. The resulting bounded field improves its forcing again, and its remainder has every rate below $\min(\gamma_e, \gamma_e^{\text{opp}})$. It therefore belongs to the dual space at $-r_-$ as well. At an ordinary end there is no zero mode, so $r_- = 0$ is included. Retain the conditions $C^* y = 0$ and, in a parametrized problem, the parameter-column annihilator. These are the same convergent pairings at every weight. Hence the kernels and cokernel duals, and therefore the finite-dimensional cokernel quotients,

are canonically identified. For unordered weight vectors insert their componentwise minimum and maximum.

Choose finitely many smooth compactly supported target columns $\rho : V \rightarrow Y$ spanning the cokernel on a compact parameter patch. They work at every weight. With $E_r = \ker(D_r + \rho)$, the identity map of fields identifies

$$\det D_r = \det E_r \otimes (\det V)^*.$$

Enlarging the stabilization gives the same map by its exact sequence. These identifications are continuous, transitive and preserve canonical orientations at invertible operators. They commute with the gauge transformations above and the shared parameter factor. The weight rectangle is contractible, so its transport has no additional choice of path or component sign.

The reference-orientation paths can initially have slower tails. Prepare them by cutting off their decaying coordinates on a common far end, retaining the exact stationary coordinates and carrying the transition data of the bundle through the preparation transport. The preparation homotopy is made at the initial small weight. A stationary endpoint is fixed; at a solution the endpoint segment has all permitted decay rates by the preceding argument. Thus the prepared path and its endpoint segments admit the full weight rectangle, even when the original path does not. Determinant transport on this rectangle and on the preparation homotopy proves compatibility with the original reference normalization. Affine-centre comparisons are made first at their common small weight, as in Proposition 7.4. The comparison of solutions then uses the full gap of this proposition; the spectrum of the original end operator is not assigned to an affine operator.

Finally consider a sequence in an already established small-weight compactification. Fix translations on its retained components. Energy exhaustion gives common tail-entry windows. The estimates above, with $r < b < \gamma_e$, bound a tail beyond R by a constant times $e^{-(b-r)R}$ in the desired weighted norm. Compact convergence therefore improves to convergence in that weight. At ordinary breaks this is a componentwise statement in the gluing charts; at spatial completion it retains the existing strong spatial tails. The small-weight forward gluing produces the same solutions, and the converse and uniqueness give the same local branches in larger weights. Common compactly supported stabilizations make the gluing exact sequences commute with the inclusions of the weighted spaces. Cutoffs, evaluation maps, stationary factors, and conversion lines are the same ordered geometric maps. Thus all boundary signs, counts and homotopies are unchanged. $\square$

APPENDIX B. RELATIVE PREPARATION AND SHIFTING

This appendix proves Proposition 2.2. The one-input statement is a version of [DFL21, Lemma 4.62] and of its family form [DFL21, Lemma 4.64], whose proof is sketched there; we give it in the topology of Section 2, jointly continuous in configurations and data, relative to the terminal subspace defined below, and on the domains of this paper. Throughout, a family is a continuous map from a compact metrizable space K to configurations on a fixed bundle, that is, to connections and sections rather than gauge equivalence classes, together with continuously varying data. No quotient by gauge transformations is taken; the lifting groups enter only through Proposition 5.4.

The terminal subspace. For an output label x , a *representative* of the critical point over x is a connection $c = B \cup \text{pr}^*b \cup B'$ on $Y_\#$, where b is a flat connection on Σ over the point of $\mathcal{M}(\Sigma)$ determined by x and B, B' are primary solutions on Y_0, Y'_0 which are products with value b on the collars $[-5/2, -2] \times \Sigma$ of their boundaries fixed in Section B.2. Every such c is a critical point of the perturbed Chern–Simons functional in the gauge orbit of a_x or

of its other lift, and the constant pair C^c built from it is equivalent to C_x under the lifting group, by a gauge transformation when c lies in the orbit of a_x . Its operator is therefore invertible, its positive element corresponds to that of C_x under conjugation, and the almost-complex gauge orientations are invariant under the bundle automorphism. The lines at c are identified with those at a_x by this automorphism. The identification agrees with transport along any path of connections from a_x to c : the almost-complex orientation of the gauge line is invariant under the full group of bundle automorphisms, of every degree and lifting class (Section 4, [KM11b, Section 5.1]), and the sign of the conversion line depends only on the lifting class (Proposition 4.2), so it is trivial for automorphisms of trivial lifting class, of any degree. For a representative in the orbit of the other lift, the relating automorphism has lifting class in the framing group, on which this sign is trivial by the last clause of Proposition 4.2, since all structures agree with the fixed germ near the odd torus; at the cut Z this is (D.1). The terminal subspace $\mathcal{T}_x$ consists of the configurations which, outside the gauge input end, agree with C^c for some representative c ; they are gauge paths ending at c followed by C^c . Proposition 2.2 is to be read with this terminal subspace. It cannot be replaced by the configurations ending at one fixed C_x : those have a single topological energy, whereas attaching a sphere of nonzero area changes the topological energy of a configuration without changing its labels (a, x) .

B.1. Reduction to a special quintuple. Let (Y, Y') be a pair of odd compression bodies with common boundary Σ , a surface with two components on each of which the bundle is odd (Paper I, Definition 3.2), with primary data whose Lagrangians $L(Y), L(Y') \subset \mathcal{M}(\Sigma)$ are embedded (Paper I, Section 3.2). For the torus domains $\Sigma = \Sigma_h \sqcup T$; for the closed-cut piece and the pairs with product target $\Sigma = \Sigma_h \sqcup \Sigma_g$. Let X_D, U_+ be the gauge region and section domain of the special quintuple [DFL21, Section 3.1]: the gauge region is $\mathbb{R} \times Y_0 \cup U_- \times \Sigma \cup \mathbb{R} \times (-Y'_0)$, the matching seam is $U_\partial \times \Sigma$, the Lagrangian arcs η_+, η'_+ carry $L(Y)$ and $L(Y')$, and the gauge input end is $(-\infty, -3] \times Y_\#$. Every flat connection on Σ is irreducible with central stabilizer.

The configuration space of Section 2 depends only on the decorated smooth manifold (bundle, matching seam, arcs with their Lagrangians, ends with their limits) and on the primary data. It does not depend on the metric, the almost-complex structure or the secondary data: perturbed flatness involves no metric, and the weights are functions of the end coordinates. Each one-input domain of Proposition 2.2(a) with finite torus length is identified with a special quintuple by a diffeomorphism of decorated manifolds, depending continuously on the lengths and widths: for $W_{a,R}$ the map ψ_R of Paper I, Definition 10.1(a), which is independent of θ and is a translation near the flat face; for the closed-cut piece with input Z the identification of Paper I, Lemma 10.18(ii); for the domains defining the mixed maps of Paper I, Lemma 9.7, a decorated diffeomorphism which carries end coordinates to end coordinates. It is the composite of a diffeomorphism supported in a compact set, which realizes by the isotopy extension theorem the compactly supported changes of the seam placement and of the neck lengths by which a domain of (II) differs from the special quintuple, and of a diffeomorphism which on each mixed end is the identity in the end coordinate times a diffeomorphism of the cross-section. The second factor accounts for the translation-invariant data on the mixed ends, which that lemma allows to differ from those of the quintuple; an example is the width of the section interval, which is π for the polar ends of Paper I, Lemma 6.3(ii). Both factors depend continuously on the parameters of these changes, by isotopy extension with parameters. Since the weights are functions of the end coordinates, the weighted smooth topology is preserved, and the metric does not enter the configuration space. The output regions used

in part (b) are identified in Section B.5. It therefore suffices to construct the deformations on a special quintuple, for families in which the primary data vary in a compact set.

B.2. Preparation of the ends and collars. Write $C_Y = [-2, 0]_s \times \Sigma$ for the collar of $\Sigma = \partial Y$ in Y which forms part of $U_- \times \Sigma$ on the mixed ends [DFL21, Section 3.1], and $C_{Y'}$ likewise. It extends to a flat collar $[-3, 0] \times \Sigma$ whose part $[-3, -2] \times \Sigma$ in Y_0 is free of the primary perturbation, so that primary solutions are flat on it. This is arranged as follows. The primary perturbations of a compact family are supported in finitely many fixed solid tori in the interior of Y_0 , so some collar of ∂Y_0 in Y_0 avoids them, and a decorated diffeomorphism supported in Y_0 , as in Section B.1, makes this collar $[-3, -2] \times \Sigma$. The collars of ∂Y_0 used below are $[-9/4, -2] \times \Sigma$, in the normal form, and $[-5/2, -2] \times \Sigma$, in the definition of $\mathcal{T}_x$. We use the following normal form, with constants $\Theta \geq 2$ and $S \geq 3$:

- (N1) u is the constant map to x on $[S, \infty) \times [-1, 1]$;
- (N2) near η_+ and η'_+ , u is independent of the normal coordinate;
- (N3) on the mixed-end regions $|\theta| \geq \Theta$ the configuration is a stationary pair (π^*B, q) , (π'^*B', q') , where B equals pr^*b on $[-9/4, 0] \times \Sigma$, that is, on C_Y and on the collar of width $1/4$ of ∂Y_0 in Y_0 , for a flat connection b on Σ , and similarly for B' ;
- (N4) $u(s, \theta)$ is independent of s for $s \in [0, 1]$;
- (N5) on $[-1, 0] \times \mathbb{R}_\theta \times \Sigma$ the connection is the pullback of a connection $b(\theta) + \psi(\theta) d\theta$ on $\mathbb{R}_\theta \times \Sigma$ whose slices $b(\theta)$ are flat, with $\psi = 0$ for $|\theta| \geq \Theta$.

Lemma B.1. *Let $\mathcal{A}(a, x)$ be the space of smooth configurations on the special quintuple with gauge input a and output x , in the topology of Section 2. Let K be a compact metrizable space and $f : K \rightarrow \mathcal{A}(a, x)$ a continuous family, with primary and secondary data varying continuously in a compact set. Then there are constants $\Theta \geq 2$ and $S \geq 3$ and a deformation f_t , $t \in [0, 1]$, jointly continuous in (k, t) and in the data, with $f_0 = f$ and f_1 in normal form. It is the identity on the gauge input end $(-\infty, -3] \times Y_\#$. Members that are stationary on the original ends beyond Θ and S remain so, and members of $\mathcal{T}_x$ are fixed.*

Proof. *Collar gauge at the mixed ends.* Let B_q be the limiting connection of the configuration at the $+$ mixed end, a representative of q ; it is continuous in the family and flat on the flat collar. Let K be its temporal gauge along the collar, the solution of $\partial_s K = -(B_q)_s K$ on $[-3, 0]$ with $K(0) = 1$. Then K^*B_q has no ds component on $[-3, 0] \times \Sigma$ and, being flat, is a product there. Choose a smooth $\sigma_N : [-3, 0] \rightarrow [-5/2, 0]$ with $\sigma_N(s) = s$ on $[-9/4, 0]$ and $\sigma_N = 0$ near -3 , and put $k_\lambda = K(\lambda\sigma_N(s))$, extended by 1. This is a path of gauge transformations from 1 to $k = k_1$, with $k_\lambda = 1$ at the seam for every λ , and $k^*B_q = \text{pr}^*b_q$ on $[-9/4, 0] \times \Sigma$, because $k = K$ there. Apply the gauge $k_{\lambda(\theta)}$ on the $+$ mixed end, with $\lambda(\theta) = 0$ for $\theta \leq \Theta_0$ and 1 for $\theta \geq \Theta_0 + 1$, and do the same at the $-$ end. Gauge transformations preserve the matching condition, and this one is the identity at the seam. If the limit representative is already a product on $[-5/2, 0] \times \Sigma$, then $K = 1$ on $[-5/2, 0]$, which contains the range of σ_N , so $k_\lambda = 1$; in particular members of $\mathcal{T}_x$ are fixed.

Ends. On the strip output end write $u = \exp_x \xi$ in an exponential chart at x in which both Lagrangians are linear, and put $u_t = \exp_x((1 - t\beta)\xi)$ with a cutoff $\beta(s)$ equal to one for $s \geq S$. On a mixed end, beyond some $\Theta_1 \geq \Theta_0 + 1$ the configuration lies in the matching chart Ψ_{R_q} (Section A.1) centred at its stationary limit $R_q = (\pi^*k^*B_q, q)$: $c = \Psi_{R_q}(\zeta)$ with ζ in the linearized boundary conditions. Put $c_t = \Psi_{R_q}((1 - t\beta(\theta))\zeta)$ with $\beta = 1$ for $\theta \geq \Theta$. The chart is fibrewise over the base, and the linear conditions are pointwise in θ , so they are preserved; the chart gives the nonlinear matching and Lagrangian conditions. A compact family has uniformly small tails in the topology of Section 2, so common Θ_1, Θ, S exist, and the charts

depend continuously on the limits and the primary data. Members already stationary beyond Θ remain so.

Collars. Let $\phi_0 : [0, \infty) \rightarrow [0, \infty)$ be smooth and nondecreasing, with $\phi_0 = 0$ on $[0, 1]$ and $\phi_0(s) = s$ near and beyond 2, and put $\phi_t = (1 - t)\text{id} + t\phi_0$. Pull back u by ϕ_t in the normal coordinate of the seam in U_+ , and A by $(s, \theta, y) \mapsto (-\phi_t(-s), \theta, y)$ on $[-2, 0] \times \mathbb{R} \times \Sigma$; since ϕ_t is the identity near 2, the result is smooth where $[-2, 0] \times \Sigma$ meets Y_0 . Since $\phi_t(0) = 0$, the seam values, hence the matching condition, are unchanged. Do the same in the normal coordinates of η_+, η'_+ , which preserves the arc values. On the strips $[0, 2] \times [2, \infty)$ both maps act on the same coordinate and do not commute. In the coordinate $s \in [0, 2]$ across such a strip, with the seam at $s = 0$ and the arc at $s = 2$, write $\Phi_{\text{seam}}(s) = \phi_t(s)$ and $\Phi_{\text{arc}}(s) = 2 - \phi_t(2 - s)$. Take the arc pullback first and the seam pullback second, so that u is replaced by $u \circ \Phi_{\text{arc}} \circ \Phi_{\text{seam}}$. At $t = 1$ this composite maps $[0, 1]$ to the seam, because $\phi_0 = 0$ there and $\Phi_{\text{arc}}(0) = 0$, and a neighbourhood of 2 to the arc, because ϕ_0 is the identity near 2; for every t it fixes the seam and the arc pointwise. At $t = 1$, (N2), (N4) and (N5) hold; stationarity beyond Θ is preserved, and so is (N3), because pr^*b is independent of s . Members of $\mathcal{T}_x$ are independent of s on these collars and are fixed.

All three steps are explicit, continuous in the data, and supported away from the gauge input end. $\square$

B.3. Positive extension, shift and straightening. Use a variant of the reparametrizations $\Phi_{\pm}$ of [DFL21, Section 4.4]: polar about the corners for $|\theta| \geq \Theta + 1$, with distance $e^{-\theta}$ to the corner as in their formula, preserving the lines of constant θ near the seam for $|\theta| \leq \Theta$, smooth in between, where the prepared configuration is stationary, and carrying the normal collars of η_+, η'_+ onto collars of $t = \pm 1$. DFL take Φ_+ polar on the whole arms $[0, 2] \times [2, \infty)$; the present variant depends on Θ , hence on the family. In these coordinates the gauge region is $(-\infty, 0] \times Y_{\#}$ with $\{0\} \times Y_0$ and $\{0\} \times Y'_0$ removed, the section domain is $[0, \infty) \times [-1, 1]$ with the corners $\pm i$ removed, the matching seam is $\{0\} \times [-1, 1] \times \Sigma$ against $\{0\} \times [-1, 1]$, and the collars $C_Y, C_{Y'}$ of the mixed ends collapse to $\{0\} \times \partial Y_0$ and $\{0\} \times \partial Y'_0$, the collar coordinate becoming the angle about these corners. By (N3)–(N5) a prepared configuration is independent of the normal coordinate near $r = 0$, it is pr^*b with no angular component near the corners, and its seam restriction is $b(\theta(t)) + \tilde{\psi}(t) dt$ with $\tilde{\psi}(t) = \psi(\theta(t))\theta'(t)$, which vanishes near $t = \pm 1$.

Lemma B.2. *Let $\mathcal{P}$ be the space of smooth primary solutions on Y_0 which are products on the collar $[-9/4, -2] \times \Sigma$ of ∂Y_0 , and let $\mathcal{S}$ be the space of smooth flat connections on Σ whose class lies in $L(Y)$; a torus component of Σ contributes only the point of $\mathcal{M}(T)$. Restriction $\rho : \mathcal{P} \rightarrow \mathcal{S}$ is a locally trivial fibre bundle, also over a compact family of primary data. A continuous family $a : K \times [0, \infty) \rightarrow \mathcal{S}$, smooth in r with all r -derivatives continuous, with a lift $B(\cdot, 0)$ of $a(\cdot, 0)$, has a lift B with the same regularity, constant in r where a is.*

Proof. Near $a_0 = \rho(B_0)$ choose a smooth local section $[a] \mapsto \hat{B}$ of $\mathcal{P} \rightarrow L(Y)$ by a primary slice; the primary moduli space is regular and its restriction map is an embedding (Paper I, Section 3.2). Then $\rho(\hat{B})$ and a represent the same class, so $\rho(\hat{B}) = g^*a$ for a gauge g near 1 on Σ , unique up to the centre, by the surface Coulomb slice and irreducibility on both components. Extend g over the collar by $\exp(\beta \log g)$, with $\beta = 1$ on $[-9/4, -2] \times \Sigma$ and 0 at depth; this gives a local section with $\rho = a$. Thus ρ is locally trivial; its fibre is the space of gauge transformations of Y_0 which are central constants on the collar, modulo the centre. On $K \times [0, \infty)$ patch the local sections with a partition of unity that is smooth in r to a connection on the pulled-back bundle; the horizontal lift solves an ordinary differential

equation $g'g^{-1} = \xi$ pointwise on Y_0 , so it has the stated regularity and is constant where a is. $\square$

Positive extension. For a prepared family, let $\alpha(r, t)$ be the horizontal lift of $r \mapsto u(r, t)$ in flat connections on Σ , with $d_\alpha^* \partial_r \alpha = 0$ and initial value $b(\theta(t))$; on a torus component it is constant in r , since $H^1(T; \text{ad}) = 0$. It exists, is unique and depends smoothly on u , since the flat connections are irreducible; near $t = \pm 1$ it is independent of t by (N2). By Lemma B.2 choose $B(r)$ on Y_0 and $B'(r)$ on $-Y'_0$, products with values $\alpha(r, \pm 1)$ on the collars of (N3), with $B(0), B'(0)$ the limit representatives of (N3). Let A_u be the connection on $[0, \infty) \times Y_\#$ equal to $B(r)$ on Y_0 , $B'(r)$ on $-Y'_0$, and $\alpha(r, t) + \tilde{\psi}(t) dt$ on $[-1, 1] \times \Sigma$, with no dr component. By (N3)–(N5) the configuration and A_u are independent of r near $r = 0$ and agree there, so the connection $\tilde{A}$ on $\mathbb{R} \times Y_\#$ equal to A for $r \leq 0$ and to A_u for $r \geq 0$ is smooth. For $r \geq S$ it is independent of r ; write c_S for its restriction there, whose Σ -slices $\alpha(S, t)$ all represent the point determined by x .

Lemma B.3. *For $\tau \in [0, S + 4]$ put*

$$A_\tau(r, y) = \tilde{A}(r + \tau, y) \quad (r \leq 0), \quad u_\tau(r, t) = u(r + \tau, t).$$

This is a family of configurations with labels (a, x) , jointly continuous in (k, τ) and the data, beginning at the prepared family, and it translates the gauge input end by τ . At $\tau = S + 4$ the region $r \geq -4$ carries $\text{pr}^* c_S$ and $u_\tau \equiv x$. Members of $\mathcal{T}_x$ remain in $\mathcal{T}_x$ for all τ .

Proof. At the seam the slice of A_τ over t is $\alpha(\tau, t)$, which represents $u(\tau, t) = u_\tau(0, t)$; the arcs carry $u(r + \tau, \pm 1) \in L$. Near a removed corner the pair is smooth in the reparametrized coordinates, with limit the stationary pair of $B(\tau)$, whose collar value is $\alpha(\tau, 1)$, a product. In the original end coordinate it converges to this stationary pair at the exponential rate one. This exceeds the mixed-end weight, which we choose below one. Every condition imposed on that weight here and in Paper I is an upper bound, and by Proposition 5.3 the counts do not depend on the choice. The configuration therefore lies in the space of Section 2, with the mixed limit moving along the arc. Translation over a bounded range is continuous in that topology. For $\tau = S + 4$ and $r \geq -4$ we have $r + \tau \geq S$. If a member lies in $\mathcal{T}_x$, then $u \equiv x$ and its seam restriction is $\text{pr}^* b$, so $\alpha \equiv b$, $\tilde{\psi} = 0$, B, B' are constant, and $\tilde{A}$ is $\text{pr}^* c$ on $[-3, \infty) \times Y_\#$; the translates keep this form. $\square$

Straightening. At $\tau = S + 4$ contract the t -dependence of c_S . The slices $\alpha(S, t)$ are independent of t for $|t| \geq t_0$, for some $t_0 < 1$, by (N2). Choose a smooth nondecreasing $\rho : [-1, 1] \rightarrow [-1, 1]$ with $\rho(t) = t$ for $|t| \leq t_0$ and $\rho = \pm 1$ near ± 1 . For $\lambda \in [0, 1]$ let c_λ be the connection on $Y_\#$ equal to $\alpha(S, (1 - \lambda)\rho(t)) + (1 - \lambda)\tilde{\psi}(t) dt$ on $[-1, 1] \times \Sigma$ and to the lifts B_λ, B'_λ of Lemma B.2 of the paths $\lambda \mapsto \alpha(S, \pm(1 - \lambda))$, starting at $B(S), B'(S)$. Near $t = \pm 1$ the slices are the constant values of these paths, so the pieces join smoothly, and at $\lambda = 0$ the connection is c_S . With a smooth χ equal to 0 for $r \leq -5$ and to 1 for $r \geq -4$, replace the connection A_{S+4} by $A_{S+4} + \chi(r)(\text{pr}^* c_\lambda - \text{pr}^* c_S)$ and keep $u \equiv x$. This is $\text{pr}^* c_\lambda$ on $[-4, 0] \times Y_\#$ and the unchanged connection for $r \leq -5$; the interpolation lies in the gauge input end, where no boundary condition is imposed, and at $\lambda = 0$ nothing changes. The slices represent the point of x , so matching holds, and the mixed limits are stationary pairs of B_λ, B'_λ . At $\lambda = 1$ the connection $c_1 = B_1 \cup \text{pr}^* b \cup B'_1$ is independent of t on the neck, with B_1, B'_1 products on the collars of (N3). On members of $\mathcal{T}_x$ the slices $\alpha(S, \cdot)$ are constant and $\tilde{\psi} = 0$, so this step is the identity.

Collar widening. Let K_1 be the temporal gauge of $B_1 \cup \text{pr}^* b$ along the flat collar, as in the proof of Lemma B.1; it equals 1 on $[-9/4, 0]$. Choose a smooth $\sigma_W : [-3, 0] \rightarrow [-3, 0]$

with $\sigma_W(s) = s$ on $[-5/2, 0]$ and $\sigma_W = 0$ near -3 , put $g_\mu = K_1(\mu\sigma_W(s))$ on Y_0 , extended by 1, and define g'_μ on Y'_0 likewise. For $\mu \in [0, 1]$ these gauge transformations equal 1 on $[-9/4, 0] \times \Sigma$, hence near the neck and the seam, and $g_1^*B_1$ is a product on $[-5/2, -2] \times \Sigma$. Apply the gauge transformation $g_{\mu\chi(r)} \cup 1 \cup g'_{\mu\chi(r)}$ on the gauge region, with χ as above. It changes the mixed limits to the stationary pairs of $g_\mu^*B_1$ and $g'_\mu^*B'_1$, and at $\mu = 1$ the region $r \geq -4$ carries pr^*c with $c = g_1^*B_1 \cup \text{pr}^*b \cup g_1'^*B'_1$, a representative of the critical point over x in the sense above. The configuration therefore lies in $\mathcal{T}_x$. A member of $\mathcal{T}_x$ is carried to the constant pair of $g_\mu^*B \cup \text{pr}^*b \cup g'_\mu^*B'$, which is again a representative since $g_\mu = 1$ on $[-5/2, 0]$ in that case; so members of $\mathcal{T}_x$ remain in $\mathcal{T}_x$.

Proof of Proposition 2.2(a) for finite torus lengths. Concatenate Lemma B.1, Lemma B.3, the straightening and the collar widening, and transport by the identifications of Section B.1. Relative preparation of the original tails is the second step of Lemma B.1. Continuity in the data holds because every choice above (the collar gauges, the charts at stationary pairs, the lifts of Lemma B.2, the horizontal lifts, the widening gauges) depends continuously on the primary data, and the identifications depend continuously on the lengths. The deformation is the identity on, or a translation of, the gauge input end away from $[-5, -3] \times Y_\#$. Hence if the family has the form $A_0 \#_{\ell} c$ with an ordinary configuration A_0 glued at the input with a long neck ℓ , the deformation is $A_0 \#_{\ell+\tau} c_\tau$, and the determinant gluing maps of Lemma 3.3 are continuous along it. This is the coherence with ordinary gluing at the gauge input, which is the coherence asserted in Proposition 2.2. A strip glued at the output is not carried along by the deformation, since preparation truncates the output tail. A configuration with a strip glued at its output is treated as a configuration in its own right, with the new output label, and the output coherence of the orientations is proved by the two-stage shift in the proof of Proposition 5.4. $\square$

B.4. Completed torus domains. Let a compact family on $W_{a,\infty}$ be given. By Lemma 8.4(i) it is in torus normal form beyond some r_0 after a continuous family of small gauge transformations. The truncation $\alpha + (1 - t + t\chi_R)b$ on the tail, with $\chi_R = 1$ for $r \leq R$ and 0 for $r \geq R + 1$, is continuous in the topology of Section 2, whose weight is independent of r , and fixes members that are α beyond R ; at $t = 1$ every member is α beyond $R + 1$. The mixed-end limits are not changed: in torus normal form they are already α beyond r_0 , because a flat connection in the torus-fibre slice of (I12) which is C^1 -close to α equals α . Indeed its fibre part is flat, coclosed and small, hence zero since $H^1(T; \text{ad } \alpha) = 0$, and then $d_\alpha \phi_j = 0$ forces $\phi_j = 0$ since $H^0(T; \text{ad } \alpha) = 0$. Such a configuration is the extension by α of a configuration on $W_{a,R+2}$ which equals α near the flat face, and the face is part of the matching seam. The deformations of Sections B.2–B.3 preserve this condition: the collar gauges and the chart cutoffs are the identity where the configuration and its limit are α , the widening gauges are supported away from the neck and the seam, the collar pullbacks fix configurations equal to α near the seam, the seam term $\tilde{\psi}$ vanishes on the torus component, the horizontal lifts and the representatives are α on T , and the translation and the straightening bring the positive extension, which is α on T , to the neighbourhood of the face. Extending by α along the whole deformation gives Proposition 2.2(a) on $W_{a,\infty}$, ending in $\mathcal{T}_x$ with a representative which is α on T ; the constant pair is then α on the tail.

B.5. Shifting on domains with two gauge inputs. On a triangle or folded triangle at positive collapse scale, and on a finite-width joined torus domain, the outgoing closed cut writes the gauge region as an ordinary cobordism followed by an output region M . For the joined torus domain this is (I2). For the triangles and folded triangles the cut is the closed

three-manifold that separates the two gauge inputs from the output and crosses the collapsing region (Paper I, Lemma 9.11); a change of metric at fixed configuration gives it a product collar (I11). In each case M , with a smooth product collar at the cut, is identified with the complement of $(-\infty, -c_1] \times Y_\#$ in a special quintuple by a diffeomorphism of decorated manifolds, continuous in the width and in the scale on compact positive ranges, which carries the product collar of M onto $[-c_1, -3] \times Y_\#$ after a reparametrization of its normal coordinate. Such a reparametrization does not change product configurations pr^*c , and the metric on this collar may be replaced by one with a longer product collar along a path of data, since the configuration spaces do not depend on the metric (Section B.1). Its compression bodies are those of the pair (D_0, D_1) , or of the marked pair of the triangle. The cut is $\{-c_1\} \times Y_\#$, with $c_1 > 5$. The steps of Sections B.2–B.3 are supported in $r \geq -5$, away from a collar of the cut, except for the translation. Replace the translation by the localized map: with χ nondecreasing, zero near $-c_1$ and one on $[-5, \infty)$, put $F_\tau(r) = r + \tau\chi(r)$ and

$$A_\tau = F_\tau^* \tilde{A}, \quad u_\tau(r, t) = u(r + \tau, t).$$

Since $F'_\tau > 0$ this is a smooth configuration; it is a translation on $[-5, \infty)$, so the seam and the corners are treated as in Lemma B.3, and near the cut it is the identity, so it glues with the ordinary cobordism. For $\tau = S + 4$ followed by the straightening and the collar widening, the region $r \geq -4$ carries a constant pair C^c , and the configuration is $A_E \# C^c$, with A_E an ordinary cobordism configuration ending at c on the parallel copy $\{-7/2\} \times Y_\#$ of the cut, whose collar $[-4, -3] \times Y_\#$ carries pr^*c , and with the original inputs unchanged. The configuration $A_E \# C^c$ is not a product of two gauge paths. The identification above carries the product collar of the cut onto $[-c_1, -3] \times Y_\#$, so the copy $\{-7/2\} \times Y_\#$ of the cut and the collar $[-4, -3] \times Y_\#$, on which the configuration is pr^*c , lie in the product collar, as ordinary gluing at the copy requires.

Pull back the transition functions of the bundle as well. If $C_{i,\tau}$ is the parameter component of the pulled-back connection, solving $\partial_\tau v_i = -C_{i,\tau} v_i$ identifies its transitions with the initial ones. This retains the lifting classes and central cocycles. The stationary representatives are the paths of perturbed-flat mixed-end limits of the configuration, including the corner pullback, and both operator rows are transported. A common sufficiently small weight gives the family in Proposition 2.1(a); the intermediate operators need not be invertible.

This is the localized shift used in Definition 5.5, and it proves Proposition 2.2(b). Relative and continuity statements follow as in Lemma B.3.

APPENDIX C. A SIGN CALCULATION FOR PROJECTIVE COMPOSITION

For strictly commuting maps F and G of degrees f and g , the closed maps $F^r = F \text{gr}^f$ and $G^r = G \text{gr}^g$ of Section 3 compose by (3.5), since $\text{gr}^f G = (-1)^{fg} G \text{gr}^f$. Thus passing from F to F^r is compatible with composition after projectivization. The factor $(-1)^{fg}$ is scalar, although the grading operator is applied to each map.

For a strict homotopy $F - F' = dH + Hd$, where $|H| = f + 1$, the correctly converted homotopy is $H \text{gr}^f$. Indeed

$$\delta(H \text{gr}^f) = dH \text{gr}^f - (-1)^{f+1} H \text{gr}^f d = (dH + Hd) \text{gr}^f.$$

Using $H \text{gr}^{f+1}$ instead would give the wrong sign. This calculation is used when a gauge composition homotopy is translated into the Hom complex convention.

Lastly, if $F_i - G_i = \delta K_i$ and the maps are closed, a homotopy between two composites is obtained by replacing one factor at a time and using (3.2). For example, for three factors

one may take

$$K_3 F_2 F_1 + (-1)^{|G_3|} G_3 K_2 F_1 + (-1)^{|G_3|+|G_2|} G_3 G_2 K_1.$$

Applying δ telescopes to $F_3 F_2 F_1 - G_3 G_2 G_1$. This is the chain-level algebra used in Lemma 14.1.

APPENDIX D. ORDINARY INTEGRAL EXCISION

This appendix contains the ordinary gauge theory used in Section 8, with the geometry of Section 2.5: the almost-complex structures at the closed cut (Section D.1), lifting groups and gluing multiplicities (Section D.2), the proof of Proposition 2.3 (Section D.3), and the integral form of the excision theorem of Kronheimer and Mrowka (Proposition D.3). The compactification and gluing arguments are those of [KM11b, Section 3.9] with empty singular locus; their definition of a cut permits a nonseparating cut, and the odd torus satisfies the non-integral condition required there. The orientation of a self-glued count is [KM11b, Proposition 5.5], whose sign is fixed by the argument of [KM10, Lemma 2.4], and the comparison of the composite of excision and reverse excision with a product cobordism is the argument of [KM11b, Theorem 5.6]. We give in full the three points listed at the beginning of Section 8, the choices specific to the present cut data, and the lifting-group counts needed for a statement at chain level; for the remaining steps we refer to these sources and to (I10).

D.1. Almost-complex structures at the closed cut. Let $W_k : M \rightarrow M^+$ be the cobordism obtained by adding k one-handles, where $M^+ = M \#^k (S^1 \times S^2)$. The folded triangle for this compression has the closed gauge cut $Z \cong (-M) \# M^+ \# T^3$. The odd tori $T_a, T_b \subset Z$ are homologous, the lifting group at this cut is generated by $s_Z = \text{PD}[T_a] \pmod{2}$, and excision along these tori gives $Z' = (-M)^\# \sqcup (M^+)^\#$. Write $\mathbf{E} : C_G(Z) \rightarrow C_G(Z')$ for the gauge excision map and $\bar{\mathbf{E}}$ for its reverse; the tensor complex at Z' uses the right convention.

The almost-complex structure induced by the folded triangle at Z need not equal the one induced by the mixed torus domain, so a conversion at this interior cut is required. We choose the structures so that

$$(D.1) \quad c_1(J)[T_a] \equiv 0 \pmod{4}.$$

For two such structures the lifting-class sign (4.1) of the conversion is trivial on s_Z , so the conversion at Z is defined on the quotient complex.

The filling. In the handle filling of Paper I, Section 11.1,

$$B_k = (W_k \setminus \text{int } \nu(\gamma)) \natural (T^2 \times D^2),$$

the tori T_a, T_b are homologous to the torus factor of the cap. Choose the product almost-complex structure on the cap. Its first Chern class vanishes on the torus factor, hence every extension of it over B_k vanishes on T_a and T_b . The boundary sum is along a three-ball, so Mayer–Vietoris makes $H^2(B_k; \mathbb{Z}) \rightarrow H^2(T^2 \times D^2; \mathbb{Z})$ surjective, and the relative obstruction argument of Lemma 4.1, with the cap as the prescribed region, extends the product structure: the primary obstruction has zero connecting class and the secondary relative group in degree four vanishes. The other end restrictions remain free.

The excision cobordisms. Take the product germ on the region $A = T^2 \times U$ of $\bar{\mathcal{E}}$ and $\mathcal{E}$, where U is the disk polygon of Paper I, Section 10; the attachments are away from the fibres at T_a, T_b . The complement of A is $[0, 1] \times Q$, where $Q = Q_- \sqcup Q_+$, $Q_\pm = (T^2 \times I_\pm) \# M_\pm$, $M_- = -M$ and $M_+ = M^+$. By convention the separate torus domains $W_{\pm, R}$ read their input lines $L_G(a_\pm)$ in the reference almost-complex structures of $(-M)^\#$ and $(M^+)^\#$; this is unobstructed, since their gauge regions deformation retract onto the input slice. Choose

$J_{\bar{\mathcal{E}}}$ with these reference structures as its input restriction at Z' ; the two prescriptions agree on the collars of the tori of Z' , since we take these tori to be the tori T_0 of the two objects, where the object germ is the product germ. Then choose $J_{\mathcal{E}}$ with restriction $J_{\bar{\mathcal{E}}}|_Z$ at Z . By excision,

$$H^*(\bar{\mathcal{E}}, A \cup \partial_{\text{in}} \bar{\mathcal{E}}; \mathbb{Z}) \cong H^*((Q, \partial Q) \times ([0, 1], \{0\}); \mathbb{Z}) = 0,$$

and likewise $H^*(\mathcal{E}, A \cup Z; \mathbb{Z}) = 0$, so the primary and secondary relative obstructions to both extensions vanish. Make both structures cylindrical at the closed cut, keeping the product germ near the fibres at T_a, T_b . Their Chern evaluations on these fibres are then zero, which is stronger than (D.1). With these choices the orientation maps of Section 8.2 and the maps $N_{\pm}$ read their input lines in the same structures, and the composite $\bar{\mathbb{E}}\mathbb{E}$ has equal end structures at Z .

D.2. Lifting groups and gluing multiplicities. We perform the gauge excision calculation before passing to a quotient. Let $\mathcal{G}_0(X)$ be the subgroup of bundle automorphisms with trivial lifting class. If a finite group ϕ_X acts freely and orientation preservingly on the relevant zero-dimensional solution set, its quotient count is

$$(D.2) \quad I(X)^{\phi_X}(\bar{a}_0, \bar{a}_1) = \frac{1}{|\phi_X|} \sum_{a_i \in \bar{a}_i} \#\mathcal{M}(X; a_0, a_1).$$

The numerator is divisible by $|\phi_X|$: it is a sum of free orbits, each with a constant sign.

Lemma D.1. *Let X be connected and let Y be a closed hypersurface with cut manifold X_Y . Suppose the critical limits on Y have trivial stabilizer in the adjoint determinant-one gauge group. For each matched configuration on X_Y that glues, the resulting $\mathcal{G}_0(X)$ -classes form a torsor for*

$$K_Y = \ker(H^1(X; \mathbb{Z}/2) \longrightarrow H^1(X_Y; \mathbb{Z}/2)).$$

In particular their multiplicity is $|K_Y|$.

Proof. Choose representatives whose two limits agree under the fixed gluing identification. Two glued classes whose restrictions to X_Y agree as determinant-one classes differ by an automorphism whose lifting class restricts to zero on X_Y , and conversely every class of K_Y changes the gluing while leaving the cut class fixed. These classes are realized by gauge transformations: for a component Y_j of Y choose a unit section n of the adjoint bundle, which exists because the Euler class of an oriented rank-three bundle is 2-torsion and $H^3(Y_j; \mathbb{Z})$ is torsion-free, and on its normal collar use the adjoint of $\cos(\pi\rho) + \sin(\pi\rho)n$, $0 \leq \rho \leq 1$, with ρ constant near the endpoints. This is the identity at both ends of the collar, while its lift runs from $+1$ to -1 , so its lifting class is $\text{PD}[Y_j]$. The stabilizer condition makes the action free. It is transitive by comparison of the two gluing identifications: the even lifts on the components of X_Y agree at paired ends up to central signs, and changing a lift by the centre on one component changes the incident signs by the coboundary of the cut graph, so the remaining choices form the first cohomology of the cut graph over $\mathbb{F}_2$, which Mayer–Vietoris identifies with K_Y , the image of the connecting map from $H^0(Y; \mathbb{Z}/2)$. The matching and analytic assertions needed for the excision cobordisms are part of Proposition 2.3. $\square$

A finite lifting group acts freely on the solutions in the present calculation. A fixed solution would give a parallel automorphism; its restriction to an irreducible end is the identity once its lifting class there vanishes, and it is then the identity on the whole connected cobordism. An element whose class is nonzero at every end, such as the classes $P_{\pm}$ of Section D.4, restricts to s_Z at the Z ends, and ϕ_Z acts freely on the critical points of Z , so it fixes no solution either. The almost-complex orientations are preserved by the whole group, so (D.2) applies.

D.3. Self-gluing along the torus cut. Write $\mathcal{E}$ for the cobordism defining E and put $V = \mathcal{E} \cup_{Z'} \bar{\mathcal{E}}$. The local polygon description gives a nonseparating cut $\Sigma = T^2 \times \gamma$ of V ; its complement V° is connected and retains the original Z ends, and the bundle is odd on the torus fibre. The other composite $\bar{\mathcal{E}} \cup_Z \mathcal{E}$ has a separating torus cut. All components of the cobordisms that we regularize have an original cylindrical end; the caps are kept fixed. On each new three-torus end we use a fixed flat product metric and the unperturbed equation. The odd bundle on the three-torus has two determinant-one flat critical points g_1, g_2 , both acyclic and interchanged by $\text{PD}[T^2 \times \{\text{point}\}]$. Their relative grading is four modulo eight, by the proof of [KM11b, Proposition 5.2(a)]; see also [KM11b, Lemma 5.3]. Hence they have the same parity, which we denote by $|g|$. The caps $T^2 \times D^2$ of Section D.4 carry the same unperturbed end data and no secondary perturbation.

In the trace the traced factor is placed last. For a map F of degree f on a tensor end with torus basis g_j and diagonal blocks F_{jj} , the right tensor convention gives

$$(D.3) \quad \text{Tr}^r(F) = \sum_j (-1)^{(f+1)|g_j|} F_{jj}.$$

Both torus generators have parity $|g|$, so the comparison with caps of coefficients $u, v \in \{1, -1\}$ carries the factor $(-1)^{(f+1)|g|}/(uv)$, up to one fixed overall orientation. The gluing multiplicities are those of Lemma D.1: two for the nonseparating cut Σ and for the two-edge composition cycle at Z' , and one for a separating cut or a tree of caps.

Proof of Proposition 2.3. The three-torus ends. Keep the entire ordinary end functional $\text{CS}_\mathfrak{h} + f$ of Paper I fixed on the original ends, including its trajectory perturbation; only the secondary cobordism perturbations vary. On the new three-torus ends the data are the fixed flat product metric and the unperturbed equation throughout all comparison families, and the lifting action on $\{g_1, g_2\}$ is free. Every nonconstant unperturbed torus trajectory has index at least four: the two end Hessians are conjugate by an adjoint gauge transformation, so their spectral correction terms agree, and the cylinder index and energy are 8κ and $8\pi^2\kappa$ with $\kappa \in \frac{1}{2}\mathbb{Z}$. The compactifications used have total dimension at most two, so after joint regularization of the cut pieces a torus trajectory would leave a negative index for the remaining pieces, counting the data parameter once. Hence no torus trajectory perturbation is needed in this range, and regularity of the unperturbed torus trajectories is never used.

Collar-supported secondary perturbations. These are the secondary perturbations of [KM11b, Sections 3.7 and 3.9], with compact time profiles. Choose a product collar joining each component to an autonomous end, with fixed metric also in parameter families, and a compact interval J in it. Let F_Y be the image on the collar cross-section Y of the allowed lifting group, and put

$$(D.4) \quad q_s(B) = \sum_\ell c_\ell \psi_\ell(s) \mathfrak{f}_\ell(B), \quad \text{supp } \psi_\ell \subset J,$$

where the $\mathfrak{f}_\ell$ are F_Y -invariant holonomy cylinder functions on Y whose loops lie in a single slice; elements of the kernel of the restriction to F_Y impose no further condition. A countable dense family of profiles and cylinder functions, with a weighted Banach completion controlling all required norms, gives the perturbation space. The holonomy estimates of [KM11a, Propositions 3.7 and 3.15] persist, time derivatives multiplying the profile bounds. Including the fixed primary cutoff terms in q_s , one has $\frac{d}{ds}(\text{CS} + q_s)(B(s)) = -\|\dot{B}(s)\|^2 + \partial_s q_s(B(s))$, and the last term integrates to at most the L^1 norm of the time derivatives of the coefficients, uniformly under stretching. No perturbation depends on connections on both sides of a cut.

Invariant regularity. Regularity is the transversality of [KM11b, Propositions 3.8 and 3.10] for secondary perturbations, carried out invariantly under the lifting groups as in (I10). Three facts are specific to this class. First, *slice continuation*. On a compact slab of a collar, a difference of two solutions, or an adjoint field, satisfies $(\partial_s + H(s) + K(s))u = 0$, where H is formally self-adjoint on a common smooth domain, H' and K are bounded on L^2 , and $H = \pm * d_{B(s)}$ for the one-form equation; a difference of two holonomy terms has this form by integrating its derivative along the segment between the connections. For a concave weight φ with $-\varphi'' \geq c > 0$ and large λ , expanding $\|(\partial_s + H + \lambda\varphi')v\|^2$ for compactly supported v and integrating by parts gives

$$(D.5) \quad (c\lambda - C)\|e^{-\lambda\varphi}u\|^2 \leq 2\|e^{-\lambda\varphi}(\partial_s + H + K)u\|^2.$$

With φ increasing on a short slab, a solution vanishing before a slice is cut off farther into the slab; the cutoff error has strictly larger φ than a neighbourhood of the slice, and $\lambda \rightarrow \infty$ gives vanishing in that neighbourhood. Reversing s gives the other direction, and vanishing on a slice suffices, since the equation then makes all time derivatives vanish there. Second, *freeness and invariant density*. If an allowed automorphism fixes a slice of a solution, slice continuation and invariance make it fix the whole end and its limit; freeness of F_Y there forces its lifting class to vanish, and irreducibility makes its even lift central. So each collar slice has a free F_Y -orbit, and prescribing equivariant first jets on a small finite-dimensional slice and its disjoint translates, approximating by cylinder functions as in the proof of [KM11a, Proposition 3.10] and averaging gives invariant cylinder functions detecting every direction transverse to gauge; freeness ensures that the prescribed derivative survives the averaging. Third, *the cokernel*. A cokernel element (ξ, b) of (d_A^*, L_A) has $\xi = 0$ by gauge equivariance and the L^2 end condition. On the collar write b by a one-form $y(s)$ with $d_{B(s)}^*y = 0$. If the element annihilates every secondary variation, time localization, invariant density and the Coulomb condition give $y = 0$ on J ; then (D.5), local unique continuation on the unperturbed core and slice continuation on the other collars give $b = 0$, since all nonlocal terms lie in the product end collars.

Sard–Smale then gives invariant regular data. For a broken configuration with one parameter b , the joint linearization $((\xi_i)_i, \dot{b}) \mapsto (D_i\xi_i + \partial_b\mathcal{F}_i\dot{b})_i$, in which the shared parameter occurs once, is onto because the variations on the pieces are independent. Choosing regular data on fixed faces first, extending over their collars by variations supported away from them, and intersecting the residual sets for the finitely many cut patterns and the countably many relative classes gives the relative regular families of [KM11b, Proposition 3.10 and Section 3.9] in the invariant class; fixed regular caps require no variation. On a bounded neighbourhood of data the energy identity, the index-energy relation and the class calculation below make the index ranges in question finite. Regularity is open on the finite union of the low-dimensional compactifications, which have regular gluing neighbourhoods at their broken strata: nonregular solutions for convergent data would converge to a regular point or to a regular broken configuration, contradicting the local estimate. Finite sums in (D.4) are dense, also for smooth coefficient families on a compact base, so the final data may be taken to be finite sums; approximating within this open set, relative to the prescribed face collars, preserves the original end equations and all supports.

Matching, compactness and gluing. An even gauge transformation is a section of the bundle with fibre $SU(2)$ and conjugation action. The fibre is 2-connected and conjugation acts trivially on π_3 , so extending prescribed values on boundary components B_0 of a compact cobordism W is obstructed only in $H^4(W, B_0; \mathbb{Z}) \cong H_0(W, \partial W \setminus B_0; \mathbb{Z})$, which vanishes when every component has an unprescribed boundary end. Applied to the two new boundaries

of V° , with its original Z ends free, equal critical classes are represented by equal connections under the gluing identification and glue on the specified bundle; for capping, the cap representatives are fixed. Relative homotopy classes with fixed end classes form a $\mathbb{Z}$ -torsor whose generator changes the index by eight, each represented on the fixed bundle [KM11b, Sections 3.2 and 3.7], so a fixed index and tuple allow at most one class. The critical sets are finite, and in a regular family each component index is bounded below by minus the base dimension, so at fixed total dimension only finitely many index allocations occur, the shared parameter being counted once.

Gluing and compactness are those of [KM11b, Section 3.9]. The point specific to self-gluing is that a regular index-zero cut solution A has a single invertible operator D_A on the connected cut piece, including both stubs, whose inverse is patched with the uniform inverse of $\partial_s + B_\beta$ on the cylinder, B_β the invertible extended Hessian at the limit β . Compactly supported secondary terms and slice-local primary terms commute with the cutoffs, and stabilization gives the same charts and determinant maps for families that are regular as parametrized families. In the compactness argument the index exclusion above replaces regularity of trajectories on the new torus ends, also in the two-parameter corner families. In total dimension at most one the cut solutions form the broken-metric boundary; further breaking has higher codimension and bubbling is excluded by low-dimensional compactness. The local gluing charts give its collars, and finiteness over relative classes gives one neck threshold. For disjoint cuts we use the multi-parameter corner model: a simultaneous stretching of two components is the diagonal in the two gluing parameters, so the two-edge composition cycle is one separating gluing followed by one self-gluing. The cuts Σ and Z' intersect, so their comparison joins finite stretching parameters by a relative regular path. The gluing fibres are those of Lemma D.1; the action is free, as shown after it, and preserves the almost-complex orientations.

Trace signs and homotopies. Use the induced almost-complex structures at the ends, with the orientation torsors of [KM11b, Section 5.1], and apply [KM11b, Proposition 5.5] with the enlarged group exactly $K_\Sigma = \langle \text{PD}[\Sigma] \rangle$. This group restricts trivially to V° , to Σ and to the original ends, so the trace uses the full determinant-one complex on Σ , and the free orientation-preserving action recovers the raw count as twice that trace. The larger quotient is taken only after this step. The alternating sign is the analytic determinant sign given by the argument of [KM10, Lemma 2.4], the parity of spectral flow in the periodic model, so there is no independent sign choice for each generator or relative class. For $f = 0$, (D.3) is this alternating trace; moving the traced line past a map of degree f supplies the factor $(-1)^{f|g_j|}$, and conjugation by the tensor isomorphism T of Section 3 gives the same formula algebraically. The trace commutes with the Hom differential and satisfies

$$\text{Tr}^r((\Lambda_0 \otimes^r 1)F(\Lambda_1 \otimes^r 1)) = \Lambda_0 \text{Tr}^r(F)\Lambda_1,$$

so the external conversion maps can be inserted with their degrees; interior conversions follow Convention 8.1.

Finally use the parameter-first orientation of Section 5, with outward-normal-first boundary orientations. The source convention in [KM11b, Section 3.9] is fibre-first. For an r -parameter family the permutation of the determinant factors contributes $(-1)^{r \text{ ind } D}$. For a zero-dimensional one-parameter count the vertical index is -1 , so this changes the sign of the count; for the one-dimensional parametrized moduli space the vertical index is zero, so its orientation and endpoint maps are unchanged; for two parameters the permutation sign is one, and boundary faces use the one-parameter rule. It remains to include the generator-order conversion. At the induced end gradings the unconverted cobordism maps F_i are

even. Let d_0 and H_0 be the source differential and its fibre-first homotopy count, so that $F_1 - F_0 = d_0 H_0 + H_0 d_0$. Placing an odd trajectory determinant after its input generator changes d_0 to $d = d_0 \text{gr}$ and H_0 to $H_0 \text{gr}$, and the parameter-first change contributes one further minus sign. The homotopy operator in our convention is therefore

$$K = -H_0 \text{gr}, \quad dK + Kd = d_0 H_0 + H_0 d_0 = F_1 - F_0,$$

where gr is the grading operator on the appropriate source complex and d_0, H_0 are odd. On disconnected ends the whole ordered input generator orientation line gives the tensor rule (3.3), so the parametrized boundary formula agrees with (5.3). Arbitrary external conversions, including odd ones, are governed by the Hom Leibniz rule. All these choices are uniform across matrix entries, which proves the three assertions in the stated orientation convention. $\square$

D.4. Excision and reverse excision are inverse over the integers. This is the argument of [KM11b, Theorem 5.6], carried out at chain level with the lifting groups retained. The composite $V = \mathcal{E} \cup_{Z'} \bar{\mathcal{E}}$ contains, in its local model T^2 times the doubled excision polygon, the nonseparating hypersurface $\Sigma = T^2 \times \gamma$. Cut along Σ to obtain V° and cap the two new boundary components by $T^2 \times D^2$; the capped cobordism V' is a product on Z . Give V the structure J_V which is the product structure on T^2 times the doubled polygon and which, on the complementary region $[0, 2] \times Q$, is translation invariant and equal to $J_{\mathcal{E}}|_Z = J_{\bar{\mathcal{E}}}|_Z$. The two agree along $[0, 2] \times \partial Q$, because this structure has the product germ near T_a and T_b . Thus J_V is cylindrical with the product germ near Σ , and its restrictions to the two ends Z are those of $\mathcal{E}$ and $\bar{\mathcal{E}}$. The trace of Proposition 2.3(b) is taken with $J_V|_{V^\circ}$, and V' carries $J_V|_{V^\circ}$ together with the product structure on the caps. For the other composite the glued structure $J_{\bar{\mathcal{E}}} \cup_Z J_{\mathcal{E}}$ already has this form near its torus cut.

The class $a = \text{PD}_\Sigma[T^2]$ interchanges the two critical points g_1, g_2 of Σ . They must be kept separate when the caps are attached: the caps carry the trivial lifting group, since a does not extend across their meridional disks. Write $s = \text{PD}_Z[T_a] \pmod{2}$, let $\hat{s} \in H^1(V^\circ; \mathbb{Z}/2)$ be its pullback, and let m_1, m_2 be the meridional classes of the two removed cap neighbourhoods. The restrictions are

	Z_{in}	Z_{out}	Σ_{in}	Σ_{out}
$\hat{s}$	s	s	0	0
m_1	s	0	a	0
m_2	s	0	0	a
$m_1 + m_2$	0	0	a	a

The group ϕ_V has order 8: the two sides of the excision polygon give classes P_+, P_- and the disconnected intermediate cut gives a class c , with restrictions $\hat{s}, \hat{s}, m_1 + m_2$ to V° ; they are independent on V , as loops crossing the two polygon sides and the intermediate cut show. Its image on V° is $\langle \hat{s}, m_1 + m_2 \rangle$ of order 4, and its kernel is K_Σ of order 2, generated by $P_+ + P_-$, which is dual to Σ and nonzero because Σ is nonseparating. The capped product has group $\phi_{V'} = \langle \hat{s} \rangle$ of order 2.

The group $H^\circ = \langle \hat{s}, m_1, m_2 \rangle$ has order 8; it contains the restrictions to V° of both ϕ_V and $\phi_{V'}$, but is not itself the pullback of either, and it acts freely by the argument following Lemma D.1. Choose regularizing data on V° invariant under H° and put $D(\alpha, \sigma; \beta, \tau) = \# \mathcal{M}(V^\circ; (\alpha, \sigma), (\beta, \tau))$, $\sigma, \tau \in \{g_1, g_2\}$, and $\Delta(\alpha, \beta) = D(\alpha, g_1; \beta, g_1)$. The action of $m_1 + m_2$ makes the two diagonal blocks equal, and the action of m_2 identifies the off-diagonal block with $\Delta(s\alpha, \beta)$. Thus, after summing over endpoint lifts, diagonal and off-diagonal blocks have the same sum.

The cap count for the unperturbed flat product data on the three-torus ends. Complete $T^2 \times D^2$ by a product cylindrical end to $T^2 \times S$, with a flat metric on T^2 and a product metric, and let A_0 be the pullback of the odd flat connection on T^2 . Its adjoint local system on T^2 splits into the three nontrivial real sign local systems and is acyclic, so the torus Hodge Laplacian, and with it the product Hodge Laplacian on $T^2 \times S$, has a positive lower bound in every degree. For $D_0 = d_{A_0}^* \oplus \sqrt{2}d_{A_0}^+$ the identities $D_0^* D_0 = \Delta_{A_0,1}$ and $D_0 D_0^* = \Delta_{A_0,0} \oplus \Delta_{A_0,2}^+$ show that the flat cap has index zero and is regular. Call its boundary lift g_1 . Since g_2 has relative grading four modulo eight, while a change of relative class at a fixed lift changes the cap index by eight, an index-zero cap approaches only g_1 , in the zero-charge class of A_0 , where its energy vanishes and it is flat. The flat odd torus connection has a unique flat extension across the disk, because the meridian is contractible and $\pi_1(T^2 \times S) = \pi_1(T^2)$, and only g_1 admits an extending meridional determinant-one lift. Hence there is exactly one determinant-one cap class; its signed coefficient is a unit, and the reversed cap gives the dual unit. This calculation is at the full two-generator torus complex, before any lifting-group quotient: the rank-one quotient would not distinguish matching two lifts from matching two orbits.

Proposition D.2. *For the regular data of Proposition 2.3, the quotient counts for V and V' have the same absolute multiplicity. They agree up to one sign, the same for all entries.*

Proof. For V , self-gluing has multiplicity 2 and there are two equal diagonal blocks; quotienting by the group of order 8 gives $\frac{2 \cdot 2}{8} \sum_{\alpha \in \bar{a}, \beta \in \bar{b}} \Delta(\alpha, \beta)$. For V' , the caps have unit coefficients, their separating gluings have multiplicity 1, and the group has order 2, which gives $\frac{1}{2} \sum_{\alpha \in \bar{a}, \beta \in \bar{b}} \Delta(\alpha, \beta)$. If the second cap selects g_2 , the symmetry of the off-diagonal block gives the same sum. The orientation signs are those of Proposition 2.3; since the unconverted map of V° is even, they reduce to the factor $(-1)^{|g|}/(uv)$ for every entry. All group actions preserve the orientations, so no cancellation occurs inside an orbit, and the two numerical factors agree. These fractions are integer orbit counts, the division by the group order being carried out on integer counts at chain level. $\square$

Proposition D.3 (Integral excision). *With the end structures of Section D.1, form the composite $\bar{E}E$ with the right-placed conversion at Z' . It is chain homotopic to the identity of $C_G(Z)$ up to one sign. The maps $H(E)$ and $H(\bar{E})$ are isomorphisms, and E and $\bar{E}$ are chain homotopy equivalences, after shifting degrees if necessary.*

Proof. The one-dimensional gluing family identifies the counts of Proposition D.2 with the composite excision map, and the same argument for the other composite gives the reverse comparison. There the cut is the separating torus cut $T^2 \times \gamma'$, so there is no trace and the matching multiplicity is one; the lifting group has order eight, against the capped group $\langle s_-, s_+ \rangle$ of order four. On each cut piece the meridional class of the removed cap exchanges $\sigma \in \{g_1, g_2\}$ and moves one endpoint lift within its orbit, so after summing over endpoint lifts the two matched choices contribute equally, a factor two, and the factors balance: $2/8 = 1/4$. Consequently $H(E)$ and $H(\bar{E})$ are isomorphisms.

With the end structures of Section D.1, E and $\bar{E}$ have the same structure at Z , and at Z' the composite is formed with the right-placed conversion, as at the reference structures of the objects. By [KM11b, Proposition 5.4] the composite is the map of V ; its hypothesis holds because the class c dual to a component of Z' lies in ϕ_V . In terms of counts, by Proposition 2.3(a), the gluing multiplicity two at Z' is absorbed by $|\phi_V| = 8$, and the identity holds at chain level, up to an integral chain homotopy, by Proposition 2.3(c). By the independence of the almost-complex structure in Proposition 4.4, the map of V may be computed

for any almost-complex structure on V as in Lemma 4.1, after the end conversions, up to one sign for all entries. The unconverted map of V° is even, so $f = 0$ and the capped factor is $(-1)^{|g|}/(uv)$; the capped cobordism V' is a product on Z , so its map is the identity class up to one sign by Proposition 4.4. Hence the composite is chain homotopic to the identity up to one sign, and in particular it is ± 1 on every homogeneous class. By Lemma 3.1, the closed excision maps are homotopy equivalences, after shifting degrees if necessary. $\square$

The first assertion of Proposition D.3 is the only inverse-excision statement used in this paper; it enters through (8.17), on which Section 9 relies.

REFERENCES

- [Ati88] M. F. Atiyah, *New invariants of 3- and 4-dimensional manifolds*, in: The Mathematical Heritage of Hermann Weyl, Proc. Sympos. Pure Math. **48**, Amer. Math. Soc., Providence, RI, 1988, 285–299.
- [DFL21] A. Daemi, K. Fukaya and M. Lipyanskiy, *Lagrangians, $SO(3)$ -instantons and the Atiyah–Floer conjecture*, arXiv:2109.07032.
- [DFL24] A. Daemi, K. Fukaya and M. Lipyanskiy, *Lagrangians, $SO(3)$ -instantons and mixed equation*, Geom. Funct. Anal. **34** (2024), 659–732; arXiv:2109.07038.
- [Don87] S. K. Donaldson, *The orientation of Yang–Mills moduli spaces and 4-manifold topology*, J. Differential Geom. **26** (1987), 397–428.
- [Don02] S. K. Donaldson, *Floer Homology Groups in Yang–Mills Theory*, Cambridge Tracts in Math. **147**, Cambridge Univ. Press, 2002.
- [Flo88a] A. Floer, *An instanton-invariant for 3-manifolds*, Comm. Math. Phys. **118** (1988), 215–240.
- [Flo88b] A. Floer, *Morse theory for Lagrangian intersections*, J. Differential Geom. **28** (1988), 513–547.
- [FH93] A. Floer and H. Hofer, *Coherent orientations for periodic orbit problems in symplectic geometry*, Math. Z. **212** (1993), 13–38.
- [FOOO09] K. Fukaya, Y.-G. Oh, H. Ohta and K. Ono, *Lagrangian Intersection Floer Theory: Anomaly and Obstruction, Part II*, AMS/IP Stud. Adv. Math. **46.2**, Amer. Math. Soc., Providence, RI, and International Press, Somerville, MA, 2009.
- [FOOO24] K. Fukaya, Y.-G. Oh, H. Ohta and K. Ono, *Exponential decay estimates and smoothness of the moduli space of pseudoholomorphic curves*, Mem. Amer. Math. Soc. **299** (2024), no. 1500; arXiv:1603.07026.
- [KM11a] P. B. Kronheimer and T. S. Mrowka, *Knot homology groups from instantons*, J. Topol. **4** (2011), 835–918; arXiv:0806.1053v2.
- [KM11b] P. B. Kronheimer and T. S. Mrowka, *Khovanov homology is an unknot-detector*, Publ. Math. Inst. Hautes Études Sci. **113** (2011), 97–208.
- [KM10] P. B. Kronheimer and T. S. Mrowka, *Knots, sutures, and excision*, J. Differential Geom. **84** (2010), 301–364.
- [KL25] R. Kupferman and R. Leder, *Elliptic pre-complexes, Hodge-like decompositions and overdetermined boundary-value problems*, Forum Math. Sigma **13** (2025), e51; arXiv:2304.08977.
- [Mau09] S. Ma’u, *Gluing pseudoholomorphic quilted disks*, arXiv:0909.3339.
- [MWW18] S. Ma’u, K. Wehrheim and C. T. Woodward, *A_∞ functors for Lagrangian correspondences*, Selecta Math. (N.S.) **24** (2018), 1913–2002; arXiv:1601.04919v5.
- [Sca15] C. W. Scaduto, *Instantons and odd Khovanov homology*, J. Topol. **8** (2015), 744–810.
- [Sim14] T. Simčević, *A Hardy space approach to Lagrangian Floer gluing*, Ph.D. thesis, ETH Zürich, Diss. ETH No. 22118, 2014; arXiv:1410.5998.
- [Weh04] K. Wehrheim, *Uhlenbeck Compactness*, EMS Series of Lectures in Mathematics, European Mathematical Society, Zürich, 2004.
- [WW15a] K. Wehrheim and C. T. Woodward, *Orientations for pseudoholomorphic quilts*, arXiv:1503.07803.
- [WW15b] K. Wehrheim and C. T. Woodward, *Pseudoholomorphic quilts*, J. Symplectic Geom. **13** (2015), 849–904.
- [WW10] K. Wehrheim and C. T. Woodward, *Quilted Floer cohomology*, Geom. Topol. **14** (2010), 833–902.
- [Wue26a] B. J. Wuebben, *Cobordism maps and Atiyah–Floer isomorphisms for framed instanton homology I*, manuscript, version of 19 October 2026 (first version 25 September 2026), <https://github.com/bwuebben/donaldson-atiyah-floer/tree/main/af-cob-1>.
- [Yan23] H. Yang, *A homological action on sutured instanton homology*, arXiv:2306.17353.

- [Zin16] A. Zinger, *The determinant line bundle for Fredholm operators: construction, properties, and classification*, Math. Scand. **118** (2016), 203–268.