

RG domain walls and hybrid triangulations

TUDOR DIMOFTE, DAVIDE GAIOTTO AND ROLAND VAN DER VEEN

This paper studies the interplay between the $\mathcal{N} = 2$ gauge theories in three and four dimensions that have a geometric description in terms of twisted compactification of the six-dimensional (2,0) SCFT. Our main goal is to construct the three-dimensional domain walls associated to any three-dimensional cobordism. We find that we can build a variety of 3d theories that represent the local degrees of freedom at a given domain wall in various 4d duality frames, including both UV S-dual frames and IR Seiberg-Witten electric-magnetic dual frames. We pay special attention to Janus domain walls, defined by four-dimensional Lagrangians with position-dependent couplings. If the couplings on either side of the wall are weak in different UV duality frames, Janus domain walls reduce to S-duality walls, *i.e.* domain walls that encode the properties of UV dualities. If the couplings on one side are weak in the IR and on the other weak in the UV, Janus domain walls reduce to RG walls, *i.e.* domain walls that encode the properties of RG flows. We derive the 3d geometries associated to both types of domain wall, and test their properties in simple examples, both through basic field-theoretic considerations and via comparison with quantum Teichmüller theory. Our main mathematical tool is a parametrization and quantization of framed flat $SL(K)$ connections on these geometries based on ideal triangulations.

1	Introduction	140
1.1	Framed 3-manifolds and the six-dimensional dictionary	141
1.2	Non-abelian couplings	145
1.3	Irregular defects	147

1.4	Duality and RG domain walls	148
1.5	Organization	150
2	The simplest couplings	151
2.1	Big boundaries and Seiberg-Witten theory	152
2.2	3-manifolds and class \mathcal{R}	160
2.3	Half-BPS boundary conditions	168
3	Janus domain walls in 4d $\mathcal{N} = 2$ gauge theories	174
3.1	General concepts	175
3.2	Janus configurations for class \mathcal{S}	177
3.3	Seiberg-Witten walls and couplings	180
4	RG walls	182
4.1	Janus attractors	184
4.2	Trapped Goldstone bosons	188
4.3	The basic RG manifold M_0	190
4.4	Line operators	194
5	Non-abelian symmetry enhancement	197
5.1	Gluing in the RG manifold M_0	199
5.2	Matter doublets	201
5.3	Three-punctured spheres and UV theories of class \mathcal{S}	203
5.4	Small tori	206
6	Example 1: $\mathcal{N} = 2^*$ interfaces	208

RG domain walls and hybrid triangulations	139
6.1 The RG wall	209
6.2 The S-wall: $T[SU(2)]$ with manifest $SU(2) \times SU(2)$ flavor symmetry	216
6.3 Reproducing standard $T[SU(2)]$	222
6.4 S-duality of line operators	225
7 Example 2: interfaces for $SU(2)$ $N_f = 4$	227
7.1 Enhanced flavor symmetry and $SO(12)$ surprises	229
7.2 The RG wall theory	234
7.3 The S-duality wall	236
7.4 The $6j$ symbol and volumes of non-ideal tetrahedra	239
Acknowledgements	241
Appendix A Coordinates from hybrid triangulations	242
A.1 Boundary coordinates	243
A.2 Twists for small annuli	248
A.3 Tetrahedra and coordinates in the bulk	253
Appendix B Triangulation of general RG and duality manifolds	262
B.1 The construction	263
Appendix C Quantum dilogarithms	269
References	270

1. Introduction

Over the last few years, a wealth of new results and insights about supersymmetric gauge theories in three and four dimensions have followed from two main directions of research. Protected quantities such as indices and correlation functions in curved space were computed exactly, testing known dualities, and predicting new ones. Moreover, the twisted compactification of six-dimensional $(2, 0)$ superconformal field theories provided an organizing principle for very large classes of lower dimensional field theories and their BPS defects and operators. The interplay between these two ideas is striking: protected calculations in a lower dimensional field theory often manifest the theory’s six-dimensional origin. The prominent example is the twisted compactification of the six-dimensional theories on a punctured Riemann surface, which produces four dimensional theories in “class \mathcal{S} ” [1, 2]. The superconformal index and the sphere partition function of a class \mathcal{S} theory take the form of two-dimensional correlation functions on the Riemann surface [3, 4].

The situation in three dimensions is somewhat analogous. It is possible to give an algebraic definition of a vast class of three-dimensional $\mathcal{N} = 2$ SCFT’s built in such a way that protected quantities have a natural interpretation in terms of the geometry of an auxiliary three-manifold. These SCFTs have UV Lagrangian descriptions as abelian Chern-Simons-Matter theories deformed by superpotential terms, which may contain monopole operators [5, 6]. The properties of such theories, which we will denote as “class \mathcal{R} ,” loosely match what one could expect to obtain from the twisted compactification of the six-dimensional A_{K-1} theories on three-manifolds with networks of line defects.¹ The expectation was reinforced by the M-theory analysis of [8], which explains how the abelian CSM structure emerges in the IR.

However, in order to give a precise six-dimensional construction of a given class \mathcal{R} theory, several subtleties remained to be addressed. In this paper, we will attempt to clarify them. Our motivation is not just to dot all i’s and cross all t’s. The three-dimensional theories are expected to define natural boundary conditions and domain walls for the four-dimensional theories. In order to make these relations manifest, we find it necessary to first clarify the above-mentioned subtleties.

¹The class \mathcal{R} of abelian Chern-Simons-Matter theories was defined purely algebraically in [5, 7], with no reference to three-dimensional geometries or Lie algebras. It was then shown (also in [6]) that for Lie algebras $\mathfrak{g} = A_{K-1}$ one can construct three-manifold theories that belong to this class.

The six-dimensional intuition predicts that to every cobordism between two punctured Riemann surfaces, one should be able to associate a natural domain wall between the corresponding class \mathcal{S} theories. Our main result is that the same domain wall can be given not one but several descriptions as distinct class \mathcal{R} theories, coupled to either a full UV description of the four-dimensional theories, or to an effective IR Seiberg-Witten description of the same theories, or to a hybrid mixture. Different descriptions will be useful in different regions of parameter space, depending on which bulk degrees of freedom are weakly coupled.

This includes two special cases. The first, which we expected to find from the beginning, is the class of “S-duality walls,” which implement the equivalence between different S-dual descriptions of the same four-dimensional theory [9, 10]. The second, which was unexpected, is a class of “RG walls,” which implement the relation between a UV and an IR description of the same theory. The discovery of RG walls inspired the related two-dimensional work of [11]. We also sketch a general argument explaining why such domain walls for any $\mathcal{N} = 2$ theory should admit class \mathcal{R} descriptions, with no reference to six dimensions. This can be thought of as a field theory version of the arguments in [8].

There are also some ancillary mathematical payoffs to our exploration. Using S-duality walls, we find a precise 3d interpretation of the integral kernels that implement the Moore-Seiberg groupoid in Liouville theory, or, equivalently, a mapping-class-group action on quantized Fenchel-Nielsen coordinates of Teichmüller theory. From RG domain walls, we find a 3d interpretation of the eigenfunctions of geodesic length operators in quantum Teichmüller theory [12–15]. For higher-rank Lie algebras A_{K-1} , these generalize to operations in Toda theory and higher Teichmüller theory [16], respectively.

1.1. Framed 3-manifolds and the six-dimensional dictionary

The oriented three-manifolds M that label theories in class \mathcal{R} are assembled by gluing together truncated tetrahedra (Figure 1) [5, 6]. Some (not necessarily all) of the big hexagonal faces of tetrahedra are glued together in pairs, while the small triangular faces at the truncated vertices are never glued. Therefore, M acquires two kinds of boundary: a *big* boundary tiled by unglued big hexagons, and a *small* boundary tiled




Figure 1: Truncated tetrahedron.

by small triangles. One additionally requires that the small boundary components have the topology of discs, annuli, or tori. Every hole in the big boundary is either filled in by a small disc, or connected to another hole by a small annulus. One such manifold is shown in Figure 2.




Figure 2: Left: a 3-manifold that admits a decomposition into truncated tetrahedra. It has two big boundaries (a 4-punctured sphere and a 3-punctured sphere) with fixed 2d triangulation \mathbf{t} ; as well as a small disc (filling in one puncture), three small annuli, and a small torus. Topologically, this manifold has just two connected boundary components, the torus and a surface of genus two. Right: schematic of how truncated tetrahedra are assembled around a small torus boundary.

We say that a three-manifold M whose boundary is separated into big and small pieces this way is *framed*. We will only consider framed three-manifolds in this paper, and will assume that they admit finite decompositions into truncated tetrahedra.²

For example, a truncated tetrahedron itself is a framed three-manifold, whose big boundary is a sphere with four holes, and whose small boundary consists topologically of four discs that fill in the holes. Another example is a knot complement, formed by gluing together tetrahedra so that the only remaining boundary is a small torus.

Notice that the big boundary of such a 3-manifold inherits a 2d ideal triangulation: a triangulation whose edges all begin and end at holes, or “punctures.” Then the 3d SCFT associated to M , for any Lie algebra $\mathfrak{g} = A_{K-1}$, depends on

- the 2d ideal triangulation \mathbf{t} of the big boundary; and

²We call our 3-manifolds “framed” because they are perfectly suited for the study of “framed” flat connections in the sense of Fock and Goncharov [16]. This use of “framing” should not be confused with a choice of trivialization of the tangent bundle TM ; the latter will not play an important role here.

- a choice of *polarization* Π for all boundaries, compatible with this triangulation.

Thus we can denote the 3d theory as $T_K[M; \mathbf{t}; \Pi]$. Notably, the theory does *not* depend on the 3d bulk triangulation of M , as long as that triangulation is compatible with the big-boundary triangulation \mathbf{t} . Indeed, different bulk triangulations lead to mirror-symmetric descriptions of $T_K[M; \mathbf{t}; \Pi]$ as a class- \mathcal{R} theory.

Geometrically, Π is a polarization for an open subset of the symplectic moduli space of framed flat $PGL(K)$ connections on ∂M , which we will review later. For small tori, the polarization is a choice of A and B cycles. For small annuli, the polarization is canonical — since there is a canonical non-contractible A-cycle on an annulus. For big boundaries \mathcal{C} , however, there is typically an $Sp(2d, \mathbb{Z})$ worth of choices, where $2d$ is the dimension of the moduli space of flat $PGL(K)$ connections on \mathcal{C} .

One of the central themes of this paper is that *the 3d $\mathcal{N} = 2$ theory $T_K[M, \mathbf{t}, \Pi]$ can provide a half-BPS boundary condition for 4d $\mathcal{N} = 2$ theories of class \mathcal{S} in many different ways*. One of these ways was described³ in [5], and is closely related to constructions of [8]: every big boundary \mathcal{C} of M can be coupled to a 4d Seiberg-Witten theory $SW_K[\mathcal{C}]$. The Seiberg-Witten theory is the IR description of the non-abelian theory $T_K[\mathcal{C}]$ obtained by twisted compactification of the 6d A_{K-1} theory on \mathcal{C} . $SW_K[\mathcal{C}]$ has gauge group $U(1)^d$, and an $Sp(2d, \mathbb{Z})$ electric-magnetic duality group. The gauge fields of $SW_K[\mathcal{C}]$ can be used to gauge abelian flavor symmetries of $T_K[M, \mathbf{t}, \Pi]$, after which the dependence on big-boundary polarization Π is erased.⁴ Subsequently, electric BPS hypermultiplets of $SW_K[\mathcal{C}]$ can be coupled to chiral operators of $T_K[M, \mathbf{t}, \Pi]$, after which the dependence on the big-boundary triangulation \mathbf{t} is erased.

We can reproduce this 3d-4d coupling from six dimensions (Figure 3). To do so, we first choose a metric on M such that all big boundaries \mathcal{C}_i are pulled out to infinity, forming semi-infinite cylindrical regions that asymptotically take the form $\mathbb{R}_+ \times \mathcal{C}_i$. Then we must “remove the regulator” from the small boundaries of M . That is, we replace each small boundary with a codimension-two defect of the 6d theory. Each small disc, filling in a hole on a big boundary \mathcal{C}_i , becomes a semi-infinite regular line defect stretching

³In [5] mainly the $\mathfrak{g} = A_1$ case was discussed. The generalization to A_{K-1} is straightforward following [6].

⁴More precisely, different choices of polarization become related by $Sp(2d, \mathbb{Z})$ electric-magnetic duality transformations of the combined 3d-4d system. We will review how \mathbf{t} and Π are “erased” in Section 2.3.




Figure 3: Adjusting the metric on the manifold M from Figure 2 and replacing small boundaries with line defects so that the compactification of the 6d theory on $M \times \mathbb{R}^3$ produces Seiberg-Witten theories on a half-space coupled to $T_K[M]$ as a boundary condition.

out into the corresponding cylindrical region. Each small annulus is shrunk to a regular line defect connecting punctures in the cylindrical regions. And each small torus is shrunk to closed regular line defect, in such a way that the A-cycle of the torus becomes an infinitesimally small loop linking the defect.

Upon compactifying the 6d A_{K-1} theory on $M \times \mathbb{R}^3$ with this metric, and flowing to the far infrared, one obtains a (non-interacting) product of Seiberg-Witten theories $\prod_i SW_K[C_i]$ on a half-space $\mathbb{R}_+ \times \mathbb{R}^3$, all coupled to the theory $T_K[M, \mathbf{t}, \Pi]$ on the boundary \mathbb{R}^3 . The coupling turns out to be precisely the one described above — we will review how this arises in Sections 2–3. The full 3d-4d system does not depend on \mathbf{t} or on big-boundary polarization Π .

Now, for this picture to make sense, two things must be true:

- 1) The regular line defects of the 6d A_{K-1} theory must be able to end. (Otherwise the dictionary between small discs and semi-infinite defects breaks down.)
- 2) The abelian flavor symmetries of $T_K[M, \mathbf{t}, \Pi]$ associated to the A-cycles of small annuli and tori, a priori $U(1)^{K-1}$, must be enhanced⁵ to $SU(K)$. (Because closed or infinite regular defects of the 6d theory carry non-abelian $SU(K)$ flavor symmetry.) More precisely, these flavor symmetries of $T_K[M]$ must be enhanced after coupling to Seiberg-Witten theories as above.

⁵This precise statement holds for *maximal* or *full* defects of the 6d A_{K-1} theory. Otherwise one expects enhancement to non-abelian subgroups of $SU(K)$. For $K = 2$ the regular defects are unique, so there is no ambiguity.

A large part of this paper will be devoted to understanding the second point, i.e. the non-abelian enhancement of flavor symmetry and its geometric origin. Constructions from our work [17] will play a key role. We will mainly focus on the A_1 theory, obtaining enhancements from $U(1)$ to $SU(2)$. The case of general K and maximal defects can begin to be analyzed using methods of [6].

The first point will not play a further role in this paper. Although we believe it is likely that codimension-two regular defects can end on a special quarter-BPS codimension-three defect, we will not attempt to define such an object here. We will only encounter small disks organized in a very specific configuration, which will be given an alternative six-dimensional interpretation in terms of irregular defects. We will explain momentarily that a rank r irregular defect in the six-dimensional A_1 theory joining two big boundary components can be naturally “regularized” by gluing the two boundary components along a $2r$ -sided polygon with small disks at the vertices.

1.2. Non-abelian couplings

Once we know that the flavor symmetries of $T_K[M, \mathbf{t}, \Pi]$ that are associated to small annulus and torus boundaries have non-abelian enhancements, we can attempt to interpret this 3d theory as a boundary condition for a more universal 4d bulk.

Let us set $K = 2$ for concreteness. Then for every small torus boundary, $T_2[M, \mathbf{t}, \Pi]$ can provide a quarter-BPS boundary condition for 4d $\mathcal{N} = 4$ $SU(2)$ super-Yang-Mills theory $T_2[T^2]$. In particular, the $SU(2)$ flavor symmetry associated to the A-cycle of the small torus is identified with the gauge symmetry of $T_2[T^2]$. This coupling has a six-dimensional origin that is closely related to the original argument for the flavor symmetry of defects. After using the dictionary above to shrink a small torus to closed regular defect in $M \times \mathbb{R}^2$, we can choose a metric for M that pulls the closed defect out to infinity, producing cylindrical regions of the form $T^2 \times \mathbb{R}_+ \times \mathbb{R}^3$. Compactification on T^2 leads to $\mathcal{N} = 4$ SYM on a half-space, coupled to $T_2[M, \mathbf{t}, \Pi]$.

More interesting are the small annuli. Recall that a small annulus connects two punctures (or holes) on the big boundary, which is made of surfaces \mathcal{C}_i . After we couple $T_2[M, \mathbf{t}, \Pi]$ to Seiberg-Witten theories $\prod_i SW_2[\mathcal{C}_i]$, the $SU(2)$ flavor symmetry associated to an annulus should be identified with the two $SU(2)$ flavor symmetries of the Seiberg-Witten theories associated to the big boundaries \mathcal{C}_i that the annulus ends on. This flavor symmetry can then be gauged in the product of SW theories in the 4d bulk.




Figure 4: Forming a genus-two boundary from two big three-punctured spheres connected by small annuli. This equivalent to a pants decomposition.

For example, consider a closed topological boundary component $\tilde{\mathcal{C}}$ of M consisting of a collection of big three-punctured spheres connected by small annuli (Figure 4). The separation into big and small pieces is equivalent to a pair-of-pants decomposition of $\tilde{\mathcal{C}}$. Each big three-punctured sphere is associated to a free 4d half-hypermultiplet in the tri-fundamental representation of an $SU(2)^3$ flavor group. We can then identify and gauge the pairs of $SU(2)$ flavor groups associated to punctures that are connected by small annuli. What we get is precisely the non-abelian UV description of the 4d class- \mathcal{S} theory $T_2[\tilde{\mathcal{C}}]$, in a UV duality frame corresponding to the given pants decomposition of $\tilde{\mathcal{C}}$ [1]. Now $T_2[M, \mathfrak{t}, \Pi]$ should provide a half-BPS boundary condition for $T_2[\tilde{\mathcal{C}}]$! This can also be motivated with a 6d construction.

The general picture now becomes the following. To the boundary of any framed 3-manifold M we claim that we can associate a universal 4d $\mathcal{N} = 2$ theory $T_2[\partial M]$, for which $T_2[M]$ provides a half-BPS boundary condition. Each closed topological component of the boundary ∂M contributes an independent factor to $T_2[\partial M]$. Small tori contribute $\mathcal{N} = 4$ SYM factors. The remaining boundaries $\tilde{\mathcal{C}}_i$ are decomposed into big punctured surfaces connected by small annuli, and thus contribute products of Seiberg-Witten theories in which pairs of $SU(2)$ flavor groups have been identified and gauged. Notice that this can be interpreted as the class- \mathcal{S} theory $T_2[\tilde{\mathcal{C}}_i]$ after a partial flow to the IR. In coupling $T_2[M]$ to $T_2[\partial M]$, *all* the flavor symmetries of $T_2[M]$ are gauged, and chiral operators of $T_2[M]$ couple via superpotentials to BPS hypermultiplets of $T_2[\partial M]$.

Starting from this general setup, it is easy to recover other boundary conditions by sending various 4d gauge couplings to zero (and masses to infinity). We can picture this operation as “cutting” the boundary ∂M into pieces. For example, by cutting all small annuli (sending the couplings of their $SU(2)$ ’s to zero) we recover a boundary condition for a product of pure Seiberg-Witten theories. Of course, we can also leave some of the small

annuli intact, leading to boundary conditions for products of pure Seiberg-Witten theories and non-abelian UV theories.

1.3. Irregular defects

We can even go further and cut the big boundaries of M into pieces along closed polygonal edge-paths in the boundary triangulation. We require that the paths' vertices lie at small discs (not annuli). Cutting $\mathcal{C} \rightarrow \mathcal{C}_1 \cup \mathcal{C}_2$ along a polygon corresponds to taking a limit in the parameter space of $SW_2[\mathcal{C}]$ so that BPS states associated to the edges of the polygon become infinitely massive while BPS states associated to all other edges (in a triangulation of \mathcal{C}) remain light. Following the 2d dictionary of [18], the new polygonal boundaries of \mathcal{C}_1 and \mathcal{C}_2 should now be interpreted as *irregular* punctures in these surfaces. The theory $SW_2[\mathcal{C}]$ degenerates into a product $SW_2[\mathcal{C}_1] \times SW_2[\mathcal{C}_2]$, where each $SW_2[\mathcal{C}_i]$ results from compactification of the 6d A_1 theory with irregular defects. The 3d theory $T_2[M]$ provides a BPS boundary condition for the product. The polygon can be thought of as a regularization of the irregular defect joining the irregular punctures on \mathcal{C}_1 and \mathcal{C}_2 .




Figure 5: Cutting the big boundary of a tetrahedron in half along a four-sided polygon (in blue).

The simplest illustration of this last scenario is for $M = \Delta$ a single tetrahedron (Figure 5). Since the big boundary of Δ is a four-punctured sphere, the theory $T_2[\partial\Delta]$ is the Seiberg-Witten description of 4d $\mathcal{N} = 2$ $N_f = 4$ SQCD. The IR gauge group is $U(1)$. In a “strong coupling” region of the u -plane the theory has a finite collection of BPS hypermultiplets with electric, magnetic, and dyonic charge. Cutting $\partial\Delta$ in half along a 4-sided polygon corresponds to taking an Argyres-Douglas type limit, in which only two electric particles survive. They are the BPS states of two “ A_1 ” Argyres-Douglas theories [19]. The tetrahedron theory $T_\Delta = T_2[\Delta, \mathbf{t}, \Pi]$ then provides a boundary condition for a product of the two Argyres-Douglas theories — or equivalently a domain wall between one and the other. This domain wall was the basic object of study in [8]. We will discuss it further in Section 3.3.

1.4. Duality and RG domain walls

Notice that a BPS boundary condition for a product of theories $T_1 \times T_2$ is equivalent to a BPS domain wall between T_1 and T_2 . Graphically, this might be thought of as in Figure 6. Therefore, when a 3-manifold M either has topologically disjoint boundary components or a boundary that has been “cut” into pieces as above, the theory $T_K[M, \mathbf{t}, \Pi]$ should provide a BPS domain wall. In this interpretation, two particular kinds of 3-manifolds that were anticipated in [20] play a special role. (For simplicity, we’ll continue to take $K = 2$.)

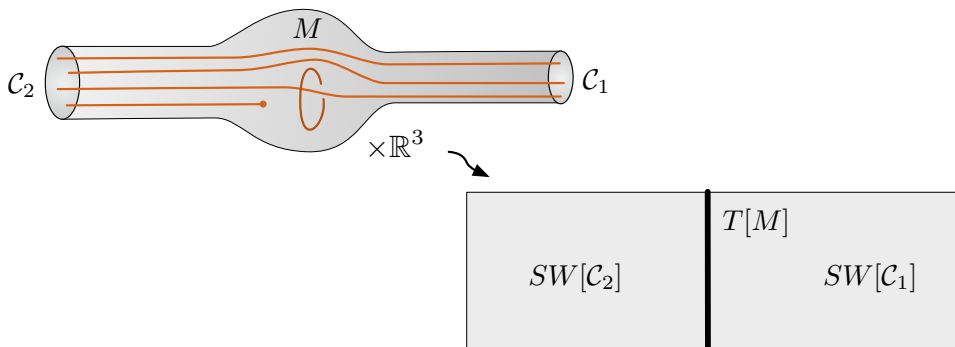


Figure 6: Reinterpreting the boundary condition of Figure 3 as a domain wall.

Let \mathcal{C} be a punctured surface, and construct a “trivial” 3-manifold $M = \mathcal{C} \times I$ that has two big boundaries \mathcal{C} , $\bar{\mathcal{C}} \simeq \mathcal{C}$, and a collection of small annuli connecting the punctures of \mathcal{C} directly to the punctures of $\bar{\mathcal{C}}$. We consider ∂M to be “cut” along these annuli. We have claimed that $T_2[M, \mathbf{t}, \Pi]$ provides a boundary condition for $SW_2[\mathcal{C}] \times SW_2[\bar{\mathcal{C}}]$, and thus a domain wall between the 4d Seiberg-Witten theory $SW_2[\mathcal{C}]$ and itself. It is a totally reflecting boundary, or a totally transparent domain wall.

To obtain something more interesting, let us choose a pants decomposition \mathbf{p} for $\bar{\mathcal{C}}$ (*i.e.* a maximal set of non-intersecting A-cycles), and form a new three-manifold $M_{\mathbf{p}}$ by shrinking the A-cycles, as in Figure 7. Specifically, although $M_{\mathbf{p}}$ and M are equivalent topologically, we have replaced the boundary $\bar{\mathcal{C}}$ in M by a network of big three-punctured spheres connected by small annuli. The 4d theory associated to this new boundary is the non-abelian class- \mathcal{S} theory $T_2[\mathcal{C}]$. We now expect the 3d theory $T_2[M_{\mathbf{p}}]$ to couple to $SW_2[\mathcal{C}] \times T_2[\mathcal{C}]$, and thus provide a natural domain wall between IR and UV descriptions of the same 4d theory — an RG domain wall!




Figure 7: Left: trivial cobordism $M = \mathcal{C} \times I$ with \mathcal{C} a big punctured torus. Right: shrinking an A-cycle on the “inner” boundary $\bar{\mathcal{C}}$ to create a new manifold $M_{\mathbf{p}}$; its inner boundary is now a big three-punctured sphere with one puncture connected to the outside and the other two connected to each other by small annuli.

We will construct RG domain walls in Section 3 purely in field theory by starting from Janus configurations where 4d parameters vary in space. In the limit that parameters vary very quickly, over a short interval, we will recover the 3d domain-wall theory $T_2[M_{\mathbf{p}}]$ coupled 4d UV and IR theories in the bulk.

The most basic property we would expect from a BPS RG domain wall is that UV bulk BPS line operators brought to the interface should match with their bulk IR description, brought to the interface from the opposite side. The correspondence between UV and IR descriptions of a BPS line defect was discussed in detail in [21], and thus we have in our hands an infinite set of constraints which the RG domain walls should obey. These constraints can be made more manageable by inserting the line defects inside a protected calculation, such as an ellipsoid partition function, or a sphere index. Alternatively, we can compare their vevs when the theory is compactified on a circle.

The correspondence between UV and IR operators will work automatically for our candidate RG domain walls, for a simple geometric reason. The UV line operators are associated to closed loops on the Riemann surface \mathcal{C} [22], and their vevs to the trace of the holonomy along the loops of a flat $PSL(2, \mathbb{C})$ connection on \mathcal{C} [21]. The 3d geometry is selected in such a way that the same holonomy can be computed along a corresponding loop on the big boundary region in terms of certain “edge coordinates.” The edge coordinates coincide with the vevs of IR line operators and the relation between holonomies and edge coordinates is known to encode the relation between UV and IR line defects [21]. The role of the 3d geometry is to allow us to transport the UV line operators from the UV end of the geometry to the IR end of the geometry, where the triangulation and edge coordinates live.




Figure 8: Shrinking A- and A'-cycles on both the inner and outer boundaries of M to create $M_{\mathbf{p},\mathbf{p}'}$. The resulting $M_{\mathbf{p},\mathbf{p}'}$ here can be thought of as the complement of a trivalent “Hopf network” in S^3 . It is the 3-manifold that gives rise to the theory $T[SU(2)]$.

A closely related modification of the trivial cobordism $M = \mathcal{C} \times I$ shrinks a set of A-cycles on both ends, according to two different pants decompositions of \mathcal{C} : \mathbf{p} and \mathbf{p}' . Call the resulting manifold $M_{\mathbf{p},\mathbf{p}'}$. Then the 3d theory $T_2[M_{\mathbf{p},\mathbf{p}'}]$ couples naturally to two copies of $T_2[\mathcal{C}]$ in different weakly coupled UV descriptions, corresponding to \mathbf{p} and \mathbf{p}' . We obtain a concrete Lagrangian formulation for the domain wall that implements the $\mathcal{N} = 2$ S-duality of [1], in direct analogy with the $\mathcal{N} = 4$ S-duality domain walls of [9].

The simplest example of such an S-duality domain wall is for 4d $\mathcal{N} = 2^*$ theory. The 3d domain-wall theory is an $\mathcal{N} = 2$ deformation of the theory called $T[SU(2)]$ in [9]. We will recover $T[SU(2)]$ as the theory corresponding to the 3-manifold in Figure 8. In fact, we will obtain several different mirror-symmetric descriptions of $T[SU(2)]$, including one whose $SU(2) \times SU(2)$ flavor symmetry is manifest.

1.5. Organization

Most of this paper will specialize to theories coming from the compactification of the A_1 six-dimensional (2,0) SCFT, *i.e.* to $K = 2$.

We begin in Section 2 by reviewing geometric properties of Seiberg-Witten theories of class \mathcal{S} and presenting their natural coupling to 3d theories of class \mathcal{R} at IR interfaces. In Section 3, we discuss properties of Janus configurations for 4d $\mathcal{N} = 2$ theories and explain (in principle) how

to extract from them 3d domain-wall theories and their bulk-boundary couplings. Generalizing the constructions of [8], we interpret Lagrangian definitions of the 3d theories in terms of triangulated 3-manifold geometries.

We apply our understanding of Janus configurations in Section 4 to build the most fundamental example of an RG domain wall, for pure $\mathcal{N} = 2$ $SU(2)$ gauge theory. We interpret it in terms of a framed 3-manifold geometry, establishing the basic geometric properties of UV (non-abelian) bulk-boundary couplings. Then in Section 5 we generalize the geometry of UV bulk-boundary couplings to any 4d (A_1) theory of class \mathcal{S} coupled to a 3d theory of class \mathcal{R} . We focus in particular on how the (a priori) abelian flavor symmetries in 3d get enhanced so that they can be identified with non-abelian symmetries in a 4d bulk.

Finally, in Sections 6 and 7 we demonstrate our various constructions in two detailed examples, describing RG walls and duality walls for 4d $SU(2)$ $\mathcal{N} = 2^*$ theory (coming from compactification on a punctured torus) and for 4d $SU(2)$ theory with $N_f = 4$ flavors of matter (coming from compactification on a four-punctured sphere).

In Appendix A we include mathematical definitions of various moduli spaces of framed flat connections used in the paper, their coordinates (including a novel definition of complexified Fenchel-Nielsen twists), and their quantization. In Appendix B we explain how to build a useful 3d triangulation for any framed 3-manifold corresponding to an RG or duality wall.

2. The simplest couplings

In this section, we seek to describe the natural couplings between a 3d theory of class \mathcal{R} and an abelian Seiberg-Witten theory of class \mathcal{S} . These couplings are associated to the big boundary of a framed 3-manifold. The main point is that at fixed values of its moduli a 4d Seiberg-Witten theory $SW_2[\mathcal{C}]$ comes with a family of natural WKB triangulations of the surface \mathcal{C} , which encode much of the 4d physics. By matching a WKB triangulation of \mathcal{C} with the big-boundary triangulation \mathbf{t} of a 3-manifold M , we can encode the physics of the coupling between $SW_2[\mathcal{C}]$ and $T_2[M, \mathbf{t}]$ — including a proper identification of gauge and flavor symmetries, and superpotential terms involving BPS hypermultiplets. We always have $K = 2$, studying theories coming from the $A_1(2, 0)$ theory in 6d.

2.1. Big boundaries and Seiberg-Witten theory

We begin by reviewing some basic properties of Seiberg-Witten theories in class \mathcal{S} , and their relation to triangulations and polarizations of surfaces. These ideas were developed in [18], following [23–25], [1, 2], and the mathematical constructions of [16].

We would like to consider the six-dimensional $A_1(2,0)$ theory in the background $\mathcal{C} \times \mathbb{R}^4$, where \mathcal{C} is a real two-dimensional oriented surface.⁶ We allow real codimension-two defects that fill all of \mathbb{R}^4 , and show up as “punctures” or “holes” on \mathcal{C} . An appropriate topological twist⁷ along \mathcal{C} allows eight supercharges to be preserved, along with an $SU(2)_R \times U(1)_r$ R-symmetry. At low energies, we expect to find an effective 4d $\mathcal{N} = 2$ theory $T_2[\mathcal{C}]$. It is a superconformal theory, and $U(1)_r$ remains non-anomalous, as long as all defects are regular, which we will assume in this discussion.

One may choose an arbitrary background metric for \mathcal{C} , though after flowing to the IR the theory $T_2[\mathcal{C}]$ only depends on its complex (or conformal) structure. If \mathcal{C} has genus g and h holes, the complex structure moduli space has complex dimension

$$(2.1) \quad d = d(g, h) := 3g - 3 + h.$$

An additional h parameters are used to prescribe boundary conditions for the complex structure at the holes. In the limits where \mathcal{C} is stretched into pairs of pants, $T_2[\mathcal{C}]$ acquires a weakly coupled Lagrangian description as an $SU(2)^d$ gauge theory, with $SU(2)^h$ flavor symmetry. The complex structure of each stretched internal leg in \mathcal{C} provides the gauge coupling for a dynamical $SU(2)$, while each global $SU(2)$ is associated with a boundary hole.

If we move onto the Coulomb branch and flow further to the infrared, $T_2[\mathcal{C}]$ can be described as a Seiberg-Witten theory $SW_2[\mathcal{C}]$, with gauge group broken to $U(1)^d$. Of course the $SU(2)^h$ flavor symmetry persists. The Seiberg-Witten curve Σ is a double cover of \mathcal{C} that is branched at exactly $4g - 4 + 2h = -2\chi(\mathcal{C})$ points. The curve depends on the d marginal couplings from the UV and h mass parameters, as well as d new Coulomb moduli.

⁶We will use both Euclidean and Lorentzian signature at various points, and only distinguish between the two when necessary. In the case of Lorentzian signature, time is always one of the directions in \mathbb{R}^4 .

⁷This twist identifies the $SO(2)$ isometry group of \mathcal{C} with the first factor in the subgroup $SO(2)_R \times SO(3)_R$ of the $SO(5)_R$ R-symmetry group of the 6d theory.

Many properties of $SW_2[\mathcal{C}]$ can be understood by considering *WKB triangulations* of \mathcal{C} [18]. These are ideal triangulations of \mathcal{C} — meaning that all edges in the triangulation start and end at holes⁸ — that depend on a point u on the Coulomb branch and a phase θ . (More generally, the triangulation depends on mass parameters m and marginal UV couplings τ ; we include these parameters in ‘ u ’.) A WKB triangulation $\mathbf{t}_{u,\theta}$ at generic (u, θ) is obtained by looking at trajectories on \mathcal{C} such that the Seiberg-Witten form λ_{SW} has constant phase θ ; each triangle is a region of \mathcal{C} that contains a unique branch point of the cover $\Sigma \rightarrow \mathcal{C}$ and whose boundaries the trajectories never cross (Figure 9). From a topological perspective, the dependence of the triangulation on (u, θ) is piecewise constant.




Figure 9: A typical WKB triangulation of a part of \mathcal{C} . The SW fibration $\Sigma \rightarrow \mathcal{C}$ has a branch point in each triangle, and each edge E_i corresponds to a constant-phase path γ_i on Σ . Here (for example) $\langle \gamma_1, \gamma_2 \rangle = +1$, while $\langle \gamma_1, \gamma_3 \rangle = 0$, etc.

At fixed u , the charge lattice $\Gamma_{\mathcal{C}}$ of $SW_2[\mathcal{C}]$ consists roughly of elements in $H_1(\Sigma, \mathbb{Z})$ that are odd under deck transformations.⁹ It includes a sublattice $\Gamma_f \sim \mathbb{Z}^h$ of flavor charges, which map to small loops around the punctures of Σ . The lattice of gauge charges $\Gamma_g = \Gamma_{\mathcal{C}}/\Gamma_f \sim \mathbb{Z}^{2d}$ is the odd part of $H_1(\bar{\Sigma}, \mathbb{Z})$, where $\bar{\Sigma}$ denotes Σ with its punctures removed. The electromagnetic skew-symmetric product $\langle \gamma, \gamma' \rangle$ is just the intersection product for 1-cycles in Σ ; it is non-degenerate on Γ_g and vanishes for any pure-flavor

⁸For these triangulations to make sense, we must assume that $h > 0$: there is always at least one hole.

⁹We refer to [21] for a detailed description of the charge lattice.

charge $\gamma \in \Gamma_f$. The central charge of a state with gauge/flavor charge $\gamma \in \Gamma_{\mathcal{C}}$ can be written as the period of the SW differential

$$(2.2) \quad Z_\gamma = \gamma \cdot (a, a_D, m) = \int_\gamma \lambda_{SW},$$

where m are the complex masses associated with the flavor symmetry. Abelian electric-magnetic duality acts by $Sp(2d, \mathbb{Z})$ transformations on the gauge lattice Γ_g , and dual transformations on (a, a_D) .

Any WKB triangulation $\mathbf{t}_{u,\theta}$ of \mathcal{C} provides a basis for $\Gamma_{\mathcal{C}}$.¹⁰ To understand this, note that any edge E of the triangulation separates two triangles, each of which contains a branch point. The unique open on \mathcal{C} that connects the branch points and crosses E lifts to a cycle $\tilde{\gamma}_E$ on Σ , and thus defines a charge γ_E . (With some care, the orientation of \mathcal{C} can be used to consistently assign an orientation to $\tilde{\gamma}_E$, and thus to fix the sign of γ_E .) It is not hard to see that

- 1) The product $\langle \gamma_E, \gamma_{E'} \rangle$ equals the (signed) number of oriented triangles shared by edges E and E' . Thus $\langle \gamma_E, \gamma_{E'} \rangle \in \{0, \pm 1, \pm 2\}$.
- 2) The sum $\gamma_v := \sum_{E \text{ touches } v} \gamma_E$ of edge charges for the edges that touch a puncture v equals twice the flavor charge associated to that puncture. Thus $\langle \gamma_v, * \rangle = 0$.

In addition, every edge E of a WKB triangulation $\mathbf{t}_{u,\theta}$ that is not the univalent edge in a degenerate triangle as in Figure 10 also corresponds to a BPS hypermultiplet Φ_E of $SW_2[\mathcal{C}]$. The hypermultiplet has charge γ_E and is stable in a region of the Coulomb branch that includes u . Geometrically, it arises because the topological cycle $\tilde{\gamma}_E \subset \Sigma$ can be realized as a unique path of minimal length on Σ , on which λ_{SW} has some constant phase.¹¹




Figure 10: A univalent edge in a degenerate triangle.

¹⁰More precisely, the WKB triangulation provides a basis for a sublattice of $\Gamma_{\mathcal{C}}$ generated by the charges of BPS particles in the theory. Sometimes this is a proper sublattice of $\Gamma_{\mathcal{C}}$.

¹¹In the M-theory construction of $SW_2[\mathcal{C}]$, where an M5 brane wraps $\Sigma \times \mathbb{R}^4$, M2 branes can end on the cycles $\tilde{\gamma}_W$ without breaking SUSY, giving rise to BPS particles in \mathbb{R}^4 .

The univalent edge E_f in a degenerate triangle is somewhat special. It carries pure flavor charge $2\gamma_f$ associated to the central puncture, in the normalization where a doublet for the $SU(2)_f$ flavor symmetry at the puncture has charges $\pm\gamma_f$. There is no BPS hypermultiplet of charge $2\gamma_f$ in the spectrum of $SW_2[\mathcal{C}]$, because there is no minimal-length path $\tilde{\gamma}_{E_f} \subset \Sigma$ that just winds around the puncture. Instead, there are *two* hypermultiplets $(\Phi_f, \tilde{\Phi}_f)$ with charges $\gamma' \pm \gamma_f$, in a doublet of $SU(2)_f$, where the charge $\gamma_{E'} = \gamma' - \gamma_f$ is naturally associated to the outer edge E' that surrounds the degenerate triangle. Notice that the full BPS spectrum of the theory will be $SU(2)_f$ invariant only for sufficiently small values of the $SU(2)_f$ mass parameter. For general values of the mass parameter, wall-crossing can break the $SU(2)_f$ multiplets. The presence of the doublet of hypers is thus a bit special. To understand it, and many other facts, we need to review the notion of flip.




Figure 11: A flip of the WKB triangulation. Here the nonzero intersection products are $\langle \gamma_a, \gamma_E \rangle = \langle \gamma_c, \gamma_E \rangle = -1$, $\langle \gamma_b, \gamma_E \rangle = \langle \gamma_d, \gamma_E \rangle = \langle \gamma_a, \gamma_d \rangle = \langle \gamma_c, \gamma_b \rangle = 1$. Note how the transformed charges after the flip continue to have the right intersection products.

As the phase θ is varied, a WKB triangulation $\mathbf{t}_{u,\theta}$ may jump. In particular, at a critical value θ^* that allows some pair of branch points to be connected by a phase- θ^* trajectory, an edge of the WKB triangulation “flips” ($E \rightarrow E'$) as in Figure 11. The BPS hypermultiplets associated to the edges E, E' on either side of the flip are a particle/anti-particle pair, with $\gamma_{E'} = -\gamma_E$; they both correspond to the cycle $\tilde{\gamma}_E$ of phase θ^* on Σ , in two opposite orientations. The charges of the four other edges of the quadrilateral involved in the flip also change by multiples of γ_E .

There is another special phenomenon, the “pop,” which will happen at the critical values in θ aligned with the phase of the mass parameter at a puncture. It helps understand why the degenerate triangle is associated to a doublet of hypermultiplets. The pop switches the sign of the flavor charge γ_f and thus exchanges the roles of the two hypers in the doublet as the BPS particle associated to the edge E' . As we vary θ by π , the pop will happen exactly once, and the triangulation at a pop always includes a degenerate

triangle at the puncture. If this is our original degenerate triangle, we are done. If not, a flip must have happened at some point to the edge E' of our original degenerate triangle. After that flip, there are two edges going into the puncture, associated to charges $\gamma_f \pm \gamma'$, and thus we have again the desired doublet of hypers in the BPS spectrum.

2.1.1. Framed flat connections. We need one more important idea to relate the combinatorics of theories of class \mathcal{R} to boundary SW theories of class \mathcal{S} .

When a theory of class \mathcal{S} is compactified on a circle S^1_β of finite radius, one obtains in the far infrared a 3d sigma model to a hyperkähler moduli space $\mathcal{P}_2(\mathcal{C})$ of real dimension $4d$. There are several nice ways to understand it. Very roughly $\mathcal{P}_2(\mathcal{C})$ is obtained by treating the d -complex-dimensional Coulomb branch of $SW_2[\mathcal{C}]$ as a real moduli space and using electric and magnetic Wilson lines around S^1_β to complexify it again. A more accurate description is that $\mathcal{P}_2(\mathcal{C})$ is the $SU(2)$ Hitchin moduli space associated to \mathcal{C} . Indeed, upon compactification on S^1_β the 6d A_1 theory becomes maximally supersymmetric 5d Yang-Mills theory. In the presence of the topological twist on \mathcal{C} , its BPS equations are Hitchin equation for a real $SU(2)$ connection A and a complex adjoint-valued one-form field φ .

The complex structure on $\mathcal{P}_2(\mathcal{C})$ is parametrized by a twistor coordinate $\zeta \in \mathbb{CP}^1$. The most useful fact for us is that when $\zeta \neq 0, \infty$ the solution to the Hitchin equations can be repackaged as a flat complex connection

$$(2.3) \quad \mathcal{A} = A + \frac{i\beta}{2}(\zeta^{-1}\varphi + \zeta\bar{\varphi}).$$

We will usually take $\zeta = e^{i\theta}$ to be a pure phase, so that $\mathcal{A} = A + i \operatorname{Re}(\zeta^{-1}\varphi)$. Then $\mathcal{P}_2(\mathcal{C})$ is *essentially* isomorphic, as a complex symplectic manifold, to the moduli space of flat $PGL(2, \mathbb{C})$ connections on \mathcal{C} . This is how we will think of it from now on. On smooth parts of the moduli space, the canonical holomorphic symplectic form is given by the Atiyah-Bott-Goldman formula [26, 27]

$$(2.4) \quad \Omega = \int_{\mathcal{C}} \operatorname{Tr} \delta\mathcal{A} \wedge \delta\mathcal{A}.$$

The smooth parts of $\mathcal{P}_2(\mathcal{C})$ have complex dimension $2d = 6g - 6 + 2h$.

The isomorphism between $\mathcal{P}_2(\mathcal{C})$ and the moduli space of flat complex connections is subject to two important caveats. One will be addressed in Section 2.1.2: we must sometimes consider lifts of flat connections to $SL(2, \mathbb{C})$, and sometimes connections halfway between $PGL(2, \mathbb{C})$ and

$SL(2, \mathbb{C})$, depending on the precise choice of UV gauge groups in $T_2[\mathcal{C}]$. The other caveat is that the isomorphism only holds when the space of flat connections is smooth.¹²

To be more precise, in the presence of punctures, the *three* real mass parameters for $SW_2[\mathcal{C}]$ on $\mathbb{R}^2 \times S^1_\beta$ associated to the flavor symmetry at each puncture control the Kähler class and complex structure of the moduli space. When complex-structure mass parameters vanish, the standard moduli space of flat $PGL(2, \mathbb{C})$ connections would develop a singularity; but the additional nonzero Kähler mass partially resolves the actual moduli space $\mathcal{P}_2(\mathcal{C})$. A simple shortcut to deal with this resolution is to supplement the complex flat connections by a choice of *framing*.

Indeed, the eigenvalues of the holonomy around each puncture are fixed in terms of the complex-structure masses. But in addition we introduce on \mathcal{C} an associated bundle of flags in \mathbb{C}^2 , and choose in the neighborhood of each puncture a flat section that is invariant under the $PGL(2, \mathbb{C})$ holonomy. Since a flag in \mathbb{C}^2 is just a complex line, this amounts to choosing an eigenline. Then $\mathcal{P}_2(\mathcal{C})$ is identified as the space of framed flat connections on \mathcal{C} :

$$(2.5) \quad \mathcal{P}_2(\mathcal{C}) \simeq \{\text{framed } PGL(2, \mathbb{C}) \text{ connections on } \mathcal{C}\}.$$

The choice of framing accomplishes the desired de-singularization of $\mathcal{P}_2(\mathcal{C})$. If the squares of the two eigenvalues $\lambda^{\pm 2}$ at a given puncture are distinct (*i.e.* $\lambda \neq \pm 1$), the framing just chooses one or the other, so the framed moduli space is a two-fold cover of the standard one. However, if the eigenvalues coincide, the choice of framing adds a \mathbb{CP}^1 to the moduli space, thereby blowing up a singularity. Given an ideal triangulation of \mathcal{C} (and a specific framing), one can define a set of coordinates for a Zariski-open patch of $\mathcal{P}_2(\mathcal{C})$ (isomorphic to $(\mathbb{C}^*)^{2d}$) that are associated to the edges of a WKB triangulation [16].¹³ Mathematically, the coordinate x_E for edge E is defined by taking the cross-ratio of four framing lines at the vertices of the quadrilateral containing E (Figure 12).¹⁴ Physically, x_E is the expectation value of a half-BPS line operator of charge γ_E that wraps S^1_β in the Seiberg-Witten theory $SW_2[\mathcal{C}]$. Thus, $\mathcal{P}_2(\mathcal{C})$ is parametrized by vevs of IR line operators.

¹²We refer again to [21] for more detailed discussion.

¹³These are a natural generalization of Thurston's shear coordinates in 2d hyperbolic geometry, thereafter studied by Fock, Penner, and others.

¹⁴Conventions for cross-ratios are as in [6]. In particular, the cross-ratios used in this paper, which are natural from the 3d perspective, are *minus* the positive coordinates of [16].

The subalgebra of the $\mathcal{N} = 2$ supersymmetry algebra preserved by the line operators is characterized by the same phase ζ that determined the complex structure of $\mathcal{P}_2(\mathcal{C})$.¹⁵




Figure 12: Defining the edge coordinate x_E as a cross-ratio of four framing flags A, B, C, D . The flags must be parallel-transported to a common point inside the quadrilateral in order to evaluate the cross-ratio.

The holomorphic edge coordinates x_E obey Poisson brackets induced by (2.4), which are simply determined in terms of the electric-magnetic product. Namely,

$$(2.6) \quad \{x_E, x_{E'}\} = \langle \gamma_E, \gamma_{E'} \rangle x_E x_{E'}.$$

In particular, the product of the edge coordinates surrounding a given puncture is a central element. It is equal to the square of the $PGL(2)$ eigenvalue (as chosen by the framing) at that puncture:

$$(2.7) \quad \prod_{E \text{ ending at } v} (-x_E) = \lambda^2.$$

This follows easily by multiplying cross-ratios as in Figure 12.

The line operators can be *quantized* by adding angular momentum to the $\mathbb{R}^3 \times S^1_\beta$ geometry [21, 28]. One considers a fibered product $\mathbb{R} \times (\mathbb{R}^2 \times_q S^1_\beta)$ such that a complex coordinate z on \mathbb{R}^2 undergoes a rotation $z \rightarrow qz$ after a turn around S^1_β . An additional R-symmetry twist allows this background to preserve half the supersymmetry. BPS line operators are then constrained to live at the origin of \mathbb{R}^2 and any point on \mathbb{R} , while wrapping S^1_β . They satisfy relations of a quantum torus algebra

$$(2.8) \quad \hat{x}_E \hat{x}_{E'} = q^{\langle \gamma_E, \gamma_{E'} \rangle} \hat{x}_{E'} \hat{x}_E,$$

where the ordering of \hat{x}_E 's is precisely the ordering of line operators along \mathbb{R} .

¹⁵In contrast to domain walls, BPS line defects typically preserve a full $SU(2)_R$ as well.

It is often convenient to use formal logarithmic variables \hat{X}_E so that $\hat{x}_E = \exp \hat{X}_E$. Then

$$(2.9) \quad [\hat{X}_E, \hat{X}_{E'}] = \langle \gamma_E, \gamma_{E'} \rangle \hbar, \quad q = e^{\hbar};$$

and now there is also a simple formula for the central elements (*i.e.* pure flavor Wilson lines)

$$(2.10) \quad \sum_{E \text{ ending at } v} (\hat{X}_E - i\pi - \frac{\hbar}{2}) = 2\Lambda,$$

where $\lambda = \exp \Lambda$ is the holonomy eigenvalue at the puncture v . The factors of \hbar (and the less significant $i\pi$'s) in this formula arise from the R-symmetry twist in the physical geometry.

2.1.2. $SL(2)$ vs $PGL(2)$. We mentioned above that the precise moduli space of (framed) flat connections isomorphic to the space of vacua $\mathcal{P}_2(\mathbb{C})$ is sometimes neither an $SL(2)$ moduli space nor a $PGL(2) \simeq PSU(2) \simeq SL(2)/\mathbb{Z}_2$ moduli space, but rather a space halfway inbetween. The subtle choice descends from choices of UV gauge groups.

The simplest example to illustrate this is the $\mathcal{N} = 2^* SU(2)$ gauge theory, *i.e.* the theory associated to a one-punctured torus. It is well known that S-duality exchanges the theory with $SU(2)$ gauge group and the theory with $PSU(2) \simeq SU(2)/\mathbb{Z}_2$ gauge group. S-duality simply exchanges the A- and B-cycles of the torus. The expectation values of Wilson loop operators for the $SU(2)$ gauge theory wrapping the circle S^1_β map to traces of the holonomy of the flat connection on the A-cycle of the one-punctured torus in the corresponding representations of $SL(2)$. The expectation value of 't Hooft loop operators map to traces of the B-cycle holonomy.

In the $\mathcal{N} = 2^* SU(2)$ gauge theory, the Wilson loop in a fundamental representation is allowed, but the fundamental 't Hooft loop is not. The converse is true for a $\mathcal{N} = 2^* PSU(2)$ gauge theory. On the other hand, the theories have the same set of BPS particles, which have electric and magnetic charges (associated to IR edge-coordinates x_E) with the same quantization as Wilson and 'tHooft loop operators of integer spin.




Figure 13: Punctured torus, with triangulation and A,B-cycles.

We arrive at the following picture:

- The moduli space of $PSL(2)$ connections on the one-punctured torus is “too small”: it admits functions that correspond to vevs of Wilson and ’tHooft loop operators of integer spin only. The vevs of these operators can be expressed as Laurent polynomials in the x_E variables for the three edges of a triangulated one-punctured torus. The intersection pairing of the γ_E cycles is even: $\langle \gamma_E, \gamma'_E \rangle = \langle \gamma'_E, \gamma''_E \rangle = \langle \gamma''_E, \gamma_E \rangle = 2$.
- The moduli space of $SL(2)$ flat connections on the one-punctured torus is “too big”: it admits functions which correspond to vevs of both fundamental Wilson and ’t Hooft loop operators. These vevs can be expressed in terms of the x_E only if we allow certain square roots of x_E monomials. The square roots create sign problems in defining things like the quantum torus algebra.
- The moduli space of vacua for the $\mathcal{N} = 2^* SU(2)$ gauge theory is an intermediate quotient of the moduli space of $SL(2)$ flat connections, where we keep track of the A-cycle holonomy in $SL(2)$, but only of the $PSL(2)$ image of the B-cycle holonomy. Some square roots of monomials of the X_E have to be allowed, but just enough to keep signs under control, and the quantum torus algebra well defined.

Simple generalizations of these statements hold for Riemann surfaces of higher genus g . The gauge group has a \mathbb{Z}_2^g center that acts trivially on matter fields, and thus can be left ungauged or gauged in various patterns. The moduli space of $SL(2)$ flat connections can be quotiented in various ways, so that one keeps track of holonomies in $SL(2)$ for a maximal set of cycles with even mutual intersection, and only in $PSL(2)$ for other cycles. The correspondence between four-dimensional choices of gauge group and two-dimensional choices of $SL(2)$ vs $PSL(2)$ can be established through the map between four-dimensional line defects and closed curves in the two-dimensional geometry [21].

2.2. 3-manifolds and class \mathcal{R}

The 3d $\mathcal{N} = 2$ theory in class \mathcal{R} associated to a triangulated 3-manifold M (of the type described in the introduction) is constructed algorithmically in a way that mimics the construction of framed $PGL(2, \mathbb{C})$ flat connections on M . We review here the main results of the construction, and the basic relations between the geometry of flat connections and 3d physics.

2.2.1. Framed flat connections in 3d. To begin, we recall that in three dimensions the choice of framing for a flat connection should be associated to the *small boundary* of M [5, 6]. Since the small boundaries have abelian fundamental group, any standard flat connection on M admits a choice of framing — *i.e.* a choice of invariant flag on the small boundary. One then considers several algebraic moduli spaces:¹⁶

$$(2.11) \quad \begin{aligned} &\bullet \mathcal{P}_2(\partial M) = \{\text{framed flat } PGL(2, \mathbb{C}) \text{ connections on } \partial M\} \\ &\bullet \tilde{\mathcal{L}}_2(M) = \{\text{framed flat } PGL(2, \mathbb{C}) \text{ connections on } M\} \\ &\bullet \mathcal{L}_2(M) = \{\text{framed flat connections on } \partial M \text{ that extend to } M\} \\ &\quad \subset \mathcal{P}_2(\partial M). \end{aligned}$$

Notice that $\mathcal{L}_2(M)$ is just the image of the natural projection $\tilde{\mathcal{L}}_2(M) \rightarrow \mathcal{P}_2(\partial M)$; often $\tilde{\mathcal{L}}_2(M)$ and $\mathcal{L}_2(M)$ are birationally isomorphic. The space $\mathcal{L}_2(M)$ is ultimately the one we are most interested in. Mathematically, a certain Zariski-open subset of $\mathcal{L}_2(M)$ is a Lagrangian submanifold inside $\mathcal{P}_2(\partial M)$.¹⁷ Physically, (this subset of) $\mathcal{L}_2(M)$ ends up describing the space of supersymmetric vacua of $T_2[M]$ on $\mathbb{R}^2 \times S^1$. We need a few more details to state this properly.

Suppose that we fix a 2d ideal triangulation \mathbf{t} for the big boundary of M , whose i -th component has (say) genus g_i and h_i holes. Suppose also that M has a small annular boundaries and t small torus boundaries. Then $\mathcal{P}_2(\partial M)$ generically has dimension $\sum_i 2d_i + 2a + 2t$, with $d_i = 3g_i - 3 + h_i$ as usual. Using the triangulation \mathbf{t} , the framing data on the boundary can be used to construct coordinates on an open patch of $\mathcal{P}_2(\partial M)$, which we can call $\mathcal{P}_2(\partial M, \mathbf{t})$, that is isomorphic to a complex torus

$$(2.12) \quad \mathcal{P}_2(\partial M, \mathbf{t}) \simeq \left[\prod_i (\mathbb{C}^*)^{2d_i} \right] \times (\mathbb{C}^*)^{2a} \times (\mathbb{C}^*)^{2t}.$$

The first factor simply corresponds to big-boundary moduli spaces, with coordinates x_E as in Section 2.1.1. The last factor is completely decoupled, and is parametrized by A- and B-cycle eigenvalues for each small torus

¹⁶A complete summary of these spaces and their coordinates is included in Appendix A.

¹⁷In fact, $\mathcal{L}_2(M) \subset \mathcal{P}_2(\partial M)$ is a subvariety with the rather special property that a canonical class in the second algebraic K -theory group of the function field of $\mathcal{P}_2(\partial M)$ vanishes when restricted to $\mathcal{L}_2(M)$. Thus it is a “ K_2 Lagrangian.” This property is reviewed in [6, 29] and follows immediately from the symplectic reduction of tetrahedron moduli spaces described below.

(which are canonically conjugate with respect to the complex symplectic form on $\mathcal{P}_2(\partial M)$).

The middle factor in (2.12) is rather more interesting. It is parameterized by a A-cycle eigenvalues along with a canonically conjugate “twist” coordinates. These are analogous to Fenchel-Nielsen length and twist coordinates in 2d hyperbolic geometry [30, 31]. Since the construction of these coordinates for framed flat connections has not been fully described before, we present it in detail in Appendix A.2.




Figure 14: Relation between annulus coordinates (length λ , twist τ) and big-boundary eigenvalues.

Note that the length coordinates for each small annulus (the A-cycle eigenvalues) provide the puncture eigenvalues at the two punctures on the big boundary where the annulus ends (Figure 14). From the perspective of the big boundary, these are fixed. Moreover, the holonomy around any puncture on the big boundary that is filled in by a small disc is defined to be unipotent. If it were not so, it would be impossible to extend flat connections from ∂M to M .




Figure 15: $\mathcal{P}_2(\partial\Delta)$

A standard example of a pair $\mathcal{L}_2(M) \subset \mathcal{P}_2(\partial M)$ occurs when $M = S^3 \setminus \mathcal{K}$ is a knot complement [32]. Then $\mathcal{P}_2(\partial M) \simeq \mathbb{C}^* \times \mathbb{C}^*$ and $\mathcal{L}_2(M)$ is cut out by a polynomial in two variables, the A-polynomial of the knot [33]. A much more fundamental example is for $M = \Delta$ a tetrahedron. The tetrahedron has a canonical boundary triangulation \mathbf{t} , leading to a phase space

$$(2.13) \quad \begin{aligned} \mathcal{P}_{\partial\Delta} := \mathcal{P}_2(\partial\Delta, \mathbf{t}) &= \{z, z', z'' \in \mathbb{C}^* \mid z z' z'' = -1\} \simeq \mathbb{C}^* \times \mathbb{C}^*; \\ \{z, z'\} &= z z', \quad \{z', z''\} = z' z'', \quad \{z'', z\} = z'' z. \end{aligned}$$

The coordinates z, z', z'' label pairs of opposite edges of the tetrahedron. The fact that the product equals -1 enforces unipotent holonomy, *cf.* (2.7).

The Lagrangian submanifold is [5, 34]

$$(2.14) \quad \mathcal{L}_\Delta := \mathcal{L}_2(\Delta, \mathbf{t}) = \{z'' + z^{-1} - 1 = 0\}.$$

Notice that the pair $\mathcal{L}_\Delta \subset \mathcal{P}_{\partial\Delta}$ would not make sense without the addition of framing! Indeed, the standard space of flat connections on a tetrahedron is trivial. In this case, all information lies in the framing, and identifies \mathcal{L}_Δ with the configuration space of four lines in \mathbb{C}^2 .

In general, in each of the complex-torus patches $\mathcal{P}_2(\partial M, \mathbf{t})$ of a boundary phase space, labelled by a boundary triangulation \mathbf{t} , we can consider the Lagrangian submanifold $\mathcal{L}_2(M, \mathbf{t}) := \mathcal{L}_2(M) \cap \mathcal{P}_2(\partial M, \mathbf{t})$. A central property of the Lagrangian pairs $\mathcal{L}_2(M, \mathbf{t}) \subset \mathcal{P}_2(\partial M, \mathbf{t})$ is that they can be glued together by *symplectic reduction*.

Namely, if one glues together two manifolds M_1, M_2 along part of a common big boundary (where their boundary triangulations $\mathbf{t}_1, \mathbf{t}_2$ agree) to obtain a new manifold M , then [34]

$$(2.15) \quad \mathcal{P}_2(\partial M, \mathbf{t}) \simeq [\mathcal{P}_2(\partial M_1, \mathbf{t}_1) \times \mathcal{P}_2(\partial M_2, \mathbf{t}_2)] // (\mathbb{C}^*)^{\# \text{ internal edges}},$$

taking a symplectic quotient with respect to a $\mathbb{C}^* \simeq GL(1, \mathbb{C})$ action for every new internal edge created during the gluing. The products of big-boundary coordinates x_E around these internal edges are the moment maps for the reduction, which must be set to 1 (since the holonomy of a flat connection around an internal edge is trivial). Similarly, every new external edge E' of ∂M formed by gluing together edges E_1, E_2 gets a \mathbb{C}^* coordinate $x_{E'} = x_{E_1} x_{E_2}$.¹⁸ Most importantly, in parallel with the reduction of boundary phase spaces, the Lagrangian $\mathcal{L}_2(M, \mathbf{t})$ is obtained by pulling the product $\mathcal{L}_2(M_1, \mathbf{t}_1) \times \mathcal{L}_2(M_2, \mathbf{t}_2)$ through the quotient in (2.15).

As an application of the gluing formula, one can fully decompose M into tetrahedra Δ_i , so that

$$(2.16) \quad \mathcal{P}_2(\partial M, \mathbf{t}) = \left[\prod_i \mathcal{P}_{\partial\Delta_i} \right] // (\mathbb{C}^*)^{\# \text{ internal edges in } M},$$

while $\mathcal{L}_2(M, \mathbf{t})$ is the reduction of a product of canonical Lagrangians (2.14). Generically (that is, for suitably refined triangulations) the resulting pair $\mathcal{L}_2(M, \mathbf{t}) \subset \mathcal{P}_2(\partial M, \mathbf{t})$ is independent of how M is decomposed. The very

¹⁸If new small annuli or tori are created during the gluing, the \mathbb{C}^* coordinates associated to them in $\mathcal{P}_2(\partial M, \mathbf{t})$ are also Laurent monomials of the \mathbb{C}^* coordinates in $\mathcal{P}_2(\partial M_1, \mathbf{t}_1) \times \mathcal{P}_2(\partial M_2, \mathbf{t}_2)$. See Appendix A.

special symplectic properties of edge-coordinates that underlie the reductions (2.15)–(2.16) were first studied in 3d hyperbolic geometry by Neumann and Zagier [35] (and later Neumann [36]), following work of Thurston [37].

2.2.2. The 3d theories. Now let us take a 3-manifold M with big-boundary triangulation \mathbf{t} . Let $d = \sum_i d_i$ be the dimension of the phase spaces associated to big boundaries, and let a, t denote the number of small annuli and tori. Let us also choose a polarization Π for $\mathcal{P}_2(\partial M, \mathbf{t}) \simeq (\mathbb{C}^*)^{2d} \times (\mathbb{C}^*)^{2a} \times (\mathbb{C}^*)^{2t}$, splitting the coordinates in each factor into canonically conjugate pairs of positions and momenta, or “electric” and “magnetic” coordinates. For big boundaries, this is equivalent to a splitting of the electric-magnetic charge lattice \mathbb{Z}^{2d} from Section 2.1.

The 3d $\mathcal{N} = 2$ superconformal theory $T_2[M, \mathbf{t}, \Pi] := T_2[M, \mathbf{t}, \Pi]$ constructed in [5] turns out to have the following basic properties.

- 1) $T_2[M, \mathbf{t}, \Pi]$ has a UV Lagrangian description as an abelian Chern-Simons-matter theory; *i.e.*, it is in class \mathcal{R} .
- 2) $T_2[M, \mathbf{t}, \Pi]$ has a manifest $U(1)^{d+a+t}$ flavor symmetry as well as a $U(1)_R$ R-symmetry. Each $U(1)$ is associated to one of the electric coordinates in the polarization Π .
- 3) $T_2[M, \mathbf{t}, \Pi]$ has a chiral operator \mathcal{O}_E for every edge E on the big boundary of M .

3a. If the edge E is electric (meaning, *e.g.*, that x_E is a monomial of purely electric coordinates in the big-boundary phase space) then \mathcal{O}_E is an ordinary chiral operator, transforming with charge +1 under the $U(1)$ flavor symmetry corresponding to E .

3b. If E has nontrivial magnetic charge, then \mathcal{O}_E exists in the presence of a magnetic-monopole background for an appropriate $U(1)$ flavor symmetry.¹⁹

Note that degenerate edges as in Figure 10 are excluded from property (3).²⁰ We will eventually see that they come with doublets of chiral operators under enhanced $SU(2)$ flavor symmetries.

¹⁹Background monopole operators like this and their anomalous dimensions were recently analyzed in [38].

²⁰From a combinatorial 3d perspective, such degenerate edges are excluded because, no matter how a 3-manifold is triangulated, it is impossible for a degenerate edge to contain only electric edges (z and not z', z'') of the individual tetrahedra. Then, as discussed in [5, Sec 4.1], it is impossible to define a corresponding operator \mathcal{O}_E .

A change of polarization $\Pi \rightarrow \Pi'$ is implemented by a symplectic transformation $g \in Sp(2N, \mathbb{Z})$, with $N = d + a + t$. In the context of quantization (which we shall touch upon momentarily) this can be extended to an affine symplectic $ISp(2N, \mathbb{Z}) \simeq Sp(2N, \mathbb{Z}) \ltimes \mathbb{Z}^{2N}$ action that includes multiplicative “shifts” of \mathbb{C}^* coordinates by $-q^{1/2}$. The affine symplectic group also acts on 3d $\mathcal{N} = 2$ SCFT’s with a $U(1)_R$ symmetry [39], providing a natural way to change the polarization of $T_2[M]$:

$$(2.17) \quad T_2[M, \mathbf{t}, g \circ \Pi] = g \circ T_2[M, \mathbf{t}, \Pi].$$

Specifically, “ T -type” elements of the symplectic group add background Chern-Simons couplings for flavor symmetries (redefining flavor currents); “ S -type” elements gauge a flavor symmetry, replacing it with a new topological $U(1)_J$ global symmetry (then flowing to the IR); while affine shifts add flavor currents to the R-current.²¹ This action implies an extension of the standard property (3a) above to the more general statement (3b).

The mathematical operation of symplectic reduction also translates nicely to 3d $\mathcal{N} = 2$ SCFT’s [5]. Basically, a quotient by a \mathbb{C}^* action as in (2.15) becomes the operation of adding an operator to the $\mathcal{N} = 2$ superpotential to break a corresponding $U(1)$ symmetry, then flowing to the IR. (For this to make sense, one must make sure to use a polarization for which the moment map of the \mathbb{C}^* action is electric.) Therefore, gluing together two manifolds to form $M = M_1 \cup M_2$ translates to taking a product of theories, changing polarization if needed, and adding superpotential operators corresponding to the new internal edges created in the gluing

$$(2.18) \quad T_2[M, \mathbf{t}, \Pi] = g \circ \left(T_2[M_1, \mathbf{t}_1, \Pi_1] \times T_2[M_2, \mathbf{t}_2, \Pi_2] \right) + \left\{ W = \sum_{E_I} \mathcal{O}_{E_I} \right\}.$$

An immediate consequence of the gluing prescription is that any three-manifold theory can be glued together from elementary tetrahedron theories

$$(2.19) \quad T_\Delta := T_2[\Delta, \mathbf{t}, \Pi_z] : \begin{cases} \text{free chiral multiplet } \phi \text{ with charges } (+1, 0) \\ \text{under } U(1)_{\text{flavor}} \times U(1)_R \text{ symmetry;} \\ \text{level } -1/2 \text{ background Chern-Simons term} \\ \text{for } (A_{\text{flavor}} - A_R). \end{cases}$$

The canonical tetrahedron polarization Π_z is chosen so that z is an electric coordinate and z'' is its magnetic conjugate. Then the operator $\mathcal{O}_z = \phi$ is

²¹See the appendix of [6] for a thorough review.

associated to the pair of opposite edges of the tetrahedron in Figure 2.2.1 labelled “ z ”. By gluing $T_2[M, \mathbf{t}, \Pi]$ from a product of tetrahedron theories, one naturally finds an abelian class- \mathcal{R} Lagrangian in the UV.




Figure 16: Adding a tetrahedron (with operator ϕ) to effect a flip on ∂M .

Another easy application of the gluing rule is to describe how a theory changes when an edge of the big-boundary triangulation flips. In 3d, a flip $F_E : \mathbf{t} \rightarrow \mathbf{t}'$ on an edge E is implemented by gluing a tetrahedron onto the big boundary (Figure 16). The tetrahedron is glued along two adjacent faces, so that a new internal edge E_I is created. Working in a polarization so that E is electric (and using a polarization for the tetrahedron so that its electric edge is glued to E), we find that the theory $T_2[M, \mathbf{t}, \Pi]$ with associated operator \mathcal{O}_E transforms to

$$(2.20) \quad \text{flip} : T_2[M, \mathbf{t}', \Pi] = T_2[M, \mathbf{t}, \Pi] \times T_\Delta + \{W = \mathcal{O}_E \phi\} .$$

The internal-edge operator we have added is $\mathcal{O}_{E_I} = \mathcal{O}_E \phi$. The new theory has an operator $\mathcal{O}_{E'} = \phi$ corresponding to the flipped edge E' , with charges exactly opposite those of \mathcal{O}_E .




Figure 17: Double-flip, a trivial operation.

Note that flipping twice is a trivial operation once we flow to the IR. Indeed, given a superpotential

$$(2.21) \quad W = \mathcal{O}_E \phi + \phi \phi',$$

where we have in mind that $\phi = \mathcal{O}_{E'}$ (the first flip) and $\phi' = \mathcal{O}_{E''}$ (the second flip), we can simply integrate out the massive fundamental field ϕ . We find that $\mathcal{O}_E = \phi' = \mathcal{O}_{E''}$, reflecting the fact that the doubly-flipped edge E'' is identical to E (Figure 17).

2.2.3. Line operators. To complete the circle of 3d ideas, let us recall how the Lagrangian $\mathcal{L}_2(M, \mathbf{t})$ arises in 3d theories.

If the theory $T_2[M, \mathbf{t}, \Pi]$ is compactified on a circle S^1_β , the twisted masses associated to every $U(1)$ flavor symmetry are complexified by Wilson lines, and can directly be identified with the position coordinates in Π . The canonically conjugate momentum coordinates in Π are vevs of the complexified moment map operators for these $U(1)$'s. It is useful to couple the 3d theory supersymmetrically to a $\mathcal{N} = 2$ abelian four-dimensional gauge theory defined on half space. Then the coordinates on $\mathcal{P}_2(\partial M, \mathbf{t})$ are vevs of half-BPS flavor Wilson line operators and dual flavor 't-Hooft lines in the four-dimensional theory. They are not independent: we can bring the four-dimensional operators on the boundary, and there the vevs must lie on the Lagrangian $\mathcal{L}_2(M, \mathbf{t})$, due to the coupling to $T_2[M, \mathbf{t}, \Pi]$.

Indeed, upon compactification on a circle, the four-dimensional gauge theory reduces at low energy to a 2d $\mathcal{N} = 4$ sigma model whose target is the complex symplectic manifold $\mathcal{P}_2(\partial M, \mathbf{t})$. The boundary condition defined by coupling to $T_2[M, \mathbf{t}, \Pi]$ constrains the low energy fields to live on a Lagrangian submanifold. The boundary condition preserves $\mathcal{N} = (2, 2)$ supersymmetry in two dimension. The generating function of the Lagrangian $\mathcal{L}_2(M, \mathbf{t})$ coincides with the the effective twisted superpotential \widetilde{W} for the two-dimensional degrees of freedom at the boundary, a function of complex masses x and complex gauge scalars σ . Thus the equations for $\mathcal{L}_2(M, \mathbf{t})$ arise as [5, 20]

$$(2.22) \quad \exp \left(x \frac{\partial \widetilde{W}}{\partial x} \right) = p \Big|_{\partial \widetilde{W} / \partial \sigma = 0},$$

where p are the momenta conjugate to x .

The line operators can be “quantized” by adding angular momentum, changing the geometry to $\mathbb{R}^2 \times_q S^1_\beta$. Then the order in which line operators are brought to the boundary to act on the three-dimensional theory

$T_2[M, \mathbf{t}, \Pi]$ matters, and conjugate pairs obey $\hat{p}\hat{x} = q\hat{x}\hat{p}$. Moreover, they satisfy a Ward identity that is a quantization of the $\mathcal{L}_2(M, \mathbf{t})$ equations. For example, for a tetrahedron the Ward identity is

$$(2.23) \quad \hat{\mathcal{L}}_\Delta: \quad \hat{z}'' + \hat{z}^{-1} - 1 \simeq 0.$$

It should be viewed as the generator of a left ideal in the operator algebra. Mathematically, the quantization of Lagrangians $\mathcal{L}_2(M, \mathbf{t})$ is uniquely defined by pulling products of $\hat{\mathcal{L}}_\Delta$'s through symplectic reduction.

It is again convenient to work with logarithmic coordinates in order to properly keep track of q (or $\hbar = \log q$) corrections. In this case, the affine extension of the symplectic action on polarizations or 3d theories is crucial — affine shifts correspond to multiplying coordinates powers of q . This is easy to understand physically, since affine shifts change the R-current and there is an R-symmetry twist in the geometry $\mathbb{R}^2 \times S^1_\beta$. Further details appear in [5] and the appendix of [6].

2.3. Half-BPS boundary conditions

The properties of the theory $T_2[M, \mathbf{t}, \Pi]$ suggest a natural way to couple it to a 4d Seiberg-Witten theory $SW_2[\mathcal{C}]$ corresponding to the big boundary of M , as a 3d boundary condition.

A half-BPS boundary condition for a supersymmetric theory generally splits supermultiplets in half (according to the broken supersymmetry), giving one half Neumann boundary conditions (N b.c.) and the other half Dirichlet (D b.c.) [40, 41]. More precisely, half the bosonic fields get N b.c. and the other half D b.c., while half the fermions are set to zero.

In the present case, we are interested in superconformal boundary conditions for a 4d $\mathcal{N} = 2$ Seiberg-Witten theory, which break SUSY to $\mathcal{N} = 2$ in 3d. The boundary conditions should preserve a $U(1)_R \subset SU(2)_R$ R-symmetry and break $U(1)_r$. The breaking pattern is thus characterized by parameters $\omega \in SU(2)_R/U(1)_R \simeq \mathbb{C}\mathbb{P}^1$ and by a phase $\zeta \in U(1)_r$.

We put the theory on $\mathbb{R}^3 \times \mathbb{R}_+$, with $x^3 \geq 0$ parameterizing \mathbb{R}_+ . The standard Neumann boundary condition for gauge fields essentially gives N b.c. to the components A_\parallel of the gauge field parallel to the boundary, and D b.c. to the component A_\perp perpendicular to the boundary. Correspondingly, the complex scalar Φ in the gauge multiplet gets split roughly into real and imaginary parts, with half N b.c. and half D b.c., so that

$$(2.24) \quad \partial_3 \operatorname{Re}(\zeta^{-1}\Phi)|_\partial = 0, \quad \operatorname{Re}(\zeta^{-1}\tau\Phi)|_\partial = 0,$$

where “ $|_{\partial}$ ” denotes restriction to the boundary $x^3 = 0$ ²². Altogether, the free boundary values of the fields A_{\parallel} and $\text{Re}(\zeta^{-1}\varphi)$ compose a 3d $\mathcal{N} = 2$ gauge multiplet. In agreement with the boundary condition on adjoint scalars, the real 3d central charge of any state that transforms under bulk symmetries must be

$$(2.25) \quad Z_{3d} = \text{Re}(\zeta^{-1}Z_{4d}).$$

In addition, each hypermultiplet Φ is split into a pair of chiral/anti-chiral fields (X, Y^{\dagger}) with well-defined $U(1)_R$ charges ± 1 and identical flavor quantum numbers. More commonly, we write this as a pair of chirals (X, Y) with $R = 1$ and opposite flavor symmetry. Then the basic boundary conditions, chosen independently for each hyper, are to give N b.c. to (the bosonic fields in) X and D b.c. to Y , or vice versa. We’ll denote this as (*cf.* [17])

$$(2.26) \quad \mathcal{B}_X : \begin{array}{l} \partial_3 X|_{\partial} = 0 \\ Y|_{\partial} = 0 \end{array} \quad \text{or} \quad \mathcal{B}_Y : \begin{array}{l} \partial_3 Y|_{\partial} = 0 \\ X|_{\partial} = 0 \end{array} .$$

From a 3d perspective, the theory on \mathbb{R}^4 has a crucial $\mathcal{N} = 2$ superpotential

$$(2.27) \quad W = \int_{\mathbb{R}_+} \partial_3 X Y = - \int_{\mathbb{R}_+} X \partial_3 Y .$$

which shows how $\partial_3 X$ appears as the F-term in the chiral multiplet whose lowest component is Y and $\partial_3 Y$ as the F-term for X . This explains the complementary pairs of boundary conditions.

The basic boundary conditions just described for a bulk $\mathcal{N} = 2$ theory T_{bulk} can be deformed in the presence of a 3d $\mathcal{N} = 2$ theory T_{bdy} on the boundary. First, any flavor symmetries of T_{bdy} may be either gauged in the bulk or identified with bulk flavor symmetries. That is, we can identify 3d background vector multiplets with boundary values of the pieces of 4d vector multiplets (dynamical or background) that have N b.c.

Thereafter, we may use boundary superpotentials to couple the boundary value of any bulk chiral X or Y that has N b.c. to a chiral operator of

²²For standard Dirichlet b.c. for the gauge fields, the role of the two parts of Φ are exchanged

T_{bdy} with dual gauge/ flavor charges:

$$(2.28) \quad W_{\text{bdy}} = X|_{\partial} \cdot \mathcal{O} \quad \text{or} \quad W_{\text{bdy}} = Y|_{\partial} \cdot \mathcal{O}' .$$

Due to the superpotential (2.27), a boundary coupling $W_{\text{bdy}} = X|_{\partial} \cdot \mathcal{O}$ actually forces the D b.c. on Y to be relaxed to $Y|_{\partial} = \mathcal{O}$, and similarly for $X \leftrightarrow Y$. Altogether, the deformed b.c. for bulk hypermultiplets in the presence of T_{bdy} may be summarized as

$$(2.29) \quad \mathcal{B}_X[\mathcal{O}] : \begin{cases} \partial_3 X|_{\partial} = F_{\mathcal{O}} \\ Y|_{\partial} = \mathcal{O} \end{cases} \quad \text{or} \quad \mathcal{B}_Y[\mathcal{O}] : \begin{cases} \partial_3 Y|_{\partial} = F_{\mathcal{O}'} \\ X|_{\partial} = \mathcal{O}' \end{cases} .$$

Now, suppose that our bulk theory is $SW_2[\mathcal{C}]$ and that our boundary theory is $T_2[M, \mathbf{t}, \Pi]$, where \mathcal{C} is a big boundary of M . The properties summarized in Sections 2.1–2.2 suggest a natural way to couple the two.

Let us choose any point u on the Coulomb branch of $SW_2[\mathcal{C}]$ (again we include masses and marginal UV couplings in ‘ u ’) and an angle θ so that the WKB triangulation $\mathbf{t}_{u,\theta}$ of \mathcal{C} agrees with \mathbf{t} . We consider the bulk theory in an electric-magnetic duality frame that agrees with the polarization Π on the big boundary. Then:

- 1) We break 4d $\mathcal{N} = 2$ supersymmetry in the bulk to a 3d $\mathcal{N} = 2$ sub-algebra characterized by the phase $\zeta = e^{i\theta}$ (and $\omega \in \mathbb{CP}^1$) and impose the basic boundary conditions (2.24) for the (electric) bulk vector multiplets.
- 2) We use the surviving 3d $\mathcal{N} = 2$ $U(1)$ gauge multiplets at the boundary to gauge d corresponding $U(1)$ flavor symmetries of $T_2[M, \mathbf{t}, \Pi]$.
- 3) The bulk theory has BPS hypers $\Phi_E = (X_E, Y_E)$ of charge γ_E for every edge E in the WKB triangulation. We impose N b.c. for the chiral halves X_E (say) of these multiplets with charge $-\gamma_E$. Then we add a coupling $\mathcal{B}_{X_E}[\mathcal{O}_E]$ to the chiral operators \mathcal{O}_E with charge $+\gamma_E$ that must exist in $T_2[M, \mathbf{t}, \Pi]$.

The third item in this list requires a little explanation. Obviously, there is no IR Lagrangian description of SW theory that can simultaneously accommodate all BPS hypermultiplets as elementary fields. At most, we can pick a maximal set of electric BPS particles in an adequate duality frame, maybe adjust the parameters so that these electric particles are much lighter than other dyonic particles, and include them in the effective Lagrangian. Then it makes sense to give them boundary conditions and couple them to the

chiral operators for electric edges. On the other hand, heavy dyonic particles behave as line defects in the bulk SW theory. It makes sense to ask if such line defects can end supersymmetrically on the boundary, where they look like monopole operators for a flavor symmetry of the 3d boundary theory. The coupling $\mathcal{B}_{X_E}[\mathcal{O}_E]$ for the corresponding edges should be interpreted as allowing such line defect to end on \mathcal{O}_E .

Away from loci where electric particles are light, we can adopt the line-defect point of view for all edge couplings. It smoothly reduces to the Lagrangian description of the coupling in the regions of parameter space with mutually local light particles. However, Argyres-Douglas regions of parameter space [19, 42], where mutually non-local particles are light, lack a complete Lagrangian description in the bulk; then we cannot hope to fully describe a boundary condition in an elementary way.

We should also note that the charges γ_E do not exhaust the charges of all BPS particles. Rather, the charges of all BPS particles in the theory can be written as a linear combination of γ_E with non-negative coefficients [21]. In a certain sense (made precise by the BPS quivers program [43, 44], [45]) all BPS particles can be thought as bound states of the ones with γ_E charges. We should probably think that the bound states of charges $\sum_E n_E \gamma_E$ will couple to operators $\prod_E \mathcal{O}_E^{n_E}$ at the boundary. This ansatz is compatible with the results of our next section, where we demonstrate that our couplings are covariant under changes of triangulation.

Finally, in order to define couplings as above, it is not really necessary for \mathbf{t} to agree with a WKB triangulation of \mathcal{C} : any triangulation whose edges are associated to a positive basis of bulk BPS states will do. However, the WKB triangulation is especially natural, and leads to a nice interpretation of boundary conditions in terms of Janus configurations (Section 3).

The couplings we have described may, in general, break the $SU(2)$ flavor symmetries of the bulk to $U(1)$'s. The $U(1)$ subgroups should be identified with $U(1)$ flavor symmetries of $T_2[M, \mathbf{t}, \Pi]$ coming from annuli of M (that attach to holes on \mathcal{C}). We will eventually argue in Section 5 that whenever a hole of \mathcal{C} is attached to an annulus (as opposed to being capped off by a small disc), full $SU(2)$ flavor symmetry of the coupled 3d-4d system is in fact restored.

2.3.1. Dissolving polarization and triangulation. Given any three-manifold M with big boundary \mathcal{C} , or possibly a disjoint union of multiple big boundaries, the above rules define a boundary coupling between $SW_2[\mathcal{C}]$ and $T_2[M, \mathbf{t}, \Pi]$. We could call the full 3d-4d theory $\mathcal{B}_{SW}[M, \mathbf{t}, \Pi]$. Quite pleasantly, however, it turns out that in the IR the coupled system depends

neither on a choice of big-boundary triangulation nor on a choice of polarization [5]. Thus, it may unambiguously be called $\mathcal{B}_{SW}[M]$.

Let us try to understand why this is true. First, looking only at gauge multiplets, consider what happens if we act with electric-magnetic duality on a combined 3d-4d system. The bulk electric-magnetic duality group $Sp(2d, \mathbb{Z})$ simply acts by changing polarization on the boundary. This is precisely how the $Sp(2d, \mathbb{Z})$ action on 3d CFT's was obtained in [39]. Therefore, the two coupled systems

$$(2.30) \quad SW_2[\mathcal{C}] - T_2[M, \mathbf{t}, \Pi] \quad \simeq \quad g \circ SW_2[\mathcal{C}] - T_2[M, \mathbf{t}, g \circ \Pi]$$

are equivalent in the IR. In one case, bulk electric $U(1)$'s gauge boundary flavor symmetries; while in the other case bulk magnetic $U(1)$'s gauge “dual” boundary flavor symmetries.




Figure 18: Passing a duality wall through the 3d-4d system.

A nice way to picture the equivalence is by starting with $SW_2[\mathcal{C}]$ coupled to $T_2[M, \mathbf{t}, \Pi]$ on a half-space, then dualizing the bulk theory on a slightly “smaller” half-space, as in Figure 18. This is equivalent to coupling the magnetic theory on the far left to an appropriate 3d abelian Chern-Simons theory on a BPS domain wall, then coupling the domain wall to a slice of the original electric theory on its right. For example, for a standard $g = S = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$ action, the duality domain wall is “theory” of two background $U(1)$ flavor multiplets with a mixed Chern-Simons coupling $A'dA$. After flowing to the IR, the duality domain wall collides with the boundary condition, effecting a change of polarization on the boundary.

Keeping track of hypermultiplets and their couplings to chiral operators in the electric-magnetic duality requires a little extra care. For fixed triangulation \mathbf{t} , the duality action changes the set of edges that are electric, and thus the subset of couplings that can (and do) appear in a bulk-boundary superpotential.

For example, suppose that an electric hyper (X, Y) of $SW_2[\mathcal{C}]$ couples to a chiral operator \mathcal{O} of $T_2[M, \mathbf{t}, \Pi]$, with $W_{\text{bdy}} = X|_{\partial} \mathcal{O}$. Let us dualize the

system by passing an S duality wall through the bulk and colliding it with the boundary, as in Figure 18. In addition to carrying an abelian Chern-Simons theory, the duality wall has the property that a magnetic line defect of $S \circ SW_2[\mathcal{C}]$ on the left can end on a chiral operator $X|_{\text{wall}}$ on the right — put differently, that magnetic BPS particles passing through from the left become electric particles on the right. After colliding the duality wall with the boundary, we simply find that magnetic line defects in the bulk can end on a new chiral operator \mathcal{O}' in $T_2[M, \mathbf{t}, S \circ \Pi]$.

Similarly, when we collide an S -wall with the boundary, any magnetic line defect of $SW_2[\mathcal{C}]$ that can end on both the boundary and the duality wall becomes trapped. It combines with the magnetic operator \mathcal{O}_m of $T_2[M, \mathbf{t}, \Pi]$ that it ends on to form a new electric operator \mathcal{O}'_m of $T_2[M, \mathbf{t}, S \circ \Pi]$. This operator \mathcal{O}'_m is coupled to an electric hyper (X', Y') of $S \circ SW_2[\mathcal{C}]$. Such mechanisms, involving line defects, demonstrate that bulk-boundary superpotential couplings transform covariantly, as desired, under a change of polarization.

To see that the system $\mathcal{B}_{SW}[M, \mathbf{t}, \Pi]$ is independent of triangulation as well as polarization, we consider how the flip of an edge in \mathbf{t} acts on our standard couplings [17]. Suppose that the flipped edge E is electric (we can adjust the polarization to make it so), and that before the flip half of a 4d hyper (X, Y) is coupled to \mathcal{O}_E via the superpotential $W_{\text{bdy}} = X|_{\partial} \mathcal{O}_E$. After the flip, we introduce a new 3d chiral ϕ (from a tetrahedron theory) and the superpotential becomes

$$(2.31) \quad W_{\text{bdy}} = Y|_{\partial} \phi + \phi \mathcal{O}_E.$$

Note how the flip effectively switches N and D b.c. for the bulk hyper, so that the field Y , which carries (minus) the 4d charge of the flipped edge $-\gamma_{E'} = \gamma_E$ can couple to the new boundary operator $\mathcal{O}_{E'} = \phi$. Concurrently, due to the bulk superpotential (2.27) we find modified D b.c. $X|_{\partial} = \phi$. The coupling (2.31), however, has made the 3d field ϕ massive. Flowing to the IR we may integrate it out, and simply get back to the original coupling $W_{\text{bdy}} = X|_{\partial} \mathcal{O}_E$. Therefore, $\mathcal{B}_{SW}[M, \mathbf{t}, \Pi] = \mathcal{B}_{SW}[M, \mathbf{t}', \Pi]$.

There is an interesting subtlety at play here, concerning the charges of magnetic chiral operators before and after the flip. For example, in the present A_1 case the flip of an edge of the triangulation modifies the charges carried by the nearby edges in a specific way. Physically, adding a chiral field charged under a flavor symmetry modifies the quantum numbers of monopole operators for that flavor symmetry because of the quantization of fermionic zero-modes. It would be interesting to study these charge shifts in




Figure 19: Flipping the 3d-4d system.

a general theory, and to reproduce the expected transformation of charges of fundamental BPS particles, encoded by a tropical cluster mutation.

We can picture the action of a flip using duality walls, in much the same way that we understood changes of polarization. Namely, we note that the tetrahedron theory T_Δ can trivially be inserted as a wall in any theory $SW_2[C]$, coupling to some hypermultiplet (X, Y) via

$$(2.32) \quad W_{\text{wall}} = Y|_{\partial} \phi + \phi Y'|_{\partial}.$$

Here Y, Y' denote the half-hypers with Neumann b.c. on the two sides of the wall; their superpartners have Dirichlet b.c. $X|_{\partial} = \phi = X'|_{\partial}$. After integrating out ϕ in the IR, this wall becomes trivial. By inserting a T_Δ wall into a theory on a half-space (Figure 19) and colliding it with the boundary, we effectively perform a boundary flip.

We have shown that the coupled system $\mathcal{B}_{SW}[M]$ only depends on the topology of M and its separated big/small boundary. Nevertheless, it should be clear that any Lagrangian realization of this theory does require a choice of \mathbf{t} and Π . For example, if we ever want to make the 4d bulk theory infinitely weakly coupled, so that we just leave behind a dynamical 3d boundary theory, we must choose an electric-magnetic duality frame Π . These choices then appear as data for the 3d theory $T_2[M, \mathbf{t}, \Pi]$ that remains.

3. Janus domain walls in 4d $\mathcal{N} = 2$ gauge theories

In this section, we would like to review the idea of BPS Janus configurations for four-dimensional $\mathcal{N} = 2$ theories. Janus configurations provide the simplest examples of domain walls and boundary conditions associated to 3-manifolds. Moreover, in special limits, one can sometimes extract the field content of an effective 3d boundary theory from the data of a Janus

configuration — thus confirming our more abstract, combinatorial constructions of 3-manifold theories. For example, in Section 3.3 we will use a Janus configuration to rediscover the basic boundary couplings in a 3d-4d system $\mathcal{B}_{SW}[M]$. Later in Section 4 we will introduce Janus configurations for RG walls.

3.1. General concepts

A BPS Janus configuration for a four-dimensional $\mathcal{N} = 2$ gauge theory is a modification of the Lagrangian (assuming one exists) that allows the gauge couplings τ and mass parameters m to vary in an arbitrary way along a direction x^3 , preserving three-dimensional $\mathcal{N} = 2$ supersymmetry. The three-dimensional superalgebra has real central charges that are the real part of the four-dimensional central charges (2.25). (We'll set $\zeta = 1$ for simplicity.) In a supersymmetric vacuum, the 3d central charges must be constant. Then the equations

$$(3.1) \quad \partial_3 \operatorname{Re} Z_\gamma^{4d}(u(x^3), m(x^3), \tau(x^3)) = 0$$

fix the allowed evolution of the Coulomb-branch moduli u as a function of x^3 . Indeed, the number of real constraints (3.1) equals the real dimension of the Coulomb branch.

Janus configurations have a well-defined limit to a domain-wall configuration where the jumps in m, τ occurs suddenly. Furthermore, it is useful to observe that many protected quantities are homotopy invariant, unaffected by continuous deformations of the path $m(x^3), \tau(x^3)$. This includes sphere partition functions and indices, and (we expect) the infrared SCFT limits of the 3d theories themselves that become trapped on domain walls.

In [8] a special case of Janus domain walls was considered, the “R-flow” where the variation of m, τ was arranged in such a way that the relative order of the phases of the Z_γ^{4d} central charges would be constant, and the Z_γ^{4d} would go to infinity at large positive or large negative x^3 . In this and later sections we will not impose such a constraint, but we will still take inspiration from the R-flow analysis.

The notion of Janus domain wall is intimately related to the notion of S-duality wall. If we take a Janus configuration that flows from a weak coupling region to a region where an S-dual weak coupling description exists, it is natural to act with S-duality on a half-space, and have a domain wall interpolating between different weakly coupled descriptions of the theory. Up

to D-term deformations, we will have a duality wall for each element of the S-duality groupoid of the theory, and duality walls will compose appropriately by collision.

As the four-dimensional theories on the two sides are weakly coupled, we can sensibly talk about degrees of freedom living at the wall, and describe the duality wall as a specific 3d $\mathcal{N} = 2$ SCFT coupled to the two dual descriptions of the same four-dimensional theory. The canonical example is (mass deformed) $T[G]$, a 3d theory which appears on the S-duality wall between $\mathcal{N} = 2^*$ SYM with a gauge groups G and its Langlands dual ${}^L G$. For $G = SU(2)$, $T[SU(2)]$ is simply $\mathcal{N} = 4$ SQED with $\mathcal{N}_f = 2$, a self-mirror theory with two $SU(2)$ flavor symmetries, acting on the Coulomb and Higgs branches respectively. The two bulk theories couple to these two $SU(2)$ flavor symmetries.

The correspondence between Janus domain walls and S-duality walls is an exact UV statement. It is also useful to ask in general how a Janus configuration would look in the IR Seiberg-Witten description of the theory. There are a few things that can be said in complete generality (in particular, without specializing to theories of class \mathcal{S}), and we would like to point out some interesting facts and open questions.

The abelian gauge field Lagrangian simply has a profile of the IR gauge coupling determined by $u(x^3)$ and $m(x^3), \tau(x^3)$. The massive BPS particles have a more interesting behavior: for generic x^3 , they are not BPS anymore, as their mass $|Z_\gamma|$ is larger than the 3d central charge $\text{Re } Z_\gamma^{4d}$. But at the special locations where $\text{Im } Z_\gamma^{4d} = 0$, a 4d BPS particle has a chance to be trapped, and behave as a 3d BPS particle. We expect that a 4d BPS hypermultiplet of charge γ , for example, will give rise to a 3d chiral multiplet of charge $\pm\gamma$, depending on the sign of $\partial_3 \text{Im } Z_\gamma$. It would be interesting to verify this statement with a detailed calculation.

We can use these observations to argue that the 3d $\mathcal{N} = 2$ theories produced from Janus domain walls in 4d $\mathcal{N} = 2$ gauge theories should always have a mirror description as abelian Chern-Simons-matter theories, *i.e.*, theories of class \mathcal{R} ! Indeed, the IR description of the Janus configuration can be easily converted to an abelian Chern-Simons-matter description. The chiral matter arises from the trapped 4d BPS particles. We need to make each BPS particle locally electrically charged. This can be accomplished by picking an appropriate electric-magnetic duality frame in the neighborhood of the wall $\text{Im } Z_\gamma^{4d} = 0$, such that γ is electric. Changes in the electric-duality frame for the four-dimensional gauge fields are implemented by abelian duality walls, which add appropriate Chern-Simons couplings to the abelian gauge fields.

Thus the matter content and the gauge Lagrangian of the 3d theories are easily derived in terms of the x^3 dependence of the Z_γ . The latter may be hard to derive from $\tau(x^3)$ though, which was one reason for the restrictions imposed on R-flows in [8]. Furthermore, there is a final, crucial ingredient to be derived: the superpotential couplings. From the four-dimensional point of view, these should arise from non-local instanton effects that allow the chiral matter at different locations to interact. We do not know how to derive such contributions in general, and we feel it is a very interesting open problem. We owe much of our intuition on such instantonic processes to an ongoing project on Janus domain walls in 2d field theories [46].

It is easy to argue that such instanton effects must be present: they implement homotopy invariance for the Janus configuration. We can give two illuminating examples. First, consider a homotopy that interpolates between a Janus configuration where, for some γ , $\text{Im } Z_\gamma^{4d}$ is positive in a certain region, and a Janus configuration where $\text{Im } Z_\gamma^{4d}$ becomes negative and then immediately positive again in the same region, so that we have two locations with $\text{Im } Z_\gamma^{4d} = 0$, hosting two chiral particles ϕ_\pm of opposite charge. It is clear that in order to have homotopy invariance, the second configuration should include a superpotential coupling $\phi_+\phi_-$, which makes the particles disappear in the IR. Such a superpotential may arise from an instanton process, where a 4d particle moves from one of the two locations.

A second example is a homotopy that crosses a basic wall of marginal stability for four-dimensional BPS particles, interpolating between a configuration with two relevant BPS particles, of charges γ_1 and γ_2 , giving rise to two 3d chirals, and a configuration with three relevant 4d BPS particles, of charges γ_1 , γ_2 and $\gamma_1 + \gamma_2$, giving rise to three 3d chirals. As observed in [8], a natural way to insure homotopy invariance is to have a superpotential coupling of the XYZ type between the three chirals, related by the basic 3d mirror symmetry to the configuration of two chirals.

A full analysis of this setup would hopefully give the broadest possible generalization of our results. Leaving that for future work, we specialize now to four-dimensional theories in class \mathcal{S} .

3.2. Janus configurations for class \mathcal{S}

We briefly review a few facts about the 6d interpretation of the IR description of Januses, along the lines of [8].

The Seiberg-Witten curve Σ of a 4d theory $SW_K[\mathcal{C}]$ of class \mathcal{S} is naturally expressed as a K -fold branched cover of the UV curve \mathcal{C} , as discussed in Section 2.1. In turn, a Janus configuration for $SW_K[\mathcal{C}]$ amounts to varying

the cover $\Sigma \xrightarrow{K} \mathcal{C}$ in the x^3 direction, so as to sweep out a cover

$$(3.2) \quad \widetilde{M} \xrightarrow{K} M$$

of a 3-manifold M that topologically is equivalent to $\mathcal{C} \times \mathbb{R}$.

As the 4d central charges are periods of the Seiberg-Witten differential λ_{SW} , the condition (3.1) implies that the real part of $\partial_3 \lambda_{\text{SW}}$ is exact, *i.e.*

$$(3.3) \quad \partial_3 \text{Re } \lambda_{\text{SW}} = d\rho$$

for some real function ρ . In other words, the 1-form

$$(3.4) \quad v = \text{Re } \lambda_{\text{SW}} + \rho dx^3$$

is closed. Notice that v can be viewed either as a single-valued 1-form on \widetilde{M} , or a multi-valued 1-form on M .

The 1-form v parameterizes the Coulomb branch of the twisted 6d SCFT on M . Indeed, the twisting makes three of the scalar fields of the abelian 6d theory into a 1-form v' in the same cohomology class as v . Given a metric g on M , the BPS equations of the 6d theory imply that v' is harmonic

$$(3.5) \quad d *_g v' = 0,$$

i.e. v' is the harmonic representative of $[v] \in H^1(M, \mathbb{R})$. The three components of v' parameterize the deformations of the cover \widetilde{M} in the cotangent bundle T^*M . In regions of M where the cover becomes independent of x^3 , so that $\partial_3 v' = 0$, we simply have $v = v' = \text{Re } \lambda_{\text{SW}}$, and the holomorphic Seiberg-Witten form λ_{SW} on Σ can be reconstructed from its real part.

We can also understand the harmonic v' in terms of an M-theory compactification. If we wrap K M5 branes on the supersymmetric cycle $M \times \mathbb{R}^3$ in the 11d geometry $T^*M \times \mathbb{R}^5$, with any given metric g on M , then at low energies the branes may separate in the fibers of the cotangent bundle, recombining (say) into a K -fold cover \widetilde{M} of M . In order to preserve supersymmetry, this cover, just like M , must be a special Lagrangian submanifold in T^*M . Deformations of special Lagrangians, however, are precisely parameterized by harmonic 1-forms — in the present case, the 1-form v' .

If we compactify the entire setup on a circle in \mathbb{R}^3 , we can even find a non-abelian origin for v' . We end up with D4 branes wrapping $M \times \mathbb{R}^2$, which support a twisted 5d super-Yang-Mills theory. The BPS equations are 3d analogues of Hitchin equations along M , just as in Section 2.1.1. They

can be written compactly as

$$(3.6) \quad [D_i, D_j] = 0, \quad g^{ij}[D_i, D_j^\dagger] = 0,$$

where $D = d + \mathcal{A}$ is a covariant derivative on M formed from the complexified $SL(K)$ connection $\mathcal{A} = A + i \operatorname{Re} \varphi$ as in (2.3). In components, the equations read $F_A = [\varphi, \varphi]$ and $d_A \varphi = d_A * \varphi = 0$, and thus imply that the adjoint-valued 1-form φ is covariantly harmonic. If we are able to diagonalize φ locally — corresponding to a well-defined separation of the D4 branes in T^*M — then the eigenvalues of φ become a standard (multivalued) harmonic 1-form v' on M .

Now let us follow the analysis of [8] further, recast in a field-theory language, to describe the BPS content of an effective theory $T_K[M]$. First, recall that four-dimensional BPS particles are represented by (webs) of strands of 6d BPS strings. In 2d compactifications, the BPS strings follow trajectories of constant phase of λ_{SW} , ending at branch points where two sheets of the branched cover $\Sigma \rightarrow C$ meet. These trajectories lift to 1-cycles in Σ , *cf.* Figure 9.

In a 3d compactification, the BPS strings must follow trajectories on which the harmonic 1-form v' restricts to the volume form, in a given background metric. These trajectories end on the branch-lines of the cover $\widetilde{M} \rightarrow M$, swept out by the x^3 evolution of branch points on \mathcal{C} . They lift to minimum-volume, or “calibrated” 1-cycles on \widetilde{M} . In the limit where \mathcal{C} varies slowly, the minimum-volume condition reduces to the $\operatorname{Im} Z_\gamma^{4d} = 0$ condition for a trapped BPS state. The wrapped BPS strings then give rise to 3d chiral multiplets.

An important payoff of the 6d description is that the instantons that give rise to superpotential terms can also be described in terms of BPS strings that trace out minimum-volume discs in M . These discs have boundary along branch lines and along the minimum-volume trajectories that gave rise to BPS states. In *principle*, by fully analyzing the allowed configurations of 6d BPS strings, one can recover a full abelian Chern-Simons-matter description of a domain-wall theory.

In practice, in order to be able to identify BPS cycles, one must judiciously choose the profile $m(x^3), \tau(x^3)$, or the background metric on M , or both. In particular, it is useful to arrange that branch lines be well separated throughout most of M , and only coming close together briefly. Then each region of M where a pair of branch lines pass near each other supports a

single BPS chiral — and can be shown to map to one of the combinatorial tetrahedra²³ reviewed in Section 2. Of course, there may be many ways to put the branch lines in such a “nice” configuration, just as there are many ways to triangulate a 3-manifold; the different Chern-Simons matter Lagrangians obtained should all be related by 3d mirror symmetry, flowing to the same IR SCFT.

3.3. Seiberg-Witten walls and couplings

To be a little more specific, let us review the basic IR Janus configuration of [8], and show how it generalizes to a T_Δ duality wall in a 3d-4d system, with couplings as in Section 2.3.

We can model a local region of M where a pair of branch lines come close together by evolving the Seiberg-Witten curve

$$(3.7) \quad \Sigma : \quad w^2 = -z^2 + m, \quad \lambda_{SW} = w dz.$$

The curve is fibered over the z -plane, which represents an open neighborhood of \mathcal{C} . This happens to be the Seiberg-Witten curve of the A_1 Argyres-Douglas theory [19].

For fixed m , there is an obvious trajectory of constant phase that connects the branch points at $z = \pm\sqrt{m}$, and lifts to a 1-cycle γ in Σ . Thus the 4d theory $SW_2[\mathcal{C}]$ has a BPS hypermultiplet $\Phi = (X, Y)$ of central charge

$$(3.8) \quad Z_\gamma^{4d} = \frac{1}{\pi} \oint_\gamma \lambda_{SW} = \frac{2}{\pi} \int_{-\sqrt{m}}^{\sqrt{m}} \lambda_{SW} = m,$$

and mass $|m|$. The hyper is charged under a $U(1)$ flavor symmetry (which could be gauged in the global Seiberg-Witten theory $SW_2[\mathcal{C}]$).

We may form a Janus configuration by giving m a profile

$$(3.9) \quad m(x^3) = m_0 + ix^3,$$

for fixed real m_0 . This is an R-flow configuration. There is a trapped 3d BPS chiral $\phi \equiv X$ (say) at $x^3 = 0$, where $\text{Im } Z_\gamma^{4d} = 0$. Its 3d central charge is m_0 . Since $|m| \rightarrow \infty$ as $x^3 \rightarrow \pm\infty$, the 4d bulk theory completely decouples on either side of the domain wall and we are left with an effective 3d theory T_Δ . If we had set $m(x^3) = m_0 - ix^3$ instead, we would have obtained a 3d chiral $\phi' \equiv Y$ of opposite charge, and 3d central charge $-m_0$.

²³For $K > 2$, these regions maps to octahedra in the K -decomposition of a tetrahedron, as described in [6].




Figure 20: Left: the branch lines in a double cover $\widetilde{\Delta}$ of the tetrahedron. Right: schematic 2d slice of the cover $\widetilde{\Delta}$ and the minimal-volume cycle giving rise to a 3d chiral ϕ .

The Janus configuration (3.9) leads to the local geometry $\widetilde{M} \rightarrow M$ represented by a tetrahedron, with four branch points on its boundary (one on each face) connected pairwise by branch lines, as in Figure 20. It is helpful to note that the WKB triangulation of \mathcal{C} in the neighborhood of our two branch points looks like a square (see Figure 11 on page 155). Indeed, this is triangulation of the UV curve of the A_1 Argyres-Douglas theory, which is just a disc with an irregular singularity/puncture on its boundary [18]. There are two choices for how the square is triangulated, one appropriate for $\text{Im } m > 0$ and one for $\text{Im } m < 0$. The 3d tetrahedron flips the triangulation, as in Figure 16.

By introducing a more complicated profile for the Janus configuration

$$(3.10) \quad m(x^3) = m_0 + i[c - (x^3)^2],$$

which is no longer an R-flow, but still has $|m| \rightarrow \infty$ asymptotically, we find more interesting trapped 3d theories. Now for $c < 0$ the theory has no trapped 3d particles (it is impossible to satisfy $\text{Im } m = 0$); while for $c > 0$ there are two BPS chirals ϕ and ϕ' of opposite charge, because the locus $\text{Im } m = 0$ is crossed twice. In the latter case, there is also an instanton that mediates an interaction between these chirals coming from a 6d BPS string that wraps an annulus in \widetilde{M} (a disc in M), as in Figure 21. The instanton generates the 3d superpotential

$$(3.11) \quad W = \phi\phi',$$

so that ultimately, in the IR, the 3d theory is again empty. This example demonstrates the basic homotopy invariance that was discussed in Section 3.1.




Figure 21: Schematic 2d slices of the covers \widetilde{M} for flows corresponding to $m(x^3) = m_0 + i[c - (x^3)^2]$ with $c > 0$ (left) and $m(x^3) = m_0 + i \tanh(x^3)$ (right).

Finally, suppose that we do *not* send $|m| \rightarrow \infty$ asymptotically, so that we do not decouple the 4d bulk. We could have a profile

$$(3.12) \quad m(x^3) = m_0 + i \tanh(x^3).$$

We cross $\text{Im } m = 0$ once, so we get a 3d BPS chiral $\phi \simeq Y|_{x^3=0}$, but we still keep the 4d hypermultiplet in the asymptotic regions $x^3 \rightarrow \pm\infty$. Let us call the 4d hyper $\Phi = (X, Y)$ at $x^3 \ll 0$ and $\Phi' = (X', Y')$ at $x^3 \gg 0$. Now in the 6d theory there are two BPS strings that wrap annuli with one boundary on the minimal-volume cycle γ , and another boundary in an asymptotic region. In a limit where the jump in the Janus configuration happens instantaneously, we find two copies of the 4d theory $SW_2[\mathcal{C}]$ on half-spaces $x^3 < 0$ and $x^3 > 0$, along with bulk-boundary superpotential couplings generated by the string-instantons:

$$(3.13) \quad W_{\text{wall}} = X|_{x^3=0^-} \phi + \phi X'|_{x^3=0^+}.$$

This are precisely the kind of couplings we described combinatorially in Section 2.3, and specifically correspond to the T_Δ wall discussed around (2.32).

4. RG walls

We now turn to RG domain walls and the framed 3-manifolds that give rise to them. The basic manifold M_{RG} defining an RG wall for pure $SU(2)$ Seiberg-Witten theory will turn out to play a crucial role in the construction

of more general RG and S-duality walls, and the general analysis of enhanced flavor symmetry for theories of class \mathcal{R} .

One way to define an RG domain wall is by using Janus configurations. For example, in a 4d $\mathcal{N} = 2$ theory with an asymptotically free gauge group, we can fix an energy scale μ at which to observe the theory and vary the strong coupling scale $\Lambda(x^3)$ from zero to infinity relative to μ (Figure 22), while preserving the BPS condition (3.1). At $x^3 \ll 0$, with $\Lambda \ll \mu$, the theory is effectively non-abelian; while at $x^3 \gg 0$, with $\Lambda \gg \mu$, the theory will best be described as an abelian Seiberg-Witten theory far out on its Coulomb branch. We will quantify the latter claim in Section 4.1.




Figure 22: Schematic Janus configuration corresponding to an effective RG flow.

More generally, one can consider any 4d $\mathcal{N} = 2$ theory and vary a combination of marginal UV couplings $\tau(x^3)$, strong-coupling scales $\Lambda(x^3)$, and masses $m(x^3)$, in such a way that at $x^3 \ll 0$ the theory is near the origin of the Coulomb branch for a chosen set of gauge groups, while at $x^3 \gg 0$ the theory is far out on the Coulomb branch. In this case, what we are calling an RG wall might also be termed an “abelianization” wall.

In the limit that the jump in Λ (or other parameters) occurs very quickly, and effective gauge couplings become very weak at $x^3 < 0$ and $x^3 > 0$, one might hope to trap a well-defined 3d $\mathcal{N} = 2$ theory at $x^3 = 0$. The 3d theory should have both abelian and non-abelian flavor symmetries to allow it to couple to the 4d bulk. Moreover, it should also have electric and magnetic chiral operators that couple to bulk BPS hypermultiplets on the abelian side, as in Section 2.3. On the non-abelian side, we similarly expect to find additional chiral operators in non-trivial flavor multiplets that couple to the bulk hypermultiplets in the UV Lagrangian.

In Section 4.2 we will argue on general grounds that, in the infrared, the 3d RG-wall theory associated to the breaking $G \rightarrow \mathbb{T}$ of a gauge group to its maximal torus is roughly a G/\mathbb{T} sigma-model. For example, for pure $SU(2)$ Seiberg-Witten theory in the bulk, we expect a 3d $\mathbb{C}P^1$ sigma-model. In this description, however, the abelian flavor symmetries of the 3d theory are not entirely manifest.

In Sections 4.3–4.4, we will introduce a fundamental framed 3-manifold M_0 , and show that the 3d theory $T_2[M_0]$ derived from it combinatorially has all the right properties to be the RG-wall theory for pure $SU(2)$ Seiberg-Witten theory in the bulk. Later, we will look at more complex examples of theories with flavors.

4.1. Janus attractors

The best way to understand some of the expected properties of RG Janus configurations is to study a simple, concrete example. We focus here on pure $SU(2)$ $\mathcal{N} = 2$ gauge theory in four dimensions.

Recall that this is a 4d theory of class \mathcal{S} , whose UV curve \mathcal{C} is a sphere with two irregular punctures of the mildest possible type [2, 18]. One might also describe \mathcal{C} as an annulus. The square of the Seiberg-Witten differential is

$$(4.1) \quad \lambda_{\text{SW}}^2 = \left(\frac{\Lambda^2}{z} + 2u + \Lambda^2 z \right) \frac{dz^2}{z^2},$$

with strong-coupling scale Λ and Coulomb modulus $u = \langle \text{Tr } \Phi^2 \rangle$. Here the z -plane \mathcal{C} has rank 1/2 irregular singularities at $z = 0$ and $z = \infty$.

The abelian charge lattice of $SW_2[\mathcal{C}]$ consists only of gauge charges. We call the fundamental magnetic, dyonic, and electric charges γ_m , γ_d , and γ_e , with

$$(4.2) \quad \gamma_m + \gamma_d = 2\gamma_e, \quad \langle \gamma_m, \gamma_d \rangle = 2.$$

BPS states only carry even electric charge, so the BPS lattice is generated by γ_m, γ_d alone. In the u -plane, there are two stability chambers separated by (roughly) an ellipsoidal curve of marginal stability as in Figure 23, defined by the condition

$$(4.3) \quad \arg a = \arg a_D =: \theta,$$




Figure 23: The u -plane and line of marginal stability, for $\Lambda = 1$.

where a and a_D are the electric and magnetic central charges. The point on the curve of marginal stability where (4.3) holds for fixed θ lies approximately at polar angle $2\theta - \pi$.




Figure 24: WKB triangulation of the annulus \mathcal{C} , with branch points at $z = -\tilde{u} \pm \sqrt{\tilde{u}^2 - 1}$, $\tilde{u} := u/\Lambda^2$.

All WKB triangulations $t_{u,\theta}$ of the curve \mathcal{C} look identical topologically — there really exists only one triangulation for the annulus, shown in Figure 24. The two vertices of the triangulation lie on the S^1 boundaries of the annulus, in accordance with the fact that there are rank-1/2 irregular singularities there. There are two non-boundary edges in the triangulation, which separate the two branch points of λ_{SW} . If u is inside the curve of marginal stability, the two non-boundary edges are simply labelled by charges $\pm\gamma_m$ and $\pm\gamma_d$ (the precise assignment depends on θ). These charges correspond to magnetic and dyonic cycles in the double cover $\Sigma \rightarrow \mathcal{C}$. When u is outside the curve of marginal stability, the edges of the WKB triangulation can also

correspond to pairs $\pm((n + 1)\gamma_m + n\gamma_d, n\gamma_m + (n - 1)\gamma_d)$, $n \in \mathbb{Z}$, reflecting the fact that there are now an infinite number of BPS states. Notice that flips of the WKB triangulation, which do not affect its topological type, do change n .

We could distinguish “different” triangulations by holding the S^1 boundaries of \mathcal{C} fixed, and not allowing them to rotate with respect to each other. Then performing pairs of consecutive flips is equivalent to twisting the annulus by full turns.

Now let us imagine varying $\Lambda(x^3)$ from zero to infinity, while keeping $\text{Re}(e^{-i\theta}a)$ and $\text{Re}(e^{-i\theta}a_D)$ fixed, so as to build an RG Janus configuration. We allow ourselves an extra phase $\zeta = e^{i\theta}$ in selecting which 3d $\mathcal{N} = 2$ superalgebra to preserve. The central charges can be written as

$$(4.4) \quad \begin{aligned} a(u, \Lambda) &= \int_{\gamma_e} \lambda_{\text{SW}} = \Lambda f(u/\Lambda^2), \\ a_D(u, \Lambda) &= \int_{\gamma_m} \lambda_{\text{SW}} = \Lambda f_D(u/\Lambda^2) \end{aligned}$$

for some (locally) holomorphic functions f, f_D . It follows generically from this that no matter what values of a, a_D we take at $x^3 \ll 0$ (when $\Lambda \ll \mu$), after sending $\Lambda/\mu \rightarrow \infty$ we will find $|a|, |a_D| \sim \Lambda \gg \mu$. Thus an RG Janus configuration necessarily forces the theory far onto the Coulomb branch.

We can actually do much better, and identify precisely where on the Coulomb branch we land. In terms of the dimensionless variable $\tilde{u} := u/\Lambda^2$, we want to find $\tilde{u}(x^3 = \infty)$. Writing $\partial_3 a = \Lambda f'(\tilde{u}) \partial_3 \tilde{u} + f(\tilde{u}) \partial_3 \Lambda$ and similarly for a_D , we easily obtain from the Janus condition $\partial_3 \text{Re} \zeta^{-1}(a, a_D) = 0$ that

$$(4.5) \quad \partial_3 \begin{pmatrix} \text{Re} \tilde{u} \\ \text{Im} \tilde{u} \end{pmatrix} = \begin{pmatrix} -\text{Re}(\zeta^{-1} f') & \text{Im}(\zeta^{-1} f') \\ -\text{Re}(\zeta^{-1} f'_D) & \text{Im}(\zeta^{-1} f'_D) \end{pmatrix}^{-1} \begin{pmatrix} \text{Re}(\zeta^{-1} f) \\ \text{Re}(\zeta^{-1} f_D) \end{pmatrix} \partial_3 \log \frac{\Lambda}{\mu}.$$

This flow equation has an obvious fixed point given by $\text{Re}(\zeta^{-1} f(\tilde{u})) = \text{Re}(\zeta^{-1} f_D(\tilde{u})) = 0$, which turns out to be the unique fixed point in the \tilde{u} plane, and is attractive for increasing Λ . Moreover, by comparison to (4.3), this fixed point must lie precisely on the curve of marginal stability, with phase $\zeta = e^{i\theta}$! Thus as $\Lambda \rightarrow \infty$, all flows are attracted to marginal stability.

For example, when $\zeta = \pm 1$, the attractor point is at the dyon point $\tilde{u} = -1$ or $u = -\Lambda^2$. The exact flows to the dyon point are shown in Figure 25. Since f is nonvanishing here, we necessarily have $|a| \sim \Lambda |f(-1)| \sim \Lambda$ as $\Lambda \rightarrow \infty$. Thus the theory $T_2[\mathcal{C}] \rightarrow \text{SW}_2[\mathcal{C}]$ is abelianized as promised. The attractor moreover identifies a natural $Sp(2, \mathbb{Z})$ abelian duality frame for




Figure 25: Janus flows to the dyon point.

$SW_2[\mathcal{C}]$: it is the frame in which the light dyon becomes a fundamental “electric” hypermultiplet.

Some basic aspects of this RG flow attractor mechanism remain true for any domain wall that scales to infinity in a uniform way all the dimensional couplings of a generic $\mathcal{N} = 2$ theory (masses and asymptotically free gauge couplings). If we write all coupling as $\Lambda_i = c_i \Lambda (x^3)^{d_i}$ in terms of constant dimensionless ratios c_i , and the Coulomb branch parameters in terms of dimensionless parameters as $u_i = \tilde{u}_i \Lambda (x^3)^{D_i}$, we can write the periods as $Z_\gamma = \Lambda z_\gamma(\tilde{u}, c)$ and the BPS condition will force the flow to be attracted to the point $\text{Re}(\zeta^{-1} z_\gamma) = 0$.

This is a rather interesting point, which plays a key role in the recent work by Kontsevich [45] on wall-crossing. It is the unique fixed point for the gradient flow of an interesting function on the Coulomb branch. The useful role in wall-crossing follows from the observation that the central charge of all BPS particles is aligned at this point, and thus all BPS particles can be seen as marginal bound states of a basic set of stable objects.

In a sense, this property supports our strategy to build the couplings of domain wall theories to the IR Seiberg-Witten theory by specifying only the couplings to the basic set of stable objects, which is in correspondence to the edges of the WKB triangulation.

4.2. Trapped Goldstone bosons

We can use a simple thought exercise to guess what 3d degrees of freedom might be trapped on an RG domain wall. In order to isolate a 3d theory, we want to send the dynamical 4d gauge couplings to zero on both sides of the wall — for a gauge group G on the UV side and its maximal torus \mathbb{T} on the IR side. Suppose that we do this first on the UV side, $x^3 \leq 0$. Then the G gauge symmetry becomes a flavor symmetry. As we move onto the Coulomb branch on the IR side, this *flavor* symmetry is broken to \mathbb{T} , and we might expect to find trapped Goldstone bosons at the wall itself, parametrizing a G/\mathbb{T} moduli space.

We can make this a little more concrete by starting with nonabelian G gauge theory on the whole 4d space, and replacing the region $x^3 < 0$ with a half-BPS Dirichlet boundary condition for the gauge fields. This should have the same effect as sending the UV gauge coupling to zero there. As reviewed in Section 2.3, the basic supersymmetric Dirichlet boundary condition pairs the three components A_{\parallel} of the gauge field parallel to the boundary with the real part of Higgs field Φ , setting²⁴

$$(4.6) \quad A_{\parallel}|_{\partial} = 0, \quad \text{Re } \Phi|_{\partial} = 0,$$

while giving Neumann b.c. to $\text{Re } \tau_{UV}\Phi$. (For clarity, we will just set $\zeta = 1$ here.)

In order to find the effective degrees of freedom on the RG wall, we ask: *how does the UV Dirichlet boundary condition look in the IR?* We can move onto the Coulomb branch by turning on the Cartan part of the Higgs field at infinity, compatibly with the with Neumann-Dirichlet boundary conditions at $x^3 = 0$. In the 4d bulk ($x^3 > 0$), this is a standard Higgsing, and Goldstone bosons are eaten up by gauge fields to give massive W-bosons. At the boundary, however, the gauge field A_{\parallel} is frozen, so massless Goldstone bosons survive, parametrizing the desired G/\mathbb{T} .

We can see evidence of the boundary degrees of freedom if we look at the moduli space of supersymmetric vacua of the theory. In the UV, the boundary condition at infinity forces the boundary value of $\text{Re } \tau_{UV}\Phi$ to lie in a specific conjugacy class, but different points inside that G/\mathbb{T} manifold correspond to different vacua of the theory. Thus in the IR we need to have some boundary degrees of freedom which reproduce this space of vacua. The precise set of degrees of freedom depends on the choice of IR duality frame.

²⁴One could also deform the Dirichlet b.c. on the Higgs field to $\text{Re } \Phi|_{\partial} = \text{const}$, which introduces mass terms for the G/\mathbb{T} sigma models discussed below.

Consider the simplest example, $G = SU(2)$. The most obvious way to obtain a \mathbb{CP}^1 moduli space of vacua is to couple the IR theory to a 3d \mathbb{CP}^1 sigma-model, with $\mathcal{N} = 2$ supersymmetry. There is also a more subtle way. We need to remember the boundary conditions imposed by coupling the IR gauge fields to some matter: the vev of $\text{Re } a|_{\partial}$ acts as a 3d mass parameter, while $\text{Re } a_D|_{\partial}$ acts as an FI parameter

$$(4.7) \quad \text{Re } a_D|_{\partial} = \mu,$$

where μ is the moment map for the 3d flavor symmetry we are gauging. Thus if the boundary theory consisted, say, of a doublet of chiral fields q_{α} of charge -1 under the gauge group, the moment map condition together with the $U(1)$ gauge symmetry would reproduce precisely a \mathbb{CP}^1 moduli space of vacua.

In order to test this simple candidate theory we can couple it to both a non-abelian $SU(2)$ gauge theory and the abelian $U(1)$ gauge theory in order to engineer a potential RG domain wall. This chiral doublet interface between the non-abelian theory and the abelian IR theory gives rather reasonable boundary conditions. For example, we have

$$(4.8) \quad \text{Re } \tau_{UV} \Phi_{\alpha\beta}|_{\partial} = \bar{q}_{(\alpha} q_{\beta)} \quad \text{Re } a_D|_{\partial} = |q|^2,$$

which insures that q^{α} and \bar{q}^{α} are eigenvectors of $\text{Re } \tau_{UV} \Phi$ with eigenvalues $\text{Re } a_D$. As the eigenvalues of $\tau_{UV} \Phi$ are essentially the classical values of a_D , this condition agrees well with the condition that $\text{Re } a_D$ should remain constant from the UV to the IR side of a Janus.

On the other hand, in order to allow the q_{α} vevs without breaking SUSY, we need the 3d mass matrix

$$(4.9) \quad \text{Re } \Phi_{\alpha\beta}|_{\partial} - \text{Re } a \epsilon_{\alpha\beta}|_{\partial},$$

to have q_{α} as a zero eigenvector. Thus $\text{Re } a$ must coincide with an eigenvalue of $\text{Re } \Phi$ and we recover the condition that $\text{Re } a$ should remain constant from the UV to the IR side of a Janus.

If we move to other IR duality frames, which are more natural if the attractor RG fixed point is close to the monopole ($\zeta \sim i$) or dyon ($\zeta \sim 1$) points, we should do an appropriate $Sp(2, \mathbb{Z})$ transformation on the 3d degrees of freedom. This will produce a 3d GLSM: a theory of two chiral multiplets ϕ_1, ϕ_2 with an axial $U(1)$ symmetry that has been gauged. The remaining flavor symmetry is $SU(2)$ (rotating the chirals as a doublet) times a topological $U(1)_J$. The $U(1)_J$ will be coupled to the gauge fields in the

new duality frames. The abelian boundary conditions are covariant under $Sp(2, \mathbb{Z})$ transformations, and thus the new setups will work as well to mimic the desired boundary conditions.

In our later analysis based on the explicit RG manifold, we will recover these descriptions of the RG domain wall, and thus automatically test it further, by insuring the existence of appropriate couplings to the monopole and dyon particles and the correct behavior of line defects at the interface. This will provide rather robust evidence that the chiral doublet interface reflects faithfully the low energy properties of the RG domain wall.

4.3. The basic RG manifold M_0

The framed 3-manifold M_0 that represents the RG domain wall for pure $SU(2)$ theory should interpolate between an annulus \mathcal{C} “in the UV” and an annulus “in the IR” (Figure 26). In the IR, we represent \mathcal{C} as a big boundary with a WKB triangulation \mathbf{t} , coupling to abelianized Seiberg-Witten theory as discussed in Section 2.3. In the UV, however, we shrink the annulus into a small boundary, which *should* couple to the nonabelian $SU(2)$ theory. In the limit of infinite shrinking, the small annulus represents a defect, corresponding to Dirichlet b.c. for the UV gauge fields. The small annulus must be attached to two big cones at its ends, triangulated with degenerate triangles, which carry the irregular singularities of \mathcal{C} .




Figure 26: The RG manifold for pure $SU(2)$ theory, and its IR (abelian) and UV (nonabelian) boundaries.

Topologically, the 3-manifold M_0 is equivalent to $\mathcal{C} \times I$. Its boundary ∂M_0 is the union of the IR and UV boundaries, which have been “cut” along the two irregular singularities. Notice that ∂M_0 contains two small disc boundaries (on the irregular singularities) in addition to the small annulus.

To describe the space $\mathcal{P}_2(\partial M_0, \mathbf{t})$ of framed flat connections on ∂M_0 , we choose cross-ratio coordinates x_m, x_d for the magnetic and dyonic edges

on the big boundary. We also use a holonomy eigenvalue λ for the holonomy around the (oriented) A-cycle of the small annulus, and a canonically conjugate twist τ . Thus the nonvanishing Poisson brackets are

$$(4.10) \quad \{\log x_d, \log x_m\} = 2, \quad \{\log \tau, \log \lambda\} = 1.$$

Notice that by using λ rather than λ^2 as a coordinate, we have partially lifted from a moduli space of flat $PGL(2)$ connections to flat $SL(2)$ connections. It is also useful to introduce an electric coordinate $x_e = -\sqrt{x_m x_d}$.²⁵

The twist τ is described carefully in Appendix A.2. Its definition requires us to choose a path γ_τ from one end of the small annulus to the other. At each end of this path, the framing flags for a flat connection on the small annulus can be normalized by using the additional framing flags from the small discs of ∂M_0 . Then the ratio of normalizations is τ . Changing the path γ_τ by a full twist around the annulus (equivalent to twisting the entire 3-manifold, *cf.* Section 4.1) simply rescales $\tau \rightarrow \lambda^2 \tau$. Here, it will turn out to be more symmetric to choose two paths γ_1, γ_2 related by a full twist, and to set $\tau = \sqrt{\tau_1 \tau_2} = \lambda \tau_1 = \lambda^{-1} \tau_2$.




Figure 27: Triangulation of the RG manifold, showing how coordinates on the big and small annuli are related to tetrahedron parameters.

The 3-manifold M_0 can be triangulated as shown in Figure 27. First two tetrahedra are glued together along a common face to form a bipyramid.

²⁵Square roots like this can be made sense of in two ways: either by using logarithmically lifted coordinates as discussed in Appendix A, or breaking 3d symmetry slightly and imposing a positive structure as in [16].

Then two more faces (labelled \mathcal{A}) are identified to form M_0 . The small square of boundary at the front of the bipyramid gets rolled up into the small annulus of M_0 , while the big square at the back turns into the IR boundary of M_0 . Labeling the tetrahedra with edge parameters z, z', z'' and w, w', w'' as in the figure, and choosing τ paths as shown, we find that

$$(4.11) \quad \begin{aligned} x_m &= z'w', & x_d &= z''w'', & x_e &= \frac{1}{\sqrt{zw}}; \\ \lambda^2 &= \frac{z}{w}, & \tau &= \sqrt{\frac{z''}{w''} \frac{w'}{z'}} = \lambda \frac{z''}{w''}. \end{aligned}$$

The rules for obtaining λ and τ are summarized in Appendix A.3.3; basically one draws paths $\gamma_\lambda, \gamma_\tau$ on the small boundary and multiplies or divides (according to orientation) by the edge parameters on angles subtended by the paths.

In order to define a 3d theory $T_2[M_0, \mathbf{t}, \Pi]$, we must choose a polarization Π for $\mathcal{P}_2(\partial M, \mathbf{t})$ and then compare the position coordinates in Π to the positions/momenta of tetrahedra to figure out what symmetries get gauged. (Since there are no internal edges in the triangulation of M_0 , there will be no superpotential terms.) For the UV boundary, we canonically choose λ to be a position and τ its conjugate momentum. For the IR boundary, there are three natural choices of polarization:

$$(4.12) \quad \Pi_e = \begin{pmatrix} x_e \\ x_m \end{pmatrix}, \quad \Pi_m = \begin{pmatrix} x_m \\ x_e^{-1} \end{pmatrix}, \quad \Pi_d = \begin{pmatrix} x_d \\ x_e \end{pmatrix}.$$

so that the position coordinates (the top components of these vectors) are the electric, magnetic, and dyonic x 's, respectively.

In the polarization Π_e , the positions λ and x_e are just composed of tetrahedron positions z, w . So nothing gets gauged. We end up with a theory of two chirals ϕ_z, ϕ_w with charges $(-1, -1)$ under a $U(1)_e$ flavor symmetry and charges $(1, -1)$ under a $U(1)_\lambda$ flavor symmetry, which is manifestly enhanced to $SU(2)_\lambda$. This is the nonabelian symmetry associated to the small annulus. Being more careful to keep track of background Chern-Simons levels (which depend on the precise choice of conjugate momenta, here x_m, τ) and $U(1)_R$ charges (which result from a logarithmic lift of tetrahedron parameters, see Section 4.4), the full theory can be described as

(4.13)

 $T_2[M_0, \mathbf{t}, \Pi_e] :$

$$\left\{ \begin{array}{l} \text{Two chirals } \phi_z, \phi_w, \text{ with } U(2) \simeq SU(2)_\lambda \times U(1)_e \text{ flavor symmetry,} \\ \text{charges } \begin{array}{c|cc} & \phi_z & \phi_w \\ \hline e & -1 & -1 \\ \lambda & & \square \\ R & 0 & 0 \end{array}, \quad \text{abelian CS matrix } \begin{array}{c|ccc} & e & \lambda & R \\ \hline e & 1 & 0 & 1 \\ \lambda & 0 & 0 & 0 \\ R & 1 & 0 & * \end{array} \end{array} \right.$$

We do not specify an R-R Chern-Simons term, since it is not uniquely determined by the geometry.




Figure 28: Symmetries associated to boundaries and operators associated to edges, in $T_2[M_0, \mathbf{t}, *]$.

We can couple $T_2[M_0, \mathbf{t}, \Pi_e]$ to nonabelian $SU(2)$ gauge theory on one side and an *electric* $U(1)$ gauge theory on the other to obtain a full 3d-4d system with a half-BPS domain wall $\mathcal{B}_2[M_0]$. The abelian coupling is just as discussed in Section 2. In addition to gauging the $U(1)_e$ symmetry in the bulk, we note that the theory $T_2[M_0, \mathbf{t}, \Pi_e]$ has two chiral operators $\mathcal{O}_m, \mathcal{O}_d$ on which magnetic and dyonic 't Hooft lines from the bulk Seiberg-Witten theory can end. From a 3d point of view, these operators only exist in the presence of monopole flux configurations for a background $U(1)_e$ gauge field. The *nonabelian* coupling similarly gauges $SU(2)_\lambda$ in the bulk. But there exists a 3d $SU(2)_\lambda$ doublet $(q_1, q_2) \equiv (\phi_z, \phi_w)$ of chiral operators as well. The doublet is associated to either of the two degenerate edges at the end of the small annulus. Had there been any hypermultiplets in the bulk $SU(2)$ theory (there aren't here), they could have coupled to (q_1, q_2) at the wall.

Of course, we know from the analysis of Janus attractors that magnetic or dyonic duality frames, rather than an electric frame, are more natural for a weakly coupled $U(1)$ bulk theory. Correspondingly, we should look at 3d theories in polarizations Π_m or Π_d . The 3d axial $U(1)_e$ symmetry must now be gauged and replaced by topological flavor symmetries $U(1)_m$ or $U(1)_d$.

The difference between the two cases is that we gauge $U(1)_e$ at bare Chern-Simons level $+1$ or -1 . We find

(4.14a)

$T_2[M_0, \mathbf{t}, \Pi_m]$:

$$\left\{ \begin{array}{l} \text{Two chirals } \phi_z, \phi_w, \text{ with } U(1)_e \text{ gauge and } SU(2)_\lambda \times U(1)_m \text{ flavor,} \\ \text{charges } \begin{array}{c|ccc} & \phi_z & \phi_w & \mathcal{O}_m \\ \hline e & -1 & -1 & 0 \\ m & 0 & 0 & 1 \\ \lambda & & \square & - \\ \hline R & 0 & 0 & 0 \end{array}, \quad \text{abelian CS matrix } \begin{array}{c|ccc} & e & m & \lambda & R \\ \hline e & 1 & -1 & 0 & 1 \\ m & -1 & 0 & 0 & 0 \\ \lambda & 0 & 0 & 0 & 0 \\ \hline R & 1 & 0 & 0 & * \end{array} \end{array} \right.$$

(4.14b)

$T_2[M_0, \mathbf{t}, \Pi_d]$:

$$\left\{ \begin{array}{l} \text{Two chirals } \phi_z, \phi_w, \text{ with } U(1)_e \text{ gauge and } SU(2)_\lambda \times U(1)_d \text{ flavor,} \\ \text{charges } \begin{array}{c|ccc} & \phi_z & \phi_w & \mathcal{O}_d \\ \hline e & -1 & -1 & 0 \\ d & 0 & 0 & 1 \\ \lambda & & \square & - \\ \hline R & 0 & 0 & 0 \end{array}, \quad \text{abelian CS matrix } \begin{array}{c|ccc} & e & d & \lambda & R \\ \hline e & -1 & 1 & 0 & -1 \\ d & 1 & 0 & 0 & 0 \\ \lambda & 0 & 0 & 0 & 0 \\ \hline R & -1 & 0 & 0 & * \end{array} \end{array} \right.$$

These theories are basically the GLSM descriptions of the \mathbb{CP}^1 sigma-model that we anticipated in Section 4.2, aside from the bare Chern-Simons terms. By coupling (4.14a-b) to respective magnetic and dyonic Seiberg-Witten theories in the IR bulk (and the standard $SU(2)$ theory in the UV bulk) we obtain alternative descriptions of the full 3d-4d system $\mathcal{B}_2[M_0]$. Indeed, as discussed in Section 2.3.1, choices of big-boundary triangulation and polarization get dissolved after coupling to the bulk.

We note that $T_2[M_0, \mathbf{t}, \Pi_m]$ now has a standard (dynamical) monopole operator \mathcal{O}_m , which couples via a superpotential to the fundamental magnetic hypermultiplet of the bulk Seiberg-Witten theory. Similarly, $T_2[M_0, \mathbf{t}, \Pi_m]$ has a standard anti-monopole operator \mathcal{O}_d with the right charge to couple to a fundamental bulk dyon. The existence of these operators can be justified using the techniques of [47] — it requires a careful analysis of Chern-Simons terms and parity anomalies.

4.4. Line operators

As an application of the 3-manifold geometry M_0 , we can derive the Ward identities for line operators hitting the RG wall. We add an angular momentum fugacity to the 3d-4d system, as discussed briefly in Section 2.2.3, so

that generators of the line operator algebra q -commutation relations. The resulting Ward identities reproduce known relations between cross-ratio and Fenchel-Nielsen coordinates in quantum Teichmüller theory [12, 15].

To quantize, we must work in logarithmic coordinates, as reviewed in Appendix A.3.5. Logarithmic coordinates also allow an unambiguous definition of square roots. First, we express the basic relations between boundary and bulk coordinates (4.11) as

$$(4.15) \quad \hat{X}_m = \hat{Z}' + \hat{W}', \quad \hat{X}_d = \hat{Z}'' + \hat{W}'', \quad \hat{X}_e = -\frac{1}{2}(\hat{Z} + \hat{W});$$

$$\hat{\Lambda} = \frac{1}{2}(\hat{Z} - \hat{W}), \quad \hat{\mathcal{T}} = \hat{Z}'' - \hat{W}'' + \hat{\Lambda} = -\hat{Z}' + \hat{W}' - \hat{\Lambda}.$$

The LHS of these expressions should all be viewed as operators in an algebra generated by $\hat{Z}, \hat{Z}', \hat{Z}'', \hat{W}, \hat{W}', \hat{W}''$ so that

$$(4.16) \quad [\hat{Z}, \hat{Z}'] = [\hat{Z}', \hat{Z}'] = [\hat{Z}'', \hat{Z}] = \hbar,$$

and similarly for the W 's, with $\hat{Z} + \hat{Z}' + \hat{Z}'' = \hat{W} + \hat{W}' + \hat{W}'' = i\pi + \frac{\hbar}{2}$. We can unambiguously exponentiate to find

$$(4.17) \quad \hat{x}_m = \hat{z}'\hat{w}', \quad \hat{x}_d = \hat{z}''\hat{w}'', \quad \hat{x}_e = \frac{1}{\sqrt{\hat{z}\hat{w}}}, \quad \hat{\lambda} = \sqrt{\frac{\hat{z}}{\hat{w}}}, \quad \hat{\tau} = q^{\frac{1}{2}}\hat{\lambda}\frac{\hat{z}''}{\hat{w}''},$$

and the inverse relations

$$(4.18) \quad \hat{w} = \frac{1}{\hat{\lambda}\hat{x}_e}, \quad \hat{z} = \frac{\hat{\lambda}}{\hat{x}_e}, \quad \hat{z}' = q^{\frac{1}{8}}\frac{1}{\sqrt{\hat{\lambda}}}\frac{1}{\sqrt{\hat{\tau}}}\sqrt{\hat{x}_m}, \quad \hat{w}' = q^{\frac{1}{8}}\sqrt{\hat{\lambda}}\sqrt{\hat{\tau}}\sqrt{\hat{x}_m},$$

$$\hat{z}'' = q^{-\frac{1}{8}}\frac{1}{\sqrt{\hat{\lambda}}}\sqrt{\hat{\tau}}\sqrt{\hat{x}_d}, \quad \hat{w}'' = q^{-\frac{1}{8}}\sqrt{\hat{\lambda}}\frac{1}{\sqrt{\hat{\tau}}}\sqrt{\hat{x}_d},$$

where $q = e^{\hbar}$ and all the exponentiated operators q -commute: *e.g.* $\hat{x}_m\hat{x}_d = q^2\hat{x}_d\hat{x}_m$, $\hat{\lambda}\hat{\tau} = q\hat{\tau}\hat{\lambda}$, etc.

Next, we use the Ward identities for tetrahedra to relate the UV and IR sides. They are $\hat{z}'' + \hat{z}^{-1} - 1 \simeq 0$ and $\hat{w}'' + \hat{w}^{-1} - 1 \simeq 0$, where “ $\simeq 0$ ” means “annihilates a partition function.” By substituting (4.18) into these Ward identities and simplifying we find

$$(4.19a) \quad \text{Wilson:} \quad \hat{\lambda} + \hat{\lambda}^{-1} \simeq \hat{x}_e + \frac{1}{\hat{x}_e} - \frac{1}{\hat{x}_e} \hat{x}_d$$

$$\simeq -\frac{\sqrt{\hat{x}_m''} \sqrt{\hat{x}_d''}}{q^{\frac{3}{4}}} - \frac{q^{\frac{3}{4}}}{\sqrt{\hat{x}_d''} \sqrt{\hat{x}_m''}} + q^{-\frac{1}{4}} \sqrt{\hat{x}_d''} \frac{1}{\sqrt{\hat{x}_m''}}$$

$$(4.19b) \quad \text{'t Hooft:} \quad \frac{q^{-\frac{1}{8}}}{\hat{\lambda} - \hat{\lambda}^{-1}} \left(\frac{1}{\sqrt{\hat{\lambda}}} \sqrt{\hat{\tau}} - \sqrt{\hat{\lambda}} \frac{1}{\sqrt{\hat{\tau}}} \right) \simeq -q^{-\frac{1}{4}} \sqrt{\hat{x}_m},$$

$$(4.19c) \quad \text{dyonic:} \quad \frac{q^{\frac{1}{8}}}{\hat{\lambda} - \hat{\lambda}^{-1}} \left(\sqrt{\hat{\lambda}} \sqrt{\hat{\tau}} - \frac{1}{\sqrt{\hat{\lambda}}} \frac{1}{\sqrt{\hat{\tau}}} \right) \simeq -q^{\frac{1}{4}} \frac{1}{\sqrt{\hat{x}_d}}.$$

The first equation relates the fundamental (spin 1/2) $SU(2)$ Wilson line in the UV to electric and dyonic $U(1)$ line operators in the IR. The other two equations relate the spin-1/2 UV 't Hooft line and dyonic ('t Hooft-Wilson) line to IR line operators. These relations were discussed from a purely 4d point of view in [20, 21].²⁶

Note that in an honest 4d $SU(2)$ gauge theory, the basic 't Hooft (and dyonic) line operators are in the spin-1 representation rather than the spin-1/2 as above; in the spin-1 Ward identities, the square roots of $\hat{\lambda}, \hat{\tau}, \hat{x}_m, \hat{x}_d$ would disappear. Alternatively, in a $PSU(2) \simeq SU(2)/\mathbb{Z}_2$ theory we should keep spin-1/2 magnetic operators but only use spin-1 Wilson lines.

The Ward identities (4.19) do not require any choice of polarization: they are valid in *any* duality frame describing the domain wall $\mathcal{B}_2[M_0]$. On the other hand, if we wanted to write down partition functions of an isolated 3d theory on S_b^3 or $S^2 \times_q S^1$ we would need to choose a big-boundary duality frame. In polarization Π_e , the rules of [5, 34, 48] compute an S_b^3 partition function that beautifully reproduces the cross-ratio/Fenchel-Nielsen kernel of [12]:

$$(4.20) \quad \mathcal{Z}_{\hbar}(\Lambda, X_e) = e^{\frac{1}{2\hbar}(\Lambda^2 + 2X_e^2 + 2X_e(2\pi i + \hbar))} \Phi_{\hbar}\left(i\pi + \frac{\hbar}{2} + \Lambda + X_e\right)$$

$$\cdot \Phi_{\hbar}\left(i\pi + \frac{\hbar}{2} - \Lambda + X_e\right)$$

$$= e^{\frac{1}{2\hbar}(X_e^2 + 2X_e\Lambda + (X_e + \Lambda)(2\pi i + \hbar))} \frac{\Phi_{\hbar}\left(i\pi + \frac{\hbar}{2} + \Lambda + X_e\right)}{\Phi_{\hbar}\left(-i\pi - \frac{\hbar}{2} + \Lambda - X_e\right)},$$

²⁶For comparison, we should note that our edge coordinates on a big boundary are related to those used more commonly in quantum Teichmüller theory by (*e.g.*) Fock and Goncharov and in [18, 21] as $\hat{X}_{\text{FG}} = \hat{X}_{\text{here}} - i\pi - \hbar/2$, or in exponentiated form $\hat{x}_{\text{FG}} = -q^{-\frac{1}{2}} \hat{x}_{\text{here}}$. Our convention is a little more natural from a 3d perspective.

with the quantum dilogarithm Φ_{\hbar} defined in Appendix C. This wavefunction is annihilated by (4.19).

5. Non-abelian symmetry enhancement

We saw in Section 4 that the 3d theory trapped on the RG wall for 4d $SU(2)$ Seiberg-Witten theory should have a non-abelian $SU(2)_{\lambda}$ flavor symmetry. We associated this theory to a framed 3-manifold M_0 (Figure 27), constructing it using the class- \mathcal{R} rules of [5]. Despite the fact that class- \mathcal{R} constructions naively lead to theories with abelian flavor symmetry, we found a manifest enhancement $U(1)_{\lambda} \rightarrow SU(2)_{\lambda}$ for the symmetry associated to the small annulus of M_0 . We would now like to argue that such an enhancement occurs for the symmetries associated to small annuli in any 3d-4d boundary condition based on a framed 3-manifold.

The basic idea is simple. First, let M be a framed 3-manifold whose small boundary contains a small annulus, and assume that the annulus is attached to degenerate triangles on the big boundary (as in the UV part of Figure 26). Then choose a 3d triangulation \mathbf{t}_{3d} of M such that the neighborhood of the small annulus looks like the basic RG manifold M_0 . This lets us decompose $M = M' \cup M_0$ as in Figure 29, where M' contains a big annular region on its big boundary, and gluing in M_0 effectively shrinks this big annulus to a small one.

Now we may construct the 3d theory associated to M by gluing together the theories associated to M' and M_0 , in such a way that the $SU(2)_{\lambda}$ symmetry of M_0 is inherited. Physically, we are taking the theory associated to M' and colliding it with a basic RG wall (in the IR \rightarrow UV direction) to recover the theory associated to M . The RG wall provides the $SU(2)_{\lambda}$ flavor symmetry.

To be more specific, we have to choose some polarizations. Let Π_m be the magnetic polarization for ∂M_0 , so that one of the edges on its big annulus carries a position coordinate. Then $T_2[M_0, \Pi_m]$ is basically the GLSM description of a $\mathbb{C}P^1$ sigma-model, as in (4.14a), with $SU(2)_{\lambda} \times U(1)_m$ flavor symmetry. For M' , we can choose a similar polarization Π'_m for the annular part of its big boundary (mirroring the polarization Π_m of ∂M_0), and any other polarization Π away from this annulus. Then $T_2[M', \Pi \times \Pi'_m]$ contains a $U(1)'_m$ symmetry associated to its big annulus. We form the glued theory $T_2[M, \Pi]$ by taking a product of the component theories, gauging the anti-diagonal combination $U(1)_V$ of $U(1)_m \times U(1)'_m$, and adding two superpotential terms that break the diagonal of $U(1)_m \times U(1)'_m$ and the topological symmetry $U(1)_J$ associated to $U(1)_V$. A few more details will be given




Figure 29: Cutting out a basic RG manifold M_0 to isolate a small annulus, leaving behind a big annulus in a modified manifold M' .

in Section 5.1. The point, however, is that this gluing operation does not disturb the $SU(2)_\lambda$ flavor symmetry of $T_2[M_0, \Pi_m]$, which now becomes a flavor symmetry of $T_2[M, \Pi]$.

This argument shows that whenever a small annulus attaches to the centers of degenerate triangles on the big boundary of M , its flavor symmetry is promoted from $U(1)$ to $SU(2)$. If the big-boundary triangulation \mathbf{t} is such that the annulus attaches in any other way, the flavor symmetry will generally *not* be enhanced in the isolated 3d theory $T_2[M, \mathbf{t}, \Pi]$. This is easy to see by starting with a degenerate triangulation and using flips (Figure 16) to change it. In the course of flipping, an $SU(2)$ doublet q_α of chirals associated to a degenerate triangle gets split, with only one of q_1, q_2 coupling to the flip operator; thus $SU(2)$ is explicitly broken. This is explained further in Section 5.2.

Fortunately, we already know how to remedy this problem. If we couple $T_2[M, \mathbf{t}, \Pi]$ to an abelian 4d $\mathcal{N} = 2$ theory $SW_2[\partial M]$ associated to its big boundary, creating a boundary condition $\mathcal{B}_{SW}[M]$, the dependence on triangulation \mathbf{t} and polarization Π disappears. We can first implement the coupling with a choice of \mathbf{t} so that $SU(2)$ flavor symmetry is already manifest in the 3d theory alone, and preserved by the coupling, proving that the full 3d-4d system $\mathcal{B}_{SW}[M]$ has $SU(2)$ symmetry. Then, by arguments of

Section 2.3.1, the $SU(2)$ symmetry must actually be restored when coupling in a 4d duality frame specified by any other \mathfrak{t} and Π .

We arrive at the following picture. Given a framed 3-manifold M with a small annuli, we can always construct a boundary condition $\mathcal{B}_{\text{SW}}[M]$ whose flavor symmetry contains a subgroup $SU(2)^a$. If we want an isolated 3d theory instead, we must choose a big-boundary triangulation \mathfrak{t} and polarization Π . For every annulus that ends in degenerate triangles of \mathfrak{t} , the 3d theory $T_2[M, \mathfrak{t}, \Pi]$ will retain an $SU(2)$. This may not be possible for all annuli. At the “non-enhanced” annuli, the expected $SU(2)$ symmetry is broken by superpotential couplings involving (halves of) chiral doublets.

Note that once we have a system $\mathcal{B}_{\text{SW}}[M]$ with non-abelian $SU(2)^a$ symmetry, we may proceed to couple it to non-abelian (or “UV”) 4d theories as well, just as we described in the introduction. The basic RG wall was one example of this, and more will come later.

In the remainder of this section, we fill in a few of the details from above, and also comment on symmetry enhancement and breaking at small torus boundaries.

5.1. Gluing in the RG manifold M_0

We begin by spelling out some of the details of the process in Figure 29: reconstructing the theory associated to M by colliding a basic RG wall into the theory associated to a modified manifold M' .

Consider first the phase spaces associated to the big boundary of M_0 and the annular region on the big boundary of M' . They are parameterized by edge coordinates x_m, x_d and x'_m, x'_d , respectively, with

$$(5.1) \quad \{\log x_m, \log x_d\} = -\{\log x'_m, \log x'_d\} = 1.$$

Note that the bracket has opposite sign for M' due to the reversed orientation.

Two new internal edges are created in the gluing of M_0 to M' , with gluing functions $c_m = x_m x'_m$ and $c_d = x_d x'_d$. Geometrically, we want to use c_m, c_d as moment maps for a symplectic reduction, enforcing the gluing constraints $c_m = c_d = 1$. Correspondingly, the gluing of gauge theories must involve the addition of two operators to the superpotential. In order to identify these operators, it is convenient to choose polarizations for the big-boundary phase spaces so that either x_m, x'_m or x_d, x'_d are positions. (It is impossible to choose all four edge coordinates as positions, since they do not commute.) We therefore take a magnetic polarization Π_m for ∂M_0 and

a mirror “magnetic” polarization Π'_m for ∂M :

$$(5.2) \quad \Pi_m = \begin{pmatrix} x_m \\ x_e^{-1} \end{pmatrix}, \quad \Pi'_m = \begin{pmatrix} x'_m \\ x'_e \end{pmatrix},$$

with $x_e = -\sqrt{x_m x_d}$, $x'_e = -\sqrt{x'_m x'_d}$ as usual. We supplement Π'_m with some other polarization Π for the part of the boundary phase space away from the annular region. (And of course we supplement Π_m with the usual length-twist pair (λ, τ) .)

With these polarizations, the RG theory $T_2[M_0, \Pi_m]$ has $SU(2)_\lambda \times U(1)_m$ flavor symmetry and a standard (monopole) operator \mathcal{O}_m charged under $U(1)_m$; while $T_2[M', \Pi \times \Pi'_m]$ has $U(1)'_m \times (\dots)$ symmetry and must have a similar chiral operator \mathcal{O}'_m charged under $U(1)'_m$. Then the operator enforcing the gluing constraint $c_m = x_m x'_m = 1$ in the product theory $T_\times := T_2[M_0, \Pi_m] \times T_2[M', \Pi \times \Pi'_m]$ is easy to make: it is just the product $\mathcal{O}_m \mathcal{O}'_m$.

The other operator enforcing $c_d = x_d x'_d = x_e^2 x'_e{}^2 / (x_m x'_m) = 1$ is trickier. We first have to change polarization for T_\times so that this gluing function c_d is a position. This means gauging the anti-diagonal (*i.e.* vector) part $U(1)_V$ of the flavor symmetry $U(1)_m \times U(1)'_m$. Notice that $\mathcal{O}_m \mathcal{O}'_m$ is invariant under $U(1)_V$, so it survives the gauging. After the gauging, there is a new topological flavor symmetry $U(1)_J$, as well as the remaining diagonal (*i.e.* axial) part $U(1)_A$ of $U(1)_m \times U(1)'_m$. Moreover, there must be a monopole operator η_d charged under $U(1)_J$. To the gauged theory we add

$$(5.3) \quad W = \mathcal{O}_m \mathcal{O}'_m + \eta_d,$$

breaking both $U(1)_A$ and $U(1)_J$, and obtaining $T_2[M, \Pi]$. The unbroken flavor symmetry is $SU(2)_\lambda \times (\dots)$, where (\dots) came from the boundary of M' away from the annular region.

The existence of a gauge-invariant chiral operator η_d follows from immediately from the analysis of gluing in [5]. One may also deduce its existence by reconstructing the theory $T_2[M, \Pi]$ from components $T_2[M', \Pi \times \Pi'_d]$ and $T_2[M_0, \Pi_d]$, *i.e.* starting from “dyonic” polarizations. Then we would find that the operator enforcing $c_d = 1$ is easy to describe as $\mathcal{O}_d \mathcal{O}'_d$, while the operator η_m enforcing $c_m = 1$ is non-trivial:

$$(5.4) \quad W = \eta_m + \mathcal{O}_d \mathcal{O}'_d.$$

Of course, it cannot matter which initial polarizations we start with, as long as the final one is the same.

5.2. Matter doublets

Our next item of business is to justify a claim made both above and back in Section 2, that whenever the big boundary of a 3-manifold contains a degenerate triangle surrounding a hole with flavor symmetry $SU(2)_\lambda$, the associated 3d theory contains a doublet of chiral operators q_α under $SU(2)_\lambda$. This is in contrast to all other edges of the big boundary (away from degenerate triangles), which according to [5] get single chiral operators \mathcal{O}_E .

We proceed by flipping. Suppose that (M, \mathbf{t}) is a framed 3-manifold whose boundary triangulation has a degenerate triangle surrounding a hole v . The immediate neighborhood of this degenerate triangle must look like the RHS of Figure 30. Then (M, \mathbf{t}) can be obtained by flipping an edge on the boundary of another 3-manifold (M', \mathbf{t}') whose boundary triangulation looks like the LHS of Figure 30 (and is equivalent to \mathbf{t} away from the degenerate triangle that we are studying). Recall that, from a 3d perspective, “flipping” means gluing on a tetrahedron (Figure 16); thus $M = M' \cup \Delta$.




Figure 30: Flipping to create a degenerate triangle on the boundary of M . The logarithmic edge coordinates on the LHS parametrized as $X_\pm = \Lambda \pm X + i\pi$, and obey the gluing constraint $X_- + Z = 2\pi i$.

The two edges of $\partial M'$ that end at v commute. Looking at spaces of flat connections, we can assign to them logarithmic edge coordinates $X_\pm = \Lambda \pm X + i\pi$, obeying the constraint $(X_+ - i\pi) + (X_- - i\pi) = 2\Lambda$ as in (2.10) or (A.9). We use logarithmic coordinates in order to keep track of R-charges. We choose a polarization Π' for the boundary phase space such that X and Λ are position coordinates. Correspondingly, the theory $T_2[M', \mathbf{t}', \Pi']$ has a $U(1)_\lambda \times U(1)_x$ as part of its flavor symmetry group, and (since this is a non-degenerate boundary triangulation) it must also have two operators $\mathcal{O}_+, \mathcal{O}_-$

associated to our two external edges, with charges

$$(5.5) \quad T_2[M', \mathbf{t}', \Pi'] : \quad \begin{array}{c|cc} & \mathcal{O}_+ & \mathcal{O}_- \\ \hline U(1)_\lambda & 1 & 1 \\ U(1)_x & 1 & -1 \\ U(1)_R & 1 & 1 \end{array}$$

Note that our precise choice of polarization *defined* what the $U(1)_\lambda$ and $U(1)_R$ currents were in this theory. It is this choice of $U(1)_\lambda$ that will get enhanced.

To perform the flip, we add a tetrahedron theory T_Δ with its chiral ϕ and a superpotential coupling

$$(5.6) \quad W = \phi \mathcal{O}_- .$$

The superpotential kills the operator \mathcal{O}_- , replacing it with the operator ϕ , which has exactly opposite flavor charges and R-charge $R(\phi) = 2 - R(\mathcal{O}_-) = 1$. Thus

$$(5.7) \quad T_2[M, \mathbf{t}, \Pi] : \quad \begin{array}{c|cc} & \mathcal{O}_+ & \phi \\ \hline U(1)_\lambda & 1 & -1 \\ U(1)_x & 1 & 1 \\ U(1)_R & 1 & 1 \end{array}$$

Now the pair of operators $(q_1, q_2) := (\mathcal{O}_+, \phi)$ can become a doublet under an enhanced $U(1)_\lambda \rightarrow SU(2)_\lambda$. Notice that, by careful definition of flavor and R currents, they have exactly the same charges under all other symmetries of the theory. This is what we set out to show.

Geometrically, the flip adds a tetrahedron with edge parameter Z as shown. A new internal edge is created, enforcing the constraint $Z + X_- = 2\pi i$, or $Z = -\Lambda + X + i\pi$. The coordinate Z labels the circular edge of the degenerate triangle on the RHS, and clearly this (standard) edge carries the operator ϕ . The degenerate edge of the triangle on the RHS gets coordinate $X_+ + Z' + Z'' = 2\Lambda + i\pi$. Thus it exclusively carries the flavor symmetry of the hole, but (being degenerate) it cannot carry any operators on its own. In order to define the precise polarization Π for the theory $T_2[M, \mathbf{t}, \Pi]$ with (potentially) enhanced flavor symmetry, we parametrized the edges near the hole in terms of Λ and X , and took Λ and X as positions.

5.3. Three-punctured spheres and UV theories of class \mathcal{S}

The above analysis of chiral doublets and flips has a particularly nice application when big boundary are 3-punctured spheres.

A 3-punctured sphere \mathcal{C} only admits four distinct triangulations, shown in Figure 31, and they are precisely related by the kinds of flips we just described. Let us call the three punctures v_i , $i = 1, 2, 3$, with corresponding holonomy eigenvalues λ_i and flavor symmetries $U(1)_i$.

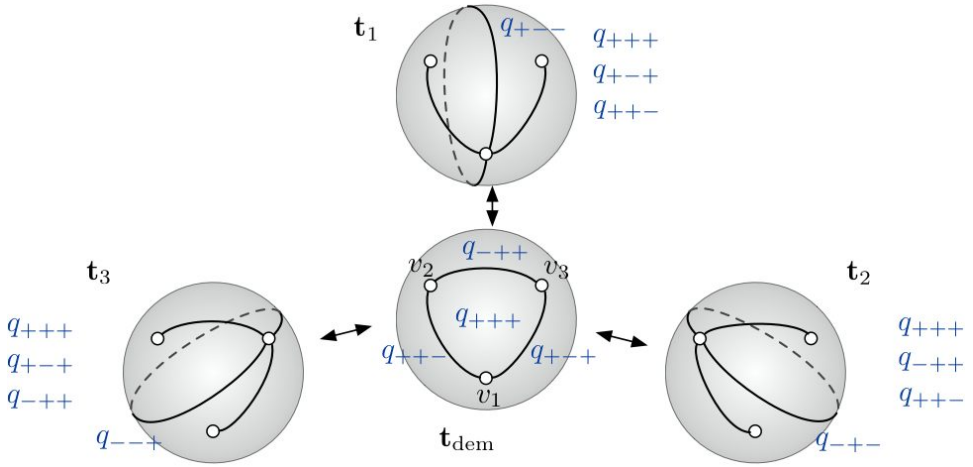


Figure 31: Flipping among the four triangulations of a 3-punctured sphere.

One of the triangulations — the democratic one, \mathbf{t}_{dem} — has two edges ending on each of the three punctures. The edge connecting v_1 and v_2 has edge parameter $-\lambda_1 \lambda_2 / \lambda_3$ (logarithmically, $\Lambda_1 + \Lambda_2 - \Lambda_3 - i\pi$), and similarly for the other edges, so that the product around any puncture v_i equals λ_i^2 . Note that all the edges commute; indeed, these edge coordinates are exclusively made from eigenvalues at holes, reflecting the fact that 3-punctured spheres carry no independent degrees of freedom for $SL(2)$ flat connections. If M is a 3-manifold with a democratically triangulated 3-punctured sphere on its boundary, we would always choose the λ_i as positions in a boundary polarization Π . Then we expect that $T_2[M, \mathbf{t}_{\text{dem}}, \Pi]$ has three chiral operators $q_{+ + -}$, $q_{+ - +}$, $q_{- + +}$ associated to three edges of \mathbf{t}_{dem} , with subscripts indicating charges under $U(1)_1 \times U(1)_2 \times U(1)_3$. Moreover, the product

$$(5.8) \quad q_{+ + +} = q_{+ + -} q_{+ - +} q_{- + +}$$

provides an important additional operator with all charges $+1$.

Now suppose that we flip the edge connecting v_2 and v_3 , to move to triangulation \mathbf{t}_1 , with two degenerate triangles. The new theory $T_2[M, \mathbf{t}_1, \Pi]$ could potentially have enhanced $U(1)_1 \times SU(2)_2 \times SU(2)_3$ flavor symmetry, and the matter content should reflect this. Indeed, the flip replaces the operator q_{-++} with a new operator q_{+--} of opposite flavor charges, leaving q_{++-} , q_{+-+} , q_{+++} undisturbed. But these four q 's can succinctly be written as $q_{+\beta\gamma}$, a bifundamental of $SU(2)_2 \times SU(2)_3$.

Similarly, flipping the other two edges of the democratic triangulation leads to theories with chiral operators $q_{\alpha\beta+}$ and $q_{\alpha+\gamma}$, filling out bifundamental multiplets under the other combinations of (potentially) enhanced flavor groups.




Figure 32: A genus-three boundary stretched into a trivalent network of tubes connected by 3-punctured spheres. The triangulations of the 3-punctured spheres dictate whether $T_2[M, \mathbf{t}]$ will have various symmetry enhancements; two possibilities are shown.

This analysis leads to a nice combinatorial picture of when a 3-manifold whose boundary contains a trivalent network of small tubes (*a.k.a.* defects) can lead to a 3d theory with enhanced flavor symmetry. At every trivalent juncture of tubes lies a 3-punctured sphere, for which we must choose one of the four triangulations of Figure 31. It is certainly most symmetric to choose triangulations \mathbf{t}_{dem} for every sphere, so that every tube ends at a bivalent hole. But then no $U(1)$ symmetries for the tubes will be enhanced! To get symmetry enhancement, we must instead choose some degenerate triangulations \mathbf{t}_i . Any tube that has *both* ends ending at a univalent hole, in the center of a degenerate triangle, will get an $SU(2)$ enhancement (Figure 32). This is never possible for all tubes at once, since none of the triangulations \mathbf{t}_i allow all three holes on a 3-punctured sphere to be univalent. Nevertheless, the “obstruction” to symmetry enhancement is minor: it comes from chiral matter that can be flipped away.

5.3.1. 4d coupling. We expect that when the boundary of M contains a trivalent network of small tubes that represents a pants decomposition of

some surface $\tilde{\mathcal{C}}$, the theory $T_2[M, \mathbf{t}]$ can be coupled to the UV theory $T_2[\tilde{\mathcal{C}}]$ of class \mathcal{S} , forming a 3d-4d boundary theory $\mathcal{B}_2[M]$. We can see this explicitly here.

First, for every three-punctured sphere \mathcal{C} in ∂M we can couple $T_2[M, \mathbf{t}]$ to a “trinion theory” $T_2[\mathcal{C}]$. The 4d trinion theory contains eight half-hypermultiplets $Q_{\alpha\beta\gamma}$ in the trifundamental representation of $SU(2)_1 \times SU(2)_2 \times SU(2)_3$ [1]. We split the $Q_{\alpha\beta\gamma}$ into two sets with opposite charges according to the chosen triangulation of \mathcal{C} . For example, for triangulation \mathbf{t}_1 , we split $Q_{\alpha\beta\gamma} \rightarrow (Q^{+\beta\gamma}, Q_{+\beta\gamma})$.²⁷ Then we introduce complementary Neumann/Dirichlet b.c. for these two subsets, and a standard boundary coupling, *e.g.*

$$(5.9) \quad \text{triangulation } \mathbf{t}_1 : \quad W_{\text{bdy}} = \sum_{\beta\gamma} Q^{+\beta\gamma} \Big|_{\partial} q_{+\beta\gamma}, \quad Q_{+\beta\gamma} \Big|_{\partial} = q_{+\beta\gamma}.$$

Let $\mathcal{B}_{\text{SW}}[M]$ denote the 3d theory coupled to free 4d bulk hypermultiplets like this. By the arguments of Section 2.3.1 the boundary condition $\mathcal{B}_{\text{SW}}[M]$ must have enhanced $SU(2)$ flavor symmetry for every single tube.

To finish, we just gauge all these $SU(2)$ symmetries of $\mathcal{B}_{\text{SW}}[M]$ in the bulk. This creates the full non-abelian theory $T_2[\tilde{\mathcal{C}}]$ of class \mathcal{S} in the bulk. It is manifestly presented in the duality frame corresponding to our chosen pants decomposition of $\tilde{\mathcal{C}}$. We obtain the desired boundary condition $\mathcal{B}_2[M]$.

5.3.2. Normalization of conformal blocks. The choices made for triangulations of 3-punctured spheres map very naturally to standard choices that are often made when normalizing conformal blocks in Liouville theory. We briefly explain how this works.

Consider the Liouville conformal block corresponding to a pants decomposition of a surface $\tilde{\mathcal{C}}$, *i.e.* to a Moore-Seiberg graph that is the skeleton of the pants decomposition. Specifying the conformal block requires taking a “square root” of the DOZZ 3-point function at every trivalent juncture. Two options for taking this square root were explored in [49, 50]. It is convenient to factor the 3-point function into pieces that look like quantum dilogarithm functions, as in [50], since such functions obey algebraic difference equations, but there is no canonical way to do this.

By the AGT correspondence [3], the conformal block is a partition function of $T_2[\tilde{\mathcal{C}}]$ on half of S_b^4 [51, 52], with certain boundary conditions on the equator. The boundary conditions can pick Neumann vs. Dirichlet for the hypermultiplets of $T_2[\tilde{\mathcal{C}}]$. Algebraically, this is precisely the choice of how

²⁷As usual, $SU(2)$ indices are raised and lowered with $\epsilon_{\alpha\beta}$.

to split the DOZZ 3-point function into quantum dilogarithms. But we also have just argued above that the choice of boundary conditions for hypermultiplets maps, geometrically, to the different triangulations of 3-punctured spheres. The two choices are one and the same.

Of course, once we glue conformal blocks together to form a complete Liouville partition function — or a partition function of $T_2[\widetilde{\mathcal{C}}]$ on a full S_b^4 — all ambiguities cancel out. In our case, it is more natural to couple the 4d theory on a half-sphere to an appropriate 3-manifold theory on the equator, and again all ambiguities must cancel out.




Figure 33: The quantum dilogarithm prefactors that multiply an S_b^3 partition function when various flips are performed. These can explicitly break Weyl symmetry $\Lambda_i \rightarrow -\Lambda_i$ in the partition functions, leading to expected breaking of $SU(2)_i$ flavor symmetries. Some additional quadratic-exponential prefactors accompany these Φ_{\hbar} functions, and correspond to background Chern-Simons terms; they can be fixed by carefully keeping track of dual twists \mathcal{T}_i .

Along these lines, we can explicitly see the effect that changing triangulations has on conformal blocks by considering S_b^3 partition functions \mathcal{Z}_b of 3d theories. Each flip multiplies \mathcal{Z}_b by a quantum dilogarithm as shown in Figure 33. These flips must simultaneously *divide* conformal blocks by the same quantum dilogarithms.

5.4. Small tori

Our last special topic in symmetry enhancement concerns the small torus boundaries of a 3-manifold M . The situation is more subtle than with small annuli.

Small torus boundaries, just like small annuli, carry an abelian $U(1)$ flavor symmetry. From our discussion above, we might expect that it is promoted to a non-abelian $SU(2)$. After all, we expect that in six dimensions the small torus boundaries are associated with regular defects that wrap closed loops in M .

However, we immediately see a problem with this proposal. For M with a small torus boundary, the three-dimensional theory $T_2[M, \Pi]$ depends on a choice of A and B cycles on the torus — implicit in the polarization Π . The $U(1) \equiv U(1)_A$ symmetry is naturally associated with the A cycle. Changing the choice of cycles corresponds to performing a usual *abelian* $Sp(2, \mathbb{Z})$ transformation on $T_2[M, \Pi]$. If the $U(1)$ is indeed promoted to $SU(2)$, the promotion should happen in all possible $Sp(2, \mathbb{Z})$ frames. But this cannot (naively) be the case.

Suppose, for example, that we act with T^k on a theory where the $U(1)_A$ is promoted to $SU(2)_A$, changing the B cycle $\gamma_B \rightarrow \gamma_B + k\gamma_A$. This adds abelian background Chern-Simons terms for $U(1)_A$, seemingly preventing enhancement in the new frame. Similarly, if we act with ST^k for large k , corresponding to new cycles $(\gamma'_A, \gamma'_B) = (-k\gamma_A - \gamma_B, \gamma_A)$, we will gauge $U(1)_A$ at Chern-Simons level k , obtaining a new theory with a topological $U(1)'_A$ symmetry that rotates the dual photon. One might hope that quantum effects would enhance $U(1)'_A$ to $SU(2)'_A$. However, for large k , the gauge group $U(1)_A$ is arbitrarily weakly coupled, so such quantum effects cannot be expected to help.

A possible way to understand the lack of enhancement directly in a 3d theory $T_2[M, \Pi]$ comes from the observation that this theory has an exactly marginal superpotential coupling for every torus cusp. This can be seen geometrically. When M is triangulated, every torus cusp leads to one redundant gluing constraint (Theorem A.2, page 260). This implies that when constructing $T_2[M, \Pi]$ one of the operators in the superpotential — where each operator is associated to a gluing constraint — does not break any flavor symmetries and automatically has R-charge equal to two. By the results of [53], this operator must be exactly marginal. We can imagine a scenario where $U(1)_A$ is enhanced to $SU(2)_A$ only at a special point in the exactly marginal parameter space. But $Sp(2, \mathbb{Z})$ transformations need not preserve this special locus. Thus, generically, $U(1)_A$ is not enhanced.

This scenario can be matched nicely with a 6d analysis. A regular defect in six dimensions carries a triplet of moment-map operators, each in the adjoint of the $SU(2)_A$ flavor symmetry group. After we twist the theory on M , with the defect wrapping a closed loop in M , two of these moment-map operators remain scalars and can be combined into a complex chiral operator

$\mu_{\mathbb{C}}$ for the effective 3d theory $T_2[M]$. It automatically has R-charge two in the standard 3d normalization: the 3d R-charge is a $U(1)_R$ subgroup of the $SU(2)_R$ R-symmetry preserved by a defect in flat space, and the triplet of moment maps has spin one under $SU(2)_R$, while the supercharges have spin $1/2$.²⁸ We can now choose a Cartan subgroup of $SU(2)_A$, and add to the 3d superpotential the $\mu_{\mathbb{C}}^3$ component of the complex moment-map operator in the direction of this Cartan. We get an exactly marginal deformation along which we break $SU(2)_A$ flavor symmetry down to $U(1)_A$. This is compatible with the constraints of [53]: the two remaining components $\mu_{\mathbb{C}}^{\pm}$ of the complex chiral moment map combine with the broken $SU(2)_A$ flavor currents to give non-protected multiplets.

We now have a consistent picture of how non-abelian symmetry enhancement/breaking could happen. It is also interesting to observe that protected quantities that are insensitive to marginal deformations — such as S_b^3 partition functions or sphere indices — always have manifest Weyl invariance $\lambda_A \rightarrow \lambda_A^{-1}$ in the parameter corresponding to a $U(1)_A$ symmetry. This supports the idea that there exists a point of enhancement in the marginal parameter space (much as in the examples of [17]), though it is certainly not a proof.

The analysis of moment maps and marginal deformations here can be extended to a theory $T_K[M]$ of K M5 branes in a straightforward fashion. This was done in [6].

6. Example 1: $\mathcal{N} = 2^*$ interfaces

We would now like to use the technology of the previous sections to construct RG walls and UV duality walls for some 4d theories of class \mathcal{S} with matter. Our approach will be to first build framed 3-manifolds representing the desired interfaces, and then to read off effective 3d interface theories (and their 4d couplings) from triangulations of the manifolds.

We begin here by looking at $\mathcal{N} = 2^*$ theory, *i.e.* four-dimensional $\mathcal{N} = 2$ theory with gauge group $SU(2)$ and an adjoint hypermultiplet $\Phi_{\alpha\dot{\alpha}}$. Let's review some of its geometric features. The complex mass of the hypermultiplet is associated with an $SU(2)_{\mu}$ flavor symmetry. To see this flavor symmetry, we must split the hypermultiplet into chirals $\Phi_{\alpha\dot{\alpha}} = (X_{\alpha\dot{\alpha}}, Y_{\alpha\dot{\alpha}})$; then

²⁸These properties of the moment-map operators can be derived by embedding the entire setup in M-theory. While dynamical branes wrap $M \times \mathbb{R}^3$ in $T^*M \times \mathbb{R}^3 \times \mathbb{R}^2$, the defect branes wrap the conormal bundle $N^*\mathcal{K} \times \mathbb{R}^3$ of a curve $\mathcal{K} \subset M$. The complex scalar moment-map operator parametrizes motion of the defect branes in \mathbb{R}^2 , and must have R-charge two since $U(1)_R$ rotates this \mathbb{R}^2 .

$SU(2)_\mu$ acts on (X, Y) as a doublet. If the mass is tuned to zero, the theory gains maximal $\mathcal{N} = 4$ supersymmetry.

As a theory of class \mathcal{S} , $\mathcal{N} = 2^*$ theory is well known to be realized by compactification on a punctured 2-torus, $\mathcal{C} = T_*^2$ [1, 2]. The $SU(2)_\mu$ flavor symmetry is associated to the puncture, while the $SU(2)_\lambda$ gauge symmetry (in given weakly coupled description) is associated to an A-cycle for the torus. If we stretch T_*^2 into a trivalent network of tubes as on the RHS of Figure 34 below, the network contains a single three-punctured sphere. Correspondingly, $T_2[T_*^2]$ should contain eight half-hypermultiplets $Q_{\alpha\dot{\alpha}\gamma}$. They combine into the traceless adjoint hyper $(X_{\alpha\dot{\alpha}}, Y_{\alpha\dot{\alpha}}) = ((Q_{\alpha\dot{\alpha}+}, Q_{\alpha\dot{\alpha}-}) - \text{trace})$ and an additional decoupled gauge-invariant hyper $\hat{\Phi} = (\hat{X}, \hat{Y}) = (Q_{\alpha}^{\alpha+}, Q_{\alpha}^{\alpha-})$. Thus $T_2[T_*^2]$ is a small extension of “standard” $\mathcal{N} = 2^*$ theory.

The usual modular group $PSL(2, \mathbb{Z})$ acts on the A and B cycles of the punctured torus, and provides a UV duality group for $T_2[T_*^2]$. In particular, the S element exchanges the A and B cycles, acting as electric-magnetic duality. The 3d degrees of freedom living on an S-duality wall — where dual $SU(2)$ gauge couplings become weak on the two sides — were extracted in a theory called $T[SU(2)]$ in [9]. We will reconstruct this theory from a framed 3-manifold M_S that implements the S cobordism. The $T[SU(2)]$ theory of [9] had a manifest $SU(2)_\lambda \times U(1)_{\lambda'}$ flavor symmetry that was enhanced to $SU(2)_\lambda \times SU(2)_{\lambda'}$ by quantum effects, which could then couple to two copies of 4d $\mathcal{N} = 2^*$ theory in the bulk. By choosing a convenient triangulation of M_S in Section 6.3, we will discover a mirror-symmetric description of $T[SU(2)]$ that has manifest $SU(2)_\lambda \times SU(2)_{\lambda'}$ already in the Lagrangian.

First, however, we will describe the more elementary $\mathcal{N} = 2^*$ RG wall.

6.1. The RG wall

The RG manifold M_{RG} for $T_2[T_*^2]$ should interpolate between a trivalent network of tubes (representing the UV) and a WKB triangulation of T_*^2 (representing the IR). These boundaries are shown in Figure 34. Modulo the action of the modular group, there is a unique triangulation of the one-punctured torus. Its three edges can be labelled by electric, magnetic, and dyonic charges of the $U(1)$ Seiberg-Witten theory in the infrared, satisfying

$$(6.1) \quad \langle \gamma_m, \gamma_e \rangle = \langle \gamma_e, \gamma_d \rangle = \langle \gamma_d, \gamma_m \rangle = 2, \quad \gamma_m + \gamma_d + \gamma_e = -\gamma_\mu,$$

where γ_μ is the flavor charge of states under a Cartan subgroup $U(1)_f \subset SU(2)_\mu$ (normalized such that a doublet has charges $\pm\gamma_\mu$). Correspondingly,

we associate parameters x_m, x_d, x_e to the edges of the triangulation, satisfying

$$(6.2) \quad \{\log x_m, \log x_e\} = \{\log x_e, \log x_d\} = \{\log x_d, \log x_m\} = 2, \\ x_m x_d x_e = -\mu^{-1}.$$

These parametrize (a partial lift) framed flat $PSL(2)$ connections on the triangulated T_*^2 , with fixed holonomy eigenvalue μ^{-1} at the puncture.²⁹

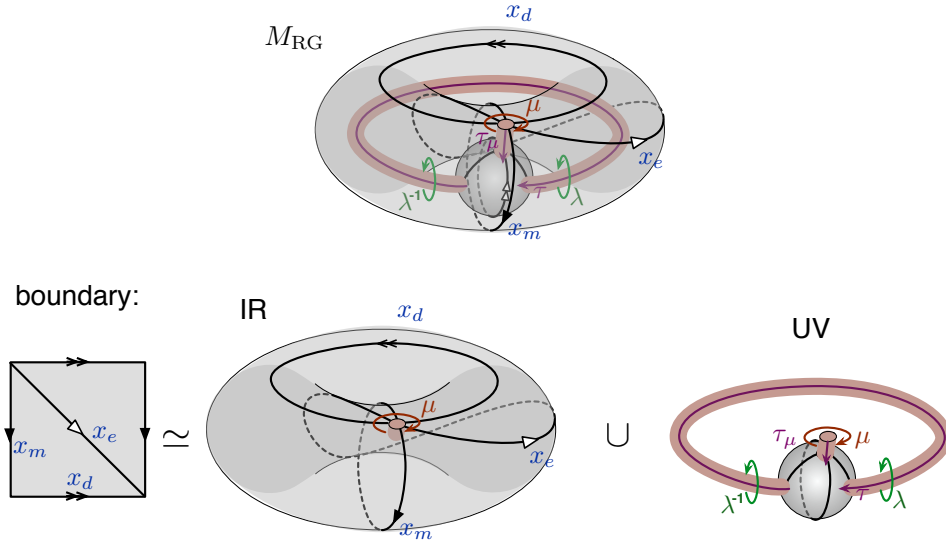


Figure 34: The RG manifold for $\mathcal{N} = 2^*$ theory (left), interpolating between IR and UV boundaries.

On the UV side, we find a big 3-holed sphere, two of whose holes are connected to each other by a small annulus. We define a pair of eigenvalue-twist coordinates (λ, τ) for this annulus, as in Section 4.3 or Appendix A.2. We will want the 3d theory $T_2[M_{RG}]$ by itself to have manifest $SU(2)_\lambda$ symmetry. Therefore (according to Section 5.3) we choose the triangulation for the 3-punctured sphere with two degenerate triangles surrounding the holes with eigenvalues $\lambda^{\pm 1}$. The third hole of the 3-punctured sphere must be connected by another small annulus to the hole on the IR boundary.

²⁹Logarithmically, the central constraint in (6.2) is $2(X_m + X_d + X_e - 3\pi i) = -2M$. As usual, this tells us how to “lift” from specifying the squared-eigenvalue at the puncture to specifying the eigenvalue itself. Dividing by two and exponentiating, we find the minus sign in (6.2).

This annulus has holonomy eigenvalue μ (from the perspective of the 3-punctured sphere, the eigenvalue is inverted), and we may also give it a dual twist coordinate τ_μ .

Altogether, the RG manifold of Figure 34 has a 6-dimensional boundary phase space, with \mathbb{C}^* coordinates $\{x_m, x_d, x_e; \lambda, \tau; \mu, \tau_\mu\}$ satisfying the Poisson brackets (6.2) as well as

$$(6.3) \quad \{\log \tau, \log \lambda\} = \{\log \tau_\mu, \log \mu\} = 1.$$

The UV coordinates (λ, τ) commute with the flavor coordinates (μ, τ_μ) and with all the external edge coordinates. The flavor twist-parameter τ_μ , however, cannot commute with all three IR edge coordinates simultaneously. We will choose³⁰ it to commute with x_e and x_m . Then one natural polarization Π_e for $\mathcal{P}_2(\partial M_{\text{RG}}, \mathfrak{t})$ is

$$(6.4) \quad \Pi_e : \quad \text{positions } (x_e, \lambda, \mu), \quad \text{conjugate momenta } (\sqrt{x_m}, \tau, \tau_\mu).$$

The simplest way to triangulate M_{RG} is to take the basic RG manifold M_0 for pure $SU(2)$ theory and identify the two circular edges at its ends. This turns the big annulus on the IR boundary into a big one-punctured torus. However, the resulting manifold is a little too “thin” at the newly identified edge — its UV and IR boundaries are pinched together there, so that they touch. We can fix this by flipping the thin edge on the outside, as in Figure 35. The edge created by the flip gets the electric coordinate x_e .

By combining the triangulation of the basic RG manifold M_0 (two tetrahedra) with the extra flip (one more tetrahedron), we arrive at the triangulation shown in Figure 36. We label the edges of the third tetrahedron with parameters f, f', f'' . The piece of small boundary that forms the small annulus of M_0 is unchanged, and we read off relations

$$(6.5a) \quad \lambda^2 = \frac{z}{w}, \quad \tau = \lambda \frac{z''}{w''}$$

³⁰We emphasize that there is no canonical choice for τ_μ , because the IR boundary does not have a triangulation with a degenerate triangle surrounding its hole (on which the flavor annulus ends). However, a canonical choice is *not necessary*! The basic reason for this is that in coupling $T_2[M_{\text{RG}}, \mathfrak{t}, \Pi]$ to 4d UV and IR theories, we will never gauge the $U(1)_\mu$ symmetry associated to the flavor annulus — we will simply identify it with flavor symmetries in the bulk. The ambiguity in choosing τ_μ translates to ambiguous background Chern-Simons couplings involving $U(1)_\mu$ gauge fields, which we do not really care about if this symmetry is not gauged.

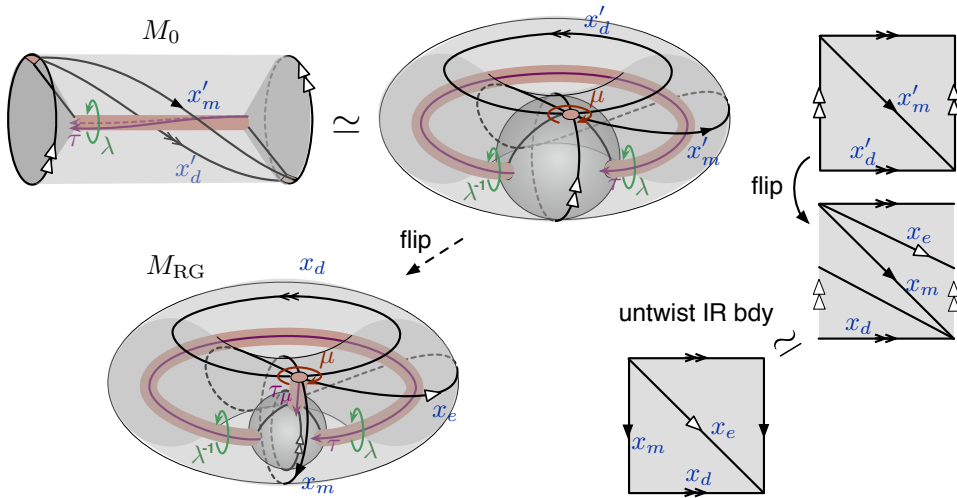


Figure 35: Identifying the ends of a basic RG manifold M_0 (for pure $SU(2)$ theory) to obtain the RG manifold for $\mathcal{N} = 2^*$. An extra flip thickens out the manifold, separating the UV and IR boundaries.

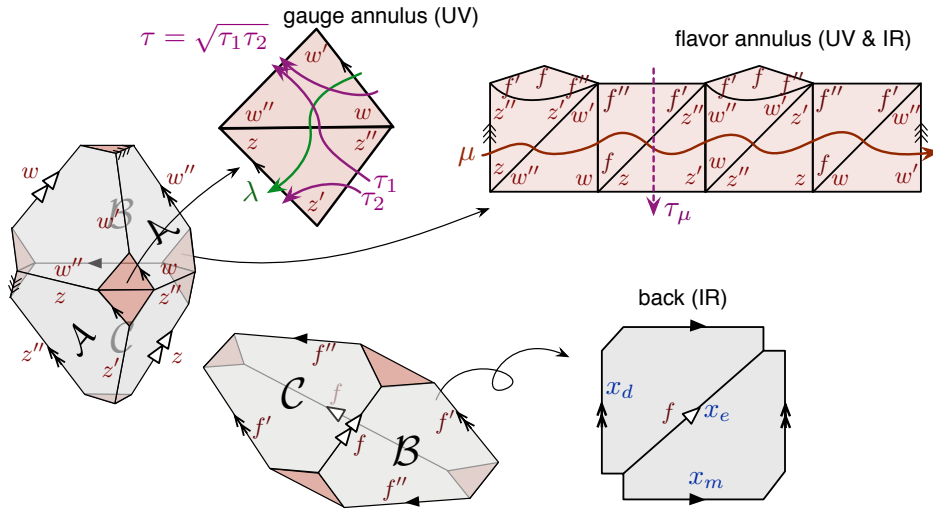


Figure 36: Triangulation of the RG manifold for $\mathcal{N} = 2^*$ theory.

just as for pure $SU(2)$ theory. The three edges on the IR boundary are changed a little bit by the flip; they get relations

$$(6.5b) \quad x_e = f, \quad x_m = x'_m f''^2 = z' w' f''^2, \quad x_d = x'_d f'^2 = z'' w'' f'^2,$$

where x'_m, x'_d denote the edge coordinates of M_0 from (4.11).

The remaining ten small triangles of the three tetrahedra come together to tile the small flavor annulus. Then we may read off the eigenvalue μ and also define a dual twist τ_μ ³¹

$$(6.5c) \quad \mu^2 = \frac{zwf^2}{z'z''w'w''} \Rightarrow \mu = -zwf; \quad \tau_\mu := \frac{f''}{\sqrt{x_m}} = \frac{1}{\sqrt{z'w'}}.$$

Notice that the expected relation $\mu^{-1} = -x_e x_m x_d$ holds, and could actually have been used to bypass the computation of μ from a path on the small annulus. Here the twist τ_μ is simply chosen as an appropriate combination of tetrahedron parameters that satisfies the desired commutation relations from above (in particular, τ_μ commutes with x_e and x_m); it could also be expressed as a certain average of paths on the flavor annulus.

6.1.1. Gauge theory and couplings. We can now build the theory $T_2[M_{\text{RG}}, \mathbf{t}, \Pi_e]$ in a fairly intuitive manner, mirroring the construction of M_{RG} . We simply take the basic RG theory $T_2[M_0, \mathbf{t}, \Pi'_e]$ from (4.13) of Section 4.3 — a theory of two chirals $q_\alpha = (\phi_z, \phi_w)$ with $SU(2)_\lambda \times U(1)'_e$ flavor symmetry — and add one more chiral ϕ_f from the third tetrahedron. There is no superpotential, since no internal edges are created in the gluing. Moreover, since the positions x_e, λ, μ in our polarization Π_e only involve tetrahedron positions z, w, f , there is no gauging. Thus, $T_2[M_{\text{RG}}, \mathbf{t}, \Pi_e]$ is basically a theory of three free chirals, whose flavor symmetry we write as $SU(2)_\lambda \times U(1)_e \times U(1)_\mu \subset U(3)$.

The flavor charges of q_α and ϕ_f can be read off by inverting the relations $\lambda^2 = z/w, \mu = -zwf$, and $x_e = f$. We find

$$(6.6) \quad \begin{array}{c|cc} & q_\alpha & \phi_f \\ \hline SU(2)_\lambda & \square & - \\ U(1)_\mu & \frac{1}{2} & 0 \\ U(1)_e & -\frac{1}{2} & 1 \end{array}$$

In addition, there is a $U(1)_R$ R-symmetry, whose charge assignment is dictated by the logarithmic form of the coordinate relations. A more careful analysis reveals that $R(q_\alpha) = 1/2$ and $R(\phi_f) = 0$.

Of course this theory of free chirals is accompanied by various background Chern-Simons terms. As usual, these can be determined from our choice of momenta in Π_e . All we really need to know, however, is that there

³¹Logarithmically, $M = Z + W + F - i\pi$ and $\mathcal{T}_\mu := -\frac{1}{2}(Z' + W')$.

is no abelian Chern-Simons term for $U(1)_\lambda$ that would prevent its enhancement to $SU(2)_\lambda$; such a term does not exist because it is absent in the basic RG theory $T_2[M_0, \mathfrak{t}, \Pi'_e]$, and we have made no changes to the polarization that involve λ or τ .

The theory $T_2[M_{\text{RG}}, \mathfrak{t}, \Pi_e]$ can couple both to non-abelian $\mathcal{N} = 2^*$ theory $T_2[T_*^2]$ and to the abelianized Seiberg-Witten theory $SW_2[T_*^2]$, by gauging the $SU(2)_\lambda$ and $U(1)_e$ flavor groups, respectively. On the abelian side, our 3d theory can couple to an electric BPS hypermultiplet $\Phi_e = (X_e, Y_e)$ via the superpotential

$$(6.7) \quad W_{\text{bdy, IR}} = X_e|_{\partial} \mathcal{O}_e, \quad \mathcal{O}_e \equiv \phi_f,$$

where \mathcal{O}_e is the standard operator associated to the electric edge on the big boundary. More interestingly, on the non-abelian side the 3d theory can couple to the adjoint hypermultiplet $\Phi_{\alpha\dot{\alpha}} = (X_{\alpha\dot{\alpha}}, Y_{\alpha\dot{\alpha}})$ via

$$(6.8) \quad W_{\text{bdy, UV}} = X^{\alpha\dot{\alpha}}|_{\partial} q_\alpha q_{\dot{\alpha}} \phi_f.$$

In this case, the three operators $q_\alpha q_{\dot{\alpha}} \phi_f$ are associated to the three-punctured sphere of ∂M_{RG} , a special case of the general structure of Section 5.3. Notice in particular that the operators $q_\alpha q_{\dot{\alpha}} \phi_f$ are neutral under the IR symmetry $U(1)_e$, and have the right $SU(2)_\lambda \times U(1)_\mu$ charges to couple to $X^{\alpha\dot{\alpha}}$.

Quite beautifully, the bulk-boundary couplings give us a full set of relations between elementary operators in the UV and the IR. Namely, after integrating out ϕ_f , the full bulk-boundary superpotential $W_{\text{bdy, IR}} + W_{\text{bdy, UV}}$ imposes the relation $X^{\alpha\dot{\alpha}}|_{\partial} q_\alpha q_{\dot{\alpha}} = X_e|_{\partial}$ between half-hypermultiplets on the two sides. In addition, the above couplings must be accompanied by modified Dirichlet b.c. for $Y_{\alpha\dot{\alpha}}$ and Y_e , which set $Y_{\alpha\dot{\alpha}}|_{\partial} = q_\alpha q_{\dot{\alpha}} \phi_f$ and $Y_e|_{\partial} = \phi_f$, in other words $Y_{\alpha\dot{\alpha}}|_{\partial} = q_\alpha q_{\dot{\alpha}} Y_e|_{\partial}$. These relations supersymmetrize nicely the boundary conditions for the vectormultiplet scalar fields.

Finally, we should note that even though $U(1)_\mu$ is not enhanced to $SU(2)_\mu$ in the isolated 3d theory $T_2[M_{\text{RG}}, \mathfrak{t}, \Pi_e]$ (we would not expect an enhancement, since the flavor annulus does not end on degenerate triangles), it must be enhanced in the full coupled 3d-4d system. The mechanism for enhancement was discussed in Section 5.

6.1.2. Line operators. We can use the 3d geometry of M_{RG} to relate coordinates for flat $SL(2)$ connections on the IR boundary to those on the UV boundary. This translates as usual to formulas relating UV and IR line operators. These formulas for $\mathcal{N} = 2^*$ theory are already well known from gauge/Liouville theory [21, 49, 50] and from Teichmüller theory (*e.g.*)

[16, 54, 55]; and were interpreted three-dimensionally in [20]. It is nice to see now that the formulas arise naturally from our triangulations as well.³²

As in Section 4.4,³³ the first step is to invert the relations between tetrahedron parameters and boundary coordinates. Both quantization and roots can be handled unambiguously by first using logarithmic relations and then exponentiating. From (6.5) we find

$$(6.9) \quad \hat{z} = iq^{\frac{1}{4}} \hat{\lambda} \sqrt{\frac{\hat{\mu}}{\hat{x}_e}}, \quad \hat{w} = iq^{\frac{1}{4}} \frac{1}{\hat{\lambda}} \sqrt{\frac{\hat{\mu}}{\hat{x}_e}}, \quad \hat{f} = \hat{x}_e,$$

$$\hat{z}'' = iq^{-\frac{1}{8}} \sqrt{\frac{\hat{x}_e}{\hat{\mu} \hat{\lambda}}} \sqrt{\hat{\tau}} \hat{\tau}_\mu, \quad \hat{w}'' = iq^{-\frac{1}{8}} \sqrt{\frac{\hat{\lambda}}{\hat{x}_e \hat{\mu}}} \frac{\hat{\tau}_\mu}{\sqrt{\hat{\tau}}}, \quad \hat{f}'' = \sqrt{\hat{x}_m} \hat{\tau}_\mu.$$

Then we can substitute these into the (quantum) tetrahedron Lagrangians $\hat{z}'' + \hat{z}^{-1} \simeq 0$, etc., and simplify, to obtain the desired relations between different sets of boundary coordinates.

The structure of the resulting equations turns out to be a little simpler if we do another flip of the edge labelled ‘ x_e ’ on the IR boundary of the RG manifold. Basically this is because the flip we used in Figure 35 to thicken the manifold twisted the triangulation of the IR boundary with respect to the UV λ -annulus; we want to flip again to undo this twist. The new flip, a standard one, adds a new tetrahedron (with parameters g, g', g'') and creates an internal edge (gluing constraint $c := fg = 1$, or $C = F + G - 2\pi i = 0$). Call the new IR boundary edge-coordinates x_e'', x_m'', x_d'' , with $\{\log x_e'', \log x_m''\} = \{\log x_m'', \log x_d''\} = \{\log x_d'', \log x_e''\} = 2$ and (upon quantization) $\hat{x}_d'' \hat{x}_m'' \hat{x}_e'' = -q^{\frac{1}{2}} \hat{\mu}^{-1}$. The bulk-boundary relations become

$$(6.10) \quad \hat{z} = -iq^{-\frac{1}{4}} \lambda \sqrt{\frac{\hat{x}_e'' \hat{\mu}}{\hat{c}}}, \quad \hat{w} = -\frac{i}{q^{\frac{1}{4}} \hat{\lambda}} \sqrt{\frac{\hat{x}_e'' \hat{\mu}}{\hat{c}}}, \quad \hat{f} = q \frac{\hat{c}}{\hat{x}_e''},$$

$$\hat{g} = \hat{x}_e'', \quad \hat{z}'' = -iq^{\frac{3}{8}} \sqrt{\frac{\hat{c}}{\hat{x}_e'' \hat{\lambda} \hat{\mu}}} \sqrt{\hat{\tau}} \hat{\tau}_\mu, \quad \hat{w}'' = iq^{\frac{3}{8}} \sqrt{\frac{\hat{c} \hat{\lambda}}{\hat{x}_e'' \hat{\mu}}} \frac{\hat{\tau}_\mu}{\sqrt{\hat{\tau}}},$$

$$\hat{f}'' = \hat{\gamma} \hat{\tau}_\mu, \quad \hat{g}'' = -\frac{q \hat{\gamma}}{\hat{x}_e''} \frac{1}{\sqrt{\hat{x}_m''}},$$

³²Interestingly, an ideal triangulation of the geometry M_{RG} (or rather the thin M'_{RG}) was also used in [56] to find the classical (un-quantized) version of relations (6.11). There, however, it was done without reference to moduli spaces of framed flat connections, or to the symplectic structures needed for quantization.

³³The full quantization procedure is summarized in Appendix A.3.5.

where $\hat{\gamma} = \hat{f}'' \sqrt{\hat{z}' \hat{w}'}$ is an operator canonically conjugate to the gluing operator \hat{c} .

Upon substituting (6.10) into the four tetrahedron equations $\hat{z}'' + \hat{z}^{-1} - 1 \simeq 0$, etc.; eliminating $\hat{\gamma}$ and $\hat{\tau}_\mu$; and setting $\hat{c} = 1$; we obtain the desired UV–IR relations. (We eliminate the conjugate $\hat{\tau}_\mu$ to the flavor symmetry parameter μ because we want to keep this as a non-dynamical flavor symmetry — we will not consider flavor ’t Hooft loops that could shift μ .) After some rearrangement, we find

$$(6.11a) \quad (\text{Wilson}) \quad \hat{\lambda} + \hat{\lambda}^{-1} \\ \simeq \frac{iq^{\frac{1}{4}}}{\sqrt{\mu \hat{x}_e''}} + \frac{\sqrt{\mu \hat{x}_e''}}{iq^{\frac{1}{4}}} - \frac{iq^{\frac{1}{4}}}{\sqrt{\mu \hat{x}_e''}} \frac{1}{\hat{x}_m''} \\ \simeq -\frac{\sqrt{\hat{x}_m''} \sqrt{\hat{x}_d''}}{q^{\frac{3}{4}}} - \frac{q^{\frac{3}{4}}}{\sqrt{\hat{x}_d''} \sqrt{\hat{x}_m''}} + q^{-\frac{1}{4}} \sqrt{\hat{x}_d''} \frac{1}{\sqrt{\hat{x}_m''}},$$

$$(6.11b) \quad \frac{q^{-\frac{1}{8}}}{\hat{\lambda} - \hat{\lambda}^{-1}} \left(\frac{1}{\sqrt{\hat{\lambda}}} \sqrt{\hat{\tau}} - \sqrt{\hat{\lambda}} \frac{1}{\sqrt{\hat{\tau}}} \right) \simeq -q^{-\frac{1}{4}} \sqrt{\hat{x}_m''},$$

$$(6.11c) \quad \frac{q^{\frac{1}{8}}}{\hat{\lambda} - \hat{\lambda}^{-1}} \left(\sqrt{\hat{\lambda}} \sqrt{\hat{\tau}} - \frac{1}{\sqrt{\hat{\lambda}}} \frac{1}{\sqrt{\hat{\tau}}} \right) \simeq -q^{\frac{1}{4}} \frac{1}{\sqrt{\hat{x}_d''}}.$$

In the IR variables \hat{x}_m and \hat{x}_d , these are identical to the Ward identities (4.19) for pure $SU(2)$ theory — the mass μ drops out. Thus is no surprise, since the RG manifold for $\mathcal{N} = 2^*$ theory is built from the basic RG manifold M_0 . The first equation in (6.11) relates the UV Wilson line to IR line operators. The remaining equations, however, do not immediately correspond to UV ’t Hooft (or dyonic) line operators; some further manipulation is required to put them in that form.

6.2. The S-wall: $T[SU(2)]$ with manifest $SU(2) \times SU(2)$ flavor symmetry

The 3-manifold M_S that implements S-duality for the non-abelian $\mathcal{N} = 2^*$ theory was described back in Figure 8 of the Introduction. It interpolates between two trivalent networks representing two different pants decompositions (related by $S \in Sp(2, \mathbb{Z})$) of the one-punctured torus T_*^2 . Geometrically, M_S is the complement of the so-called “Hopf network” in S^3 .

The boundary of the framed manifold M_S consists of two big 3-punctured spheres and three small annuli. Two of the small annuli connect a 3-punctured sphere back to itself. We will choose triangulations for the 3-punctured sphere such that both of these small annuli end on degenerate triangles, and thus carry enhanced flavor symmetries $SU(2)_\lambda$, $SU(2)_{\lambda'}$. These are the symmetries that get gauged when coupling $T_2[M_S, \mathbf{t}]$ to 4d $\mathcal{N} = 2^*$ theory in its two S-dual frames. The third small annulus forms a “bridge” between the three-punctured spheres; it will carry a $U(1)_\mu$ flavor symmetry, identified with the Cartan of the flavor group of $\mathcal{N} = 2^*$ theory in the bulk.

The phase space $\mathcal{P}_2(\partial M_S, \mathbf{t})$ is six-dimensional and can be parametrized by three pairs of eigenvalue-twist coordinates (λ, τ) , (λ', τ') , and (μ, τ_μ) , one for each small annulus. The twists τ and τ' are canonically defined (modulo multiplication by λ and λ' , respectively), and because the bridge connects two three-punctured spheres it turns out that the twist τ_μ also has a canonical definition (modulo multiplication by $\mu^{\frac{1}{2}}$). The coordinates obey

$$(6.12) \quad \{\log \tau, \log \lambda\} = \{\log \tau', \log \lambda'\} = \{\log \tau_\mu, \log \mu\} = 1,$$

with all other commutators vanishing, and there is a canonical polarization Π_S with λ, λ', μ as positions and τ, τ', τ_μ as conjugate momenta. So we want to find $T_2[M_S, \mathbf{t}, \Pi_S]$ and its bulk-boundary couplings.

The manifold M_S can be glued together from two basic RG manifolds M_0 and M'_0 . Geometrically, M_0 and M'_0 first have their ends identified, forming two solid tori (as in the first part of Figure 35), each with a three-punctured sphere and a small annulus running through its core. These are two copies of the RG manifold for $\mathcal{N} = 2^*$ theory, without the extra flip that was added in Section 6.1; the thickening from the extra flip is not needed. The boundaries of the two solid tori are then identified with an S transformation, forming a 3-sphere S^3 with the Hopf network in its interior (Figure 37).

By triangulating M_0 and M'_0 the standard way and then gluing them together, we obtain the triangulation of M_S shown in Figure 38, involving four tetrahedra. We can actually guess immediately what one Lagrangian description of $T_2[M_S, \mathbf{t}, \Pi_S]$ should be. One internal edge is formed in the gluing, enforcing the constraint that the product of magnetic IR coordinates for M_0 and M'_0 is one:

$$(6.13) \quad c := x_m x'_m = 1.$$

(Here x_m, x_d, x_e and x'_m, x'_d, x'_e will denote IR coordinates for the two basic RG manifolds, defined just as in Section 4.3.) Recall that in its magnetic polarization, the theory $T_2[M_0, \mathbf{t}, \Pi_m]$ is $U(1)_e$ gauge theory at level +1,

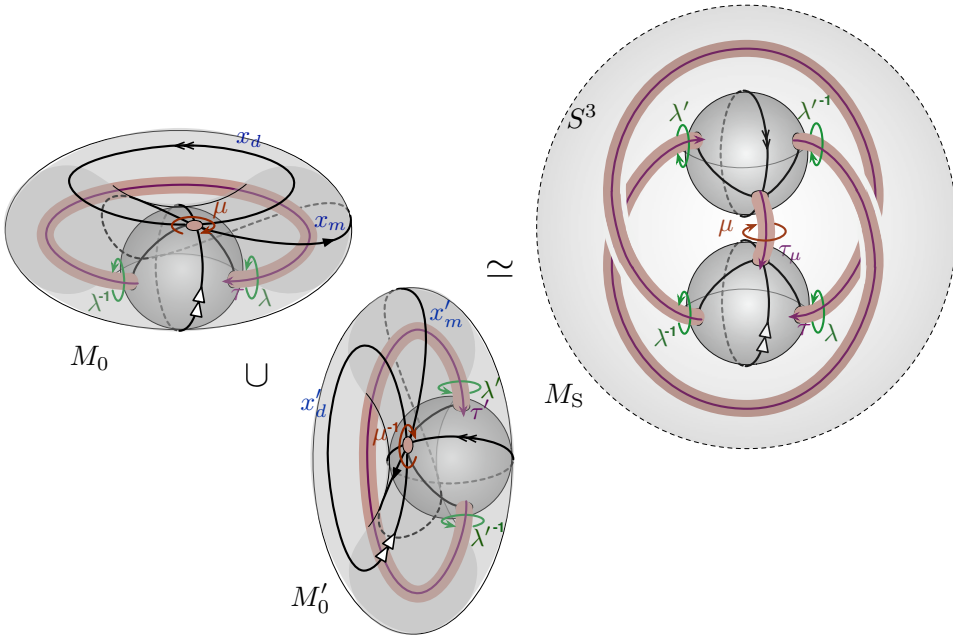


Figure 37: Gluing together two solid tori, each representing an RG manifold for $\mathcal{N} = 2^*$ theory (with no extra thickening), to form the S-duality manifold.

coupled to a doublet of chirals q_α (both with $U(1)_e$ charge +1); and this theory has a monopole operator \mathcal{O}_m transforming under the topological $U(1)_m$ flavor symmetry. Therefore, to form $T_2[M_S, \mathbf{t}, \Pi_S]$ we simply take a product $T_2[M_0, \mathbf{t}, \Pi_m] \times T_2[M'_0, \mathbf{t}', \Pi'_m]$ and add a superpotential

$$(6.14) \quad W = \mathcal{O}_m \mathcal{O}'_m$$

to break the diagonal of the flavor symmetry $U(1)_m \times U(1)'_m$ and enforce the gluing. Of course $SU(2)_\lambda \times SU(2)'_\lambda$ flavor symmetry, acting on doublets q_α and q'_α , is inherited. It turns out that to make the third flavor symmetry $U(1)_\mu$ manifest, we should also gauge the anti-diagonal of $U(1)_m \times U(1)'_m$, under which $\mathcal{O}_m \mathcal{O}'_m$ is invariant; then $U(1)_\mu$ is its associated topological symmetry.

To justify the claim about gauging the anti-diagonal of $U(1)_m \times U(1)'_m$, we should be a bit more careful about the relation between coordinates in $\mathcal{P}_2(\partial M_S, \mathbf{t})$ and the coordinates of M_0, M'_0 , as well as individual tetrahedra. For the two independent “gauge” annuli of M_0 and M'_0 we have the usual




Figure 38: A triangulation of M_S obtained by triangulating M_0 and M'_0 .

relations, cf. (4.11),

$$(6.15) \quad \lambda^2 = \frac{z}{w}, \quad \tau = \lambda \frac{z''}{w''}; \quad \lambda'^2 = \frac{r}{s}, \quad \tau' = \lambda' \frac{r''}{s''}.$$

The remaining twelve small triangles of our four tetrahedra come together to tile the “bridge” annulus of M_S , as shown in the figure. We can find μ from a path around the A-cycle of this annulus,

$$(6.16) \quad \mu = -z w s'' r'' = -\frac{x'_d}{x_e{}^2} \simeq -\frac{x'_e{}^2}{x_d},$$

where the two expressions on the RHS are equivalent modulo the gluing constraint $c = x_m x'_m = z' w' r' s' = 1$. We could also easily obtain μ by noting that $\mu^{\pm 1}$ must be the coordinates of the circular edges on the three-punctured spheres, labelled by $\blacktriangleright\blacktriangleright$ and $\triangleright\triangleright$ in the figures. Finally, a canonical τ_μ coordinate can be defined by first drawing a straight line³⁴ connecting the circular edge on one three-punctured sphere to the circular edge on the other (dashed in Figure 38), and then averaging the path-coordinates for *sixteen* paths evenly distributed to the left and right of this line. Using the geometric dictionary of Appendix A.2, this translates to an averaging of normalizations of flags along the bridge annulus. A short calculation produces

$$(6.17) \quad \tau_\mu = \sqrt{\frac{w'' z''}{r'' s''}} = \sqrt{\frac{x_d}{x'_d}}.$$

It is the relation (6.16), involving both position ($x_m = -x_e^2/x_d$) and momentum (x_e) coordinates of the M_0 manifolds, that tells us to gauge the anti-diagonal of $U(1)_m \times U(1)'_m$ in the product of M_0 theories and to identify $U(1)_\mu$ with the new topological symmetry.

6.2.1. Bulk-boundary couplings. Now that we have the full map between tetrahedron coordinates and the M_S boundary phase space, we can actually give a slightly simpler description of $T_2[M_S, \mathbf{t}, \Pi_S]$. It corresponds roughly to gluing together two *electric* theories $T_2[M_0, \mathbf{t}, \Pi_e]$, which had no dynamical gauge symmetry to begin with. Ultimately, it does not matter whether we glue together the magnetic theories or the electric theories, or individual tetrahedra: the *only* difference in these procedures is how the $ISp(8, \mathbb{Z})$ transformation to the final polarization Π_S is implemented — how it is broken up into elementary generators. The very fact that the group $ISp(8, \mathbb{Z})$ acts on 3d theories guarantees that after a flow to the IR the results will be the same.

The slightly simpler theory involves two doublets of chirals q_α and q'_α where both the $U(1)_e$ rotating q_α and the $U(1)'_e$ rotating q'_α are gauged,

³⁴There is some freedom in choosing this line, but different choices simply rescales $\tau_\mu \rightarrow \mu^{\pm 1} \tau_\mu$, provided that we use the averaged τ_μ as described below!

with a *mixed* Chern-Simons coupling. In detail:

(6.18)

$T_2[M_S, \mathbf{t}, \Pi_S]$:

$$\left\{ \begin{array}{l} U(1)_e \times U(1)'_e \text{ gauge theory coupled to 4 chirals } q_\alpha, q'_\alpha, \\ \text{with } SU(2)_\lambda \times SU(2)'_\lambda \times U(1)_\mu \text{ flavor symmetry and } U(1)_R, \\ \text{and a superpotential } W = \mathcal{O}_c \end{array} \right.$$

There are several useful things to notice about this theory. It has a gauge-invariant monopole operator \mathcal{O}_c associated to the $U(1)_A$ gauge group that is a diagonal of $U(1)_e \times U(1)'_e$, and this operator breaks a topological flavor symmetry that enforces the gluing of M_S . The other topological symmetry, associated to the anti-diagonal of $U(1)_e \times U(1)'_e$ is our flavor symmetry $U(1)_\mu$.

The theory also has two monopole operators η and η' for $U(1)_e$ and $U(1)'_e$, respectively, which are not fully gauge invariant: η is charged under $U(1)'_e$, and vice versa. This is actually ideal, because it allows us to create good bulk-boundary couplings! Let $\Phi_{\alpha\dot{\alpha}} = (X_{\alpha\dot{\alpha}}, Y_{\alpha\dot{\alpha}})$ and $\Phi'_{\alpha\dot{\alpha}} = (X'_{\alpha\dot{\alpha}}, Y'_{\alpha\dot{\alpha}})$ denote the adjoint hypermultiplets of two S-dual copies of $\mathcal{N} = 2^*$ theory. Then we may introduce $SU(2)_\lambda \times SU(2)'_\lambda \times U(1)_\mu$ -invariant couplings

$$(6.19) \quad \begin{aligned} W_{\text{bdy}} &= X^{\alpha\dot{\alpha}}|_{\partial} q_\alpha q_{\dot{\alpha}} \eta + X^{\alpha\dot{\alpha}'}|_{\partial} q'_\alpha q'_{\dot{\alpha}} \eta'; \\ Y_{\alpha\dot{\alpha}}|_{\partial} &= q_\alpha q_{\dot{\alpha}} \eta, \quad Y'_{\alpha\dot{\alpha}}|_{\partial} = q'_\alpha q'_{\dot{\alpha}} \eta'. \end{aligned}$$

The “dressed” monopole operators $q_\alpha q_{\dot{\alpha}} \eta$ and $q'_\alpha q'_{\dot{\alpha}} \eta'$ are fully invariant under the 3d gauge group $U(1)_e \times U(1)'_e$, and are associated geometrically to the two three-punctured spheres in ∂M_S , as per Section 5.3. With boundary conditions (6.19), the 3d $U(1)_\mu$ flavor symmetry, identified with the Cartans of the bulk flavor symmetry in two $\mathcal{N} = 2^*$ theories, is promoted to $SU(2)_\mu$. We may complete our construction of an S-duality domain wall by gauging $SU(2)_\lambda \times SU(2)'_\lambda$ in the bulk.

6.3. Reproducing standard $T[SU(2)]$

The theory $T_2[M_S, \mathfrak{t}, \Pi_S]$ of Section 6.2 is a 3d mirror of (*i.e.* is infrared dual to) the standard 3d S-wall theory known as $T[SU(2)]$. It has a beautiful feature: its $SU(2)_\lambda \times SU(2)'_\lambda$ flavor symmetry is classically manifest. In contrast, in the standard description of $T[SU(2)]$, only $SU(2)_\lambda \times U(1)'_\lambda$ is manifest; the $U(1)'_\lambda$, being a topological symmetry, gets promoted to $SU(2)'_\lambda$ by quantum corrections. Of course, the classical $SU(2)_\lambda \times SU(2)'_\lambda$ symmetry comes at a price: the superpotential of our theory in (6.18) must contain a monopole operator.

The equivalence between $T_2[M_S, \mathfrak{t}, \Pi_S]$ above and standard $T[SU(2)]$ can be seen by applying a sequence of elementary dualities. These dualities correspond geometrically to tetrahedron rotations (replacing a $U(1)$ gauge theory with a charged chiral with a free chiral, or vice versa) and 2–3 moves (replacing an XYZ model with $N_f = 1$ SQED). We could discuss the sequence of dualities at the level of Lagrangians, but it is much nicer to understand the duality geometrically. To this end, we will exhibit a different 3d triangulation of M_S that produces standard $T[SU(2)]$ theory (almost) directly. Since two labelled triangulations of a framed 3-manifold must be related by a sequence of 2–3 moves and tetrahedron rotations,³⁵ this amounts to a proof that $T_2[M_S, \mathfrak{t}, \Pi_S] \simeq T[SU(2)]$.

Consider the triangulation in Figure 39. First an octahedron is constructed by gluing together four tetrahedra around a central edge. This is the same octahedron as the one glued together from M_0 and M'_0 in Figure 38, but with a different triangulation. Then two front faces (marked \mathcal{A}) and two back faces (marked \mathcal{D}) are glued to form the manifold M_S .

We expect to find a theory of four chirals $\phi_z, \phi_w, \phi_r, \phi_s$, with a superpotential

$$(6.20) \quad W = \phi_z \phi_w \phi_r \phi_s$$

enforcing the gluing. Moreover, some of the flavor symmetries that leave (6.20) invariant should get gauged in the process of changing polarization to match the canonical Π_S . To determine what gets gauged, we express

³⁵This is not quite a mathematical theorem since, since we require our triangulations to satisfy a few non-degeneracy constraints. In the presence of these constraints, a full result about relating 3d triangulations is not yet known. Nevertheless, in the present case it is very easy to find the sequence of non-degenerate 2–3 moves (and tetrahedron rotations) relating the different triangulations of M_S ; it boils down to 2–3 moves on the octahedron that were discussed in [5, Sec. 4.5].

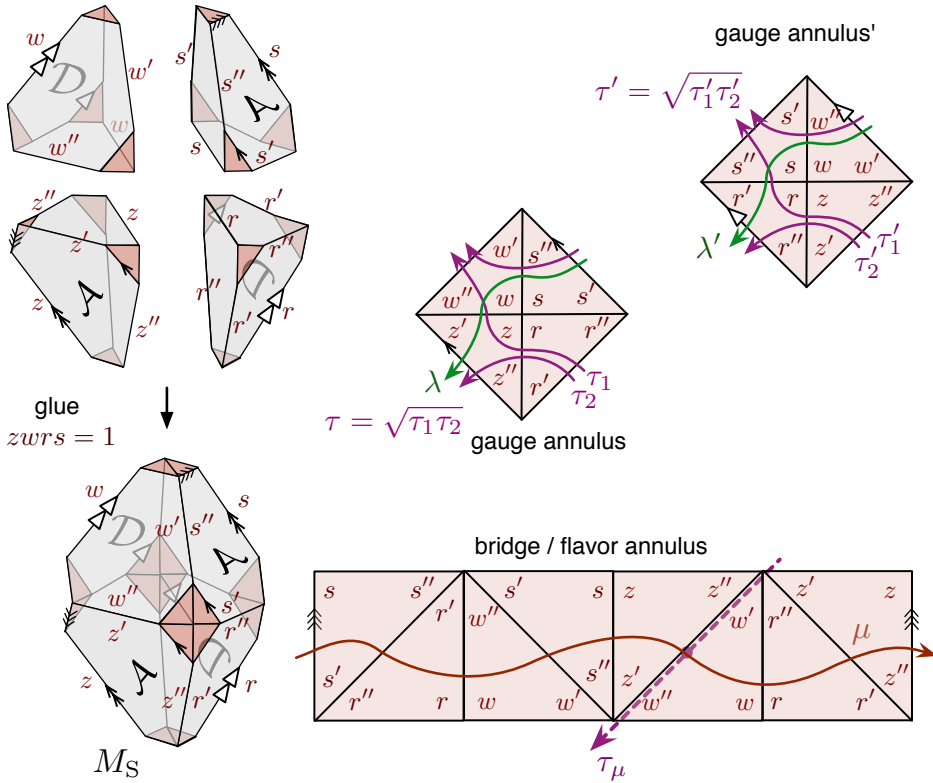


Figure 39: Gluing four tetrahedra a different way to form M_S .

the positions λ, λ', μ in terms of tetrahedron parameters using the standard rules, and find

$$(6.21) \quad \lambda^2 = \frac{s'' z'}{w}, \quad \lambda'^2 = \frac{w'' r'}{s}, \quad \mu^2 = \frac{rw}{sz}.$$

Similarly, we may compute the conjugate twist coordinates

$$(6.22) \quad \tau = \frac{\lambda}{\lambda'^2} \frac{z}{s}, \quad \tau' = \frac{\lambda'}{\lambda^2} \frac{r}{w}, \quad \tau_\mu^2 = \frac{r'' w''}{s'' z''}.$$

The relation for μ^2 to tetrahedron positions indicates that $U(1)_\mu$ is the symmetry that rotates the chirals with charges $(-\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, -\frac{1}{2})$. The other two

$U(1)$ flavor symmetries get gauged, so that $U(1)_\lambda$, and $U(1)'_\lambda$ become topological symmetries. Unraveling everything, we find a $U(1)_e \times U(1)'_e$ dynamical gauge theory with

$$(6.23) \quad \begin{array}{c|cccc} & \phi_z & \phi_s & \phi_r & \phi_w \\ \hline e & 1 & -1 & 0 & 0 \\ e' & 0 & 0 & 1 & -1 \\ \mu & -\frac{1}{2} & -\frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ R & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \end{array}$$

$$\text{CS level matrix: } \begin{array}{c|cc|cc} & e & e' & \mu & \lambda & \lambda' \\ \hline e & 0 & -1 & 0 & 2 & 0 \\ e' & -1 & 0 & 0 & 0 & 2 \\ \mu & 0 & 0 & -\frac{1}{2} & 0 & 0 \\ \lambda & 2 & 0 & 0 & 1 & -2 \\ \lambda' & 0 & 2 & 0 & -2 & 1 \end{array}$$

This is basically two copies of $N_f = 1$ SQED, with a mixed Chern-Simons term coupling the $U(1)_e$ and $U(1)'_e$ photons. We also see from the Chern-Simons (*i.e.* FI) couplings that the topological symmetries $U(1)_\lambda$ and $U(1)'_\lambda$ are directly associated to $U(1)_e$ and $U(1)'_e$. Evidently, the nonabelian enhancement to $SU(2)_\lambda \times SU(2)'_\lambda$ must be a quantum phenomenon.

To get from this theory — which must be mirror to $T_2[M_S, \mathfrak{t}, \Pi_S]$ — to standard $T[SU(2)]$ we perform one *non-geometric* 2–3 move. We isolate the first copy of SQED, with gauge group $U(1)_e$, and dualize it to an XYZ model. This duality is non-geometric because it does not correspond to an actual 2–3 move performed on the tetrahedra of M_S ; nevertheless, it makes perfect sense in gauge theory.

In SQED \rightarrow XYZ duality [47], the gauge-invariant meson $\phi_z\phi_s$ is replaced by a single fundamental chiral M . Moreover, the monopole and anti-monopole operators $\eta, \tilde{\eta}$ of $U(1)_e$ become fundamental chirals. They have charges ± 2 under $U(1)_\lambda$ and ∓ 1 under the remaining gauge group $U(1)'_e$ due to the mixed Chern-Simons terms. We find, then, a new superpotential

$$(6.24) \quad W \rightarrow \phi_r M \phi_w + \eta M \tilde{\eta},$$

where the second term is the standard XYZ contribution. After shifting the $U(1)_\lambda$ flavor current by a multiple of the gauge current $U(1)'_e$, we kill the background Chern-Simons couplings for $U(1)_\lambda$ and obtain a gauge theory

with charges

$$(6.25) \quad \begin{array}{c|ccccc} & \phi_r & \phi_w & \tilde{\eta} & \eta & M \\ \hline e' & 1 & -1 & 1 & -1 & 0 \\ \hline \lambda & 1 & -1 & -1 & 1 & 0 \\ \hline \mu & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & -1 \end{array}.$$

This is the $\mathcal{N} = 2$ deformation of the $T[SU(2)]$ theory of [9]. Now there is a classical enhancement $U(1)_\lambda \rightarrow SU(2)_\lambda$, while $U(1)'_\lambda$ remains a topological symmetry, with quantum enhancement.

6.4. S-duality of line operators

We can use the triangulation of the S-wall geometry M_S to reproduce the map between Wilson and 't Hooft line operators in $\mathcal{N} = 2^*$ theory and its S-dual. From the current perspective, this is actually the easiest way to identify how the 't Hooft operator should look in terms of (λ, τ) . Mathematically, we are finding the action of the mapping class group of the one-punctured torus on Fenchel-Nielsen coordinates.

The S transformation on line operators could be obtained directly by “gluing together” two copies of the RG-wall relations (6.11). This was done in [20]. But it is actually easier to work directly with tetrahedra.

Following Appendix A.3.5, we first quantize and invert relations (6.15)–(6.17), along with the gluing relation $c = z'w'r's'$ and its canonical conjugate $\gamma = 1/\sqrt{r''s''}$:

$$(6.26) \quad \hat{z} = iq^{\frac{1}{4}}\sqrt{\hat{\mu}}\hat{\lambda}\hat{\gamma}, \quad \hat{w} = iq^{\frac{1}{4}}\frac{\sqrt{\hat{\mu}}\hat{\gamma}}{\hat{\lambda}}, \quad \hat{r} = iq^{\frac{1}{4}}\frac{\hat{\lambda}'}{\sqrt{\hat{c}\hat{\mu}}}\hat{\gamma}, \quad \hat{s} = \frac{iq^{\frac{1}{4}}}{\hat{\lambda}'\sqrt{\hat{c}\hat{\mu}}}\hat{\gamma},$$

$$\hat{z}'' = \frac{1}{q^{\frac{1}{8}}\sqrt{\hat{\lambda}}}\frac{\hat{\tau}_\mu\sqrt{\hat{\tau}}}{\hat{\gamma}}, \quad \hat{w}'' = \frac{\sqrt{\hat{\lambda}}}{q^{\frac{1}{8}}}\frac{\hat{\tau}_\mu}{\hat{\gamma}\sqrt{\hat{\tau}}}, \quad \hat{r}'' = \frac{1}{q^{\frac{1}{8}}\sqrt{\hat{\lambda}'}}\frac{\sqrt{\hat{\tau}'}}{\hat{\gamma}}, \quad \hat{s}'' = \frac{\sqrt{\hat{\lambda}'}}{q^{\frac{1}{8}}}\frac{1}{\hat{\gamma}\sqrt{\hat{\tau}'}}.$$

We then take the left ideal generated by $\hat{z}'' + \hat{z}^{-1} - 1 \simeq 0$ (etc.) for each tetrahedron, rewrite it in terms of $\hat{\lambda}, \hat{\tau}, \hat{\lambda}', \hat{\tau}', \hat{\mu}, \hat{\tau}'_\mu, \hat{c}, \hat{\gamma}$, eliminate $\hat{\gamma}$, and set $\hat{c} \rightarrow 1$. To simplify further, we eliminate two out of three of the twist

operators $\hat{\tau}, \hat{\tau}', \hat{\tau}_\mu$, obtaining

$$(6.27a) \quad \hat{\lambda}' + \hat{\lambda}'^{-1} \simeq \frac{\mathfrak{s}(\hat{\lambda}\sqrt{\hat{\mu}}/(iq^{\frac{1}{4}}))}{\mathfrak{s}(\hat{\lambda})} q^{-\frac{1}{8}} \sqrt{\hat{\lambda}} \frac{1}{\sqrt{\hat{\tau}}} + \frac{\mathfrak{s}(iq^{\frac{1}{4}}\hat{\lambda}'\sqrt{\hat{\mu}})}{\mathfrak{s}(\hat{\lambda})} q^{-\frac{1}{8}} \frac{1}{\sqrt{\hat{\lambda}}} \sqrt{\hat{\tau}}$$

$$(6.27b) \quad \hat{\lambda} + \hat{\lambda}^{-1} \simeq \frac{\mathfrak{s}(\hat{\lambda}'/(iq^{\frac{1}{4}}\sqrt{\hat{\mu}}))}{\mathfrak{s}(\hat{\lambda}')} q^{-\frac{1}{8}} \sqrt{\hat{\lambda}'} \frac{1}{\sqrt{\hat{\tau}'}} + \frac{\mathfrak{s}(iq^{\frac{1}{4}}\hat{\lambda}'\sqrt{\hat{\mu}})}{\mathfrak{s}(\hat{\lambda}')} q^{-\frac{1}{8}} \frac{1}{\sqrt{\hat{\lambda}'}} \sqrt{\hat{\tau}'}$$

where

$$(6.28) \quad \mathfrak{s}(x) := x - x^{-1}.$$

The first relation sets the spin-1/2 Wilson loop on one side of the S-wall equal to the spin-1/2 't Hooft loop on the other, and the second does the opposite. They appear in [50] and (with slightly different normalization) in [49].

Note that the relations are invariant separately under $(\hat{\lambda}, \hat{\tau}) \mapsto (\hat{\lambda}^{-1}, \hat{\tau}^{-1})$. This Weyl invariance reflects the enhanced symmetry $SU(2)_\lambda \times SU(2)_{\lambda'}$. On the other hand, Weyl invariance in $\hat{\mu}$ is (predictably) absent. To get from one equation to the other we send $(\hat{\lambda}, \hat{\tau}) \leftrightarrow (\hat{\lambda}', \hat{\tau})$ along with $\hat{\mu} \leftrightarrow \hat{\mu}^{-1}$.

If we eliminate $\hat{\tau}, \hat{\tau}'$ from our left ideal but not $\hat{\tau}_\mu$, we get a third independent equation associated to “bridge” between the two loops of the Hopf network M_S :

$$(6.29) \quad \mathfrak{s}(q^{\frac{1}{4}}\sqrt{\hat{\mu}}) \left(\hat{\lambda}^2 + \frac{1}{\hat{\lambda}^2} + \frac{\hat{\mu}}{q^{\frac{1}{2}}} + \frac{q^{\frac{1}{2}}}{\hat{\mu}} \right) \frac{q^{\frac{1}{4}}}{\sqrt{\hat{\mu}}} \frac{1}{\hat{\tau}_\mu} - \mathfrak{s}(\hat{\mu}) \left(\hat{\lambda} + \frac{1}{\hat{\lambda}} \right) \left(\hat{\lambda}' + \frac{1}{\hat{\lambda}'} \right) \\ + \mathfrak{s}(q^{-\frac{1}{4}}\sqrt{\hat{\mu}}) \left(\hat{\lambda}'^2 + \frac{1}{\hat{\lambda}'^2} + q^{\frac{1}{2}}\hat{\mu} + \frac{1}{q^{\frac{1}{2}}\hat{\mu}} \right) q^{\frac{1}{4}}\sqrt{\hat{\mu}}\hat{\tau}_\mu \simeq 0.$$

This equation has the following physical interpretation. So far we have considered the 3d theory $T_2[M_S, \mathfrak{t}, \Pi_S]$ to be a duality wall between two copies of 4d $\mathcal{N} = 2^*$ theory. However, if we additionally gauge the $SU(2)_\mu$ flavor symmetry in the bulk (which is enhanced from $U(1)_\mu$ after the usual coupling to bulk hypers), we obtain a 3d boundary condition for the 4d theory of class \mathcal{S} associated to a closed genus-2 surface. This surface is the complete topological boundary ∂M_S of the S-manifold, so the 4d theory is $T_2[\partial M_S]$, and the full boundary condition is the complete/canonical $\mathcal{B}_2[M_S]$ (*cf.* Section 5.3). Equation (6.29) is a Ward identity involving the dynamical (spin-1) 't Hooft operator of $SU(2)_\mu$, brought to the boundary.

The three Equations (6.27), (6.29) annihilate the index and ellipsoid partition functions of $T_2[M_S, \mathbf{t}, \Pi_S]$, *i.e.* the index and ellipsoid partition functions of $T[SU(2)]$. These are well known expressions, having been obtained both via supersymmetric localization [48, 57] and relations to Liouville theory and quantum groups (*cf.* [15, 50]). Now we can obtain the obtain partition functions as $SL(2)$ Chern-Simons wavefunctions of M_S , using methods of [5, 7, 34]. For example, using the triangulation from Section 6.2, the ellipsoid partition function is written as

$$(6.30) \quad \mathcal{Z}_{\hbar}[M_S, \mathbf{t}, \Pi_S] = \int dRdS e^{\frac{1}{\hbar} \left[\frac{1}{2}(\Lambda^2 + \Lambda'^2 - M^2) + 2RS + M(R-S) + (i\pi + \frac{\hbar}{2})(R+S) \right]} \times \Phi_{\hbar}(-R \pm \Lambda) \Phi_{\hbar}(-S \pm \Lambda').$$

It is equivalent to the more standard partition function of $T[SU(2)]$ from [48, 57]. Here $(\lambda, \lambda', \mu') = \exp(\Lambda, \Lambda', M)$, and the respective operators $\hat{\tau}, \hat{\tau}', \hat{\tau}'_{\mu}$ act by shifting these three variables: $\hat{\tau}f(\Lambda) = f(\Lambda + \hbar)$, etc.

7. Example 2: interfaces for $SU(2)$ $N_f = 4$

Our last set of examples includes the RG and duality walls for 4d $\mathcal{N} = 2$ $SU(2)$ theory with $N_f = 4$ flavors of hypermultiplet matter.

The $SU(2)$ $N_f = 4$ theory has a rank-1 six-dimensional realization based on a four-punctured sphere \mathcal{C} . This construction only manifests an $SU(2)_a \times SU(2)_b \times SU(2)_c \times SU(2)_d$ subgroup of the full $SO(8)$ flavor group of the four-dimensional theory. The full flavor symmetry is only recovered in the infrared. Correspondingly, generic boundary conditions and domain walls for $SU(2)$ $N_f = 4$ built from three-manifolds will only preserve the $SU(2)_a \times SU(2)_b \times SU(2)_c \times SU(2)_d$ subgroup of the flavor symmetry. Special examples such as S-duality or RG walls, though, should preserve the full $SO(8)$ flavor symmetry once coupled to the bulk. This expected symmetry enhancement will provide beautiful consistency checks throughout this section. Momentarily, we will discuss carefully the possible flavor symmetry enhancements of the various theories we will encounter in this section.

In order to couple with a UV copy of $SU(2)$ $N_f = 4$, a framed 3-manifold must include two three-punctured spheres, joined by a small annulus which produces the flavor symmetry $SU(2)_{\lambda}$ to be gauged in the bulk. The remaining punctures of the three-punctured spheres must be endpoints of annuli associated to appropriate $SU(2)$ flavor subgroups. We can readily describe the 3-manifolds M_{RG} and M_{F} which correspond to an RG domain wall and to

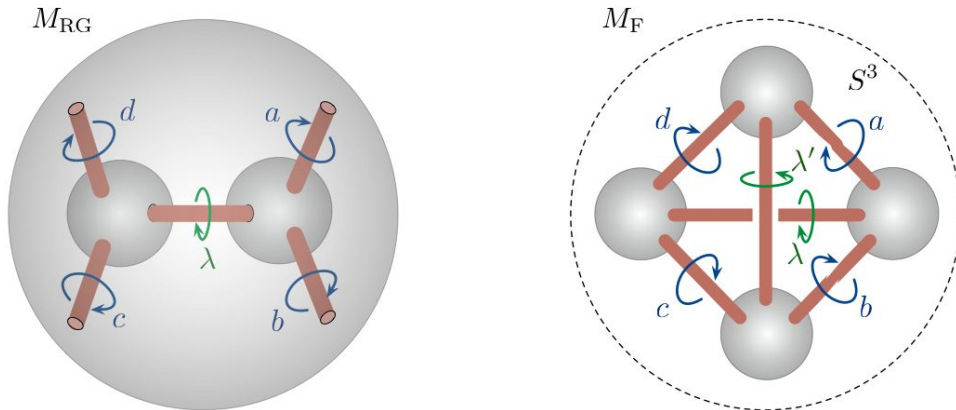


Figure 40: The RG and S-duality manifolds for $SU(2)$ $N_f = 4$ theory.

an S-duality wall. The former must have a big boundary composed of a four-punctured sphere and two three-punctured spheres. We will have an $SU(2)_\lambda$ small annulus joining the three-punctured spheres and the four $SU(2)_{a,b,c,d}$ small annuli joining the three-punctured spheres and the four-punctured sphere. We can envision the four-punctured sphere as the boundary of a solid ball, with the three-punctured spheres and the annuli carved out of it (Figure 40). The boundary phase space $\mathcal{P}_2(\partial M_{\text{RG}})$ is twelve-dimensional and it includes the holonomy and twist parameters for the five annuli and an extra pair of edge coordinates from the four-punctured sphere.

In order to obtain the 3-manifold M_{F} for an S-duality wall, we can glue together two RG domain wall geometries, permuting the four punctures accordingly. The two balls glue into a solid three-sphere, with a network of annuli and three-punctured sphere carved out of it. The network has the topology of a tetrahedral graph, with a pair of opposite edges labeled by the to-be-gauged $SU(2)_\lambda$ and $SU(2)_{\lambda'}$ symmetries and other opposite pairs labelled by $SU(2)_a$ and $SU(2)_c$ or $SU(2)_b$ and $SU(2)_d$ respectively. The boundary phase space $\mathcal{P}_2(\partial M_{\text{F}})$ is again twelve-dimensional and is parameterized by holonomy and twist coordinates for all six annuli in the network. We use the subscript ‘F’ for this manifold because its $SL(2)$ Chern-Simons wavefunction coincides with the fusion kernel in Liouville theory.

Of course, the actual flavor symmetries of the corresponding three-dimensional theories will depend on the choices of triangulations of the big boundaries. In order to identify useful choices of triangulations, we should consider the full flavor symmetry of the problem.

7.1. Enhanced flavor symmetry and $SO(12)$ surprises

7.1.1. UV considerations. The four $SU(2)$'s that are manifest in six dimensions combine pairwise into a block-diagonal $SO(4) \times SO(4)$ subgroup of the infrared $SO(8)$. Actually, all three possible pairwise combinations are OK: they correspond to $SO(4) \times SO(4)$ subgroups related by triality. The three basic representations of $SO(8)$ decompose in terms of doublet representations of the $SU(2)_a \times SU(2)_b \times SU(2)_c \times SU(2)_d$ flavor subgroup as

$$(7.1) \quad \begin{aligned} 8_v &= 2_a \otimes 2_b + 2_c \otimes 2_d & 8_s &= 2_a \otimes 2_d + 2_b \otimes 2_c \\ 8_c &= 2_a \otimes 2_c + 2_b \otimes 2_d \end{aligned}$$

Correspondingly, the fundamental hypermultiplets in the standard S-duality frame decompose into two trinion blocks

$$(7.2) \quad 2_\lambda \otimes 8_v = 2_\lambda \otimes 2_a \otimes 2_b + 2_\lambda \otimes 2_c \otimes 2_d$$

The theory has two other S-duality frames where the fundamental hypers transform in 8_s and 8_c , which correspond to the other two pair of pants decompositions of the four-punctured sphere.

We can ask how to insure that a 3d theory associated to some three-manifold will give rise to a boundary condition or domain wall that preserves the full $SO(8)$ flavor symmetry. If we consider a coupling to the UV description of the theory, the triangulations of three-punctured spheres (partially) determine which symmetries are present. At the very least, we can try to pick triangulations to preserve all flavor symmetry: $U(1)_\lambda \times SU(2)_a \times SU(2)_b \times SU(2)_c \times SU(2)_d$. This choice has the price that the $SU(2)_\lambda$ symmetry to be gauged only appears in the infrared, after coupling to the bulk hypers. One may then hope to find a description of the isolated 3d theory where the $SU(2)_a \times SU(2)_b \times SU(2)_c \times SU(2)_d$ is explicitly extended to $SO(8)$, in such a way that the operators used in the coupling to 4d hypers form a vector 8_v of $SO(8)$, which is coupled to an $SO(8)$ -invariant half of the bulk hypermultiplets.

In our concrete examples, however, we will use a different mechanism to show the full $SO(8)$ invariance of the interface. One can pick the three-punctured sphere triangulations to preserve an $SU(2)_\lambda \times SU(2)_a \times U(1)_b \times SU(2)_c \times U(1)_d$, *i.e.* the $U(2) \times U(2)$ subgroup of $SO(4) \times SO(4)$. The splitting of the bulk hypermultiplets $Q_{\alpha\beta\ell}$ and $Q_{\ell\gamma\delta}$ into halves actually preserves more than $U(2) \times U(2)$: we can combine the halves into a fundamental of $U(4) \subset SO(8)$ and an anti-fundamental (both doublets of the gauge group).

It is then natural to look for 3d theories whose $U(2) \times U(2)$ symmetry group is correspondingly enhanced to a full $U(4)$, in such a way that the operators we couple to the hypermultiplets form a fundamental of $SU(4)$ and the bulk-boundary superpotential coupling manifestly preserves $SU(4)$ as well. If that is the case, the standard promotion of $U(1)_b$ and $U(1)_d$ to $SU(2)_b$ and $SU(2)_d$ (upon coupling to the bulk) will insure that $U(4)$ is promoted to a full $SO(8)$.

7.1.2. IR considerations. If we consider the coupling to an IR Seiberg-Witten description of the theory, we can make flavor symmetry enhancements manifest only if the spectrum of BPS particles is organized in representations of the desired flavor symmetry group. In order to understand why this is a subtle requirement, observe that the BPS spectrum for generic values of the mass parameters has no reason to form full representations of $SO(8)$, as the mass parameters themselves break the $SO(8)$ flavor symmetry. Making all mass parameters small (compared to the Coulomb branch parameter) makes the BPS spectrum very intricate and the association to WKB triangulations unpredictable.

A simple compromise is to give all four (UV) hypers the same large mass, which preserves $U(4)$. Then the BPS spectrum organizes itself in $U(4)$ representations. For example, the Coulomb branch has three singular loci. At one locus we find 4 light BPS particles in a fundamental representation of $SU(4)$, which are identified with fundamental hypers in the UV theory. The other two loci are basically the usual monopole and dyon points for the pure $SU(2)$ theory, each with a light BPS particle that is a $SU(4)$ singlet. The BPS spectrum is then robust against adding some small $U(4)$ breaking mass, and we can easily find WKB triangulations whose edges carry appropriate $U(4)$ quantum numbers.

It is useful to start from a “square” triangulation, where the four punctures are the vertices of a square. We add two diagonal edges between the b and d vertices on opposite sides of the sphere to complete the triangulation. Then we flip opposite edges of the square to get degenerate triangles around a and c , so that $SU(2)_a$ and $SU(2)_b$ can (potentially) be preserved in a 3d theory (Figure 41).

The degenerate triangles will be associated to two doublets of 3d operators that have the correct quantum numbers to couple to the 4d BPS particles in a fundamental of $SU(4)$. The operators associated to the diagonal edges do not carry $SU(4)$ charge. This is the triangulation of the four sphere that we will use to realize the RG domain wall. If we can present the three-dimensional theory and the couplings to bulk fields in a way that




Figure 41: A triangulation of the IR boundary of M_{RG} that can preserve $SU(2)_a \times SU(2)_c$.

preserves $U(4)$ rather than just $SU(2)_a \times U(1)_b \times SU(2)_c \times U(1)_d$, we will recover $SO(8)$ in the IR. We will do so in the next section.

7.1.3. Surprises. In the case of the S-duality wall, say between the “ 8_v ” and “ 8_s ” duality frames, we can simply triangulate the three-spheres to preserve $SU(2)_\lambda \times SU(2)_{\lambda'} \times SU(2)_a \times U(1)_b \times SU(2)_c \times U(1)_d$ (see Figure 44 below). The corresponding theory has a chance to preserve $U(4)$ and thus give a domain wall with $SO(8)$ flavor symmetry. The possibility of $SO(8)$ enhancement upon coupling to the bulk hypers has a very entertaining consequence. Suppose that we ignore the fact that at some point we will want to gauge the $SU(2)_\lambda \times SU(2)_{\lambda'}$ flavor symmetries, and we simply use the three-dimensional theory to define a boundary condition for four trinion theories. The tetrahedral graph has a high degree of symmetry, and the final boundary condition should treat democratically all vertices and all pairs of edges of the tetrahedron.

Even before we couple the three-dimensional theory to the bulk hypers, the tetrahedral symmetry combines with symmetry enhancements in a neat way. Both the subgroups $SU(2)_a \times U(1)_b \times SU(2)_c \times U(1)_d$ and $SU(2)_\lambda \times SU(2)_{\lambda'} \times U(1)_b \times U(1)_d$ want to be promoted to $U(4)$ s. This is possible only if the overall symmetry of the 3d theory is enhanced to $SU(4) \times SU(4)$.

Once we couple to hypers, the flavor enhancement is striking. The $SU(2)$ flavor symmetries associated to any two pairs of opposite edges should belong to some $SO(8)$. We have three $SO(4) \simeq SU(2) \times SU(2)$ groups, and each pair of $SO(4)$ ’s must sit block-diagonal in an $SO(8)$ flavor group. This is possible, provided that the flavor group of the boundary condition is actually enhanced to $SO(12)$! Notice that the bulk theory of four trinions indeed has a $SO(12)$ flavor symmetry (subgroup of the full $USp(32)$): we can assemble




Figure 42: Triangulations of 3-punctured spheres in the tetrahedral graph that should lead to 3d theories $T_2[M_F]$ with $SU(4) \times SU(4)$ and $U(6)$ enhanced symmetry.

them all into a (pseudo real, 32 dimensional) chiral spinor of $SO(12)$. Indeed, such a spinor decomposes under $SO(4) \times SO(4) \times SO(4)$ as

$$(7.3) \quad (32) = 2_L \otimes 2_L \otimes 2_L + 2_L \otimes 2_R \otimes 2_R + 2_R \otimes 2_L \otimes 2_R + 2_R \otimes 2_R \otimes 2_L.$$

Notice that each $SU(2)$ group links two trinions, and the whole structure reproduces the tetrahedral network. The 3d geometry conjecturally represents a boundary condition for the theory of 32 half-hypermultiplets which preserves this $SO(12)$ subgroup of their full $USp(32)$ flavor symmetry. This is a beautiful object, whose existence was proven in [17].

The $SO(12)$ symmetry enhancement suggests a second interesting way to pick the triangulations of the three-punctured spheres. The first choice we mentioned should give a 3d theory with $SU(4) \times SU(4) \sim SO(6) \times SO(6)$ symmetry, block-diagonal inside $SO(12)$. A second choice is to pick the triangulations so to preserve the $SU(2)$'s at three edges surrounding a single face of the tetrahedron. These are the three right-chiral $SU(2)$'s in $SO(4)^3$. All these $SU(2)$'s can fit inside an $SU(6)$ subgroup of $SO(12)$. Furthermore, the spinor of $SO(12)$ decomposes as $(1) + (15) + (\bar{15}) + (1)$ under $SU(6)$. All the 4d half-hypers that receive a modified Dirichlet b.c. happen to sit inside one (15) and one singlet of $SU(6)$. Correspondingly, we expect the 3d SCFT to have an $SU(6)$ flavor symmetry in the IR.

It is interesting to observe that the authors of [58] propose a candidate theory that should be related to this second option. The proposal is based on the identification between the ellipsoid partition function of

the 3d theory and the fusion kernel in Liouville theory. Up to a prefactor³⁶, the partition function coincides with the expression one would get from a 3d $SU(2)$ gauge theory with six quark flavors, which has explicit $U(6)$ flavor symmetry. Looking at the fugacities in the partition function, denoted there as α_i , with $i = 1, 2, 3, 4, s, t$, such that the four vertices of the tetrahedron are attached to the defects with label $(1, 2, s)$, $(3, 4, s)$, $(2, 3, t)$, $(1, 4, t)$ and (s, t) , $(1, 3)$, $(2, 4)$ are opposite pairs, one can see explicitly the $SU(2)_s \times SU(2)_2 \times SU(2)_3$ block-diagonal subgroups of $U(6)$, as expected. The other fugacities give charge $1/2$ to four quarks, $-1/2$ to two of them, cyclically.

The 3d $SU(2)$ gauge theory with six quark flavors has meson operators in the (15) of the $SU(6)$ flavor symmetry. These chiral operators may well be coupled to the (15) half-hypermultiplets in the bulk. A monopole operator of the 3d $SU(2)$ gauge theory will only carry a charge under the $U(1)$ diagonal subgroup of $U(6)$. This means that it will have charge 1 under all three $U(1)_t \times U(1)_1 \times U(1)_4$. This may couple to the remaining half-hypermultiplet in a singlet of $SU(6)$. The existence of the correct set of chiral operators is a strong suggestion that the theory proposed by the authors of [58] is truly a mirror description of the (abelian!) class \mathcal{R} theory associated with our 3d geometry, in the \mathbb{Z}_3 -symmetric boundary triangulation.

There is another well known formula for the fusion kernel in Liouville theory that inspires a 3d field theory interpretation. This is given, *e.g.* in [58, Eqn 2.17]. Translated directly, the formula suggests a 3d $N_f = 4$ SQED, possibly with an extra superpotential coupling which breaks both the topological and axial flavor symmetries. This can be done, say, by adding the basic monopole and anti-monopole operators. Looking up the charge assignments of monopoles in [47] we verify that the axial fugacity has been fixed so that the monopoles are marginal.

Notice that the residual flavor symmetry of the $N_f = 4$ SQED is $SU(4)_L \times SU(4)_R$, i.e. $SO(6) \times SO(6)$, just as we expect from the first triangulation of 3-punctured spheres discussed above. Indeed, the ellipsoid partition function of the theory suggests that $SU(2)_s \times SU(2)_t$ is block diagonal inside $SU(4)_R$, i.e. an $SO(4)$ block in $SO(6)$, and the same for $SU(2)_1 \times SU(2)_3$.

³⁶Modulo a typo in [58], this prefactor corresponds to flipping the triangulations of 3-punctured spheres that should preserve $SU(6)$ back to democratic ones, using the rules of Figure 33. The actual fusion kernel in its most standard normalization corresponds to all democratic triangulations.

Next, we provide the gory details of the triangulated manifolds and the corresponding abelian CSM descriptions of the RG and S-duality kernels.

7.2. The RG wall theory

We can triangulate the RG manifold M_{RG} with six tetrahedra in such a way that the induced big-boundary triangulation \mathbf{t} is compatible with $SU(2)_\lambda \times SU(2)_a \times U(1)_b \times SU(2)_c \times U(1)_d$. We show the triangulation in Figure 43. Although the triangulation looks complicated, and the careful determination of the various couplings will take a little work, the result will be amazingly simple, and can be described right away. Indeed, the 3d theory basically coincides with the RG theory for the pure $SU(2)$ gauge theory, together with four auxiliary chiral fields whose only purpose is to glue together at the boundary the four electrically charged BPS particles in the $SU(4)$ fundamental and the corresponding UV hypermultiplets, much as we saw happen for the $N = 2^*$ theory.

The triangulation is simply built upon the basic RG manifold M_0 for pure $SU(2)$ theory. Its two tetrahedra are assigned parameters z and w . Two more tetrahedra (parameters r, s) are attached to the ends of M_0 , initially identifying only edges and not faces, in such a way that two three-punctured spheres get trapped in the gluing. These extra tetrahedra can be seen as halves of the basic M_0 building block. Subsequently, the half- M_0 tetrahedra are folded back onto the full M_0 core, making the face identifications labeled \mathcal{B} and \mathcal{C} . This leads to a geometry M'_{RG} that is almost correct, but has two thin edges \triangleright and $\triangleright\triangleright$, just like the thin edge in the $\mathcal{N} = 2^*$ RG manifold (Figure 35). The thin edges are thickened out by adding two more tetrahedra (parameters f and g), thus flipping the IR boundary triangulation precisely in the manner of Figure 41.

The triangulation has no internal edges and a brief inspection shows a natural polarization of the IR boundary such that no dynamical gauge fields are needed: the theory consists of six free chiral fields! The eigenvalues and dual twists are related to tetrahedron parameters as

$$\begin{aligned}
 (7.4) \quad \lambda^2 &= \frac{z}{w}, & a^2 &= \frac{f}{r}, & (-b)^2 &= zwr f \\
 c^2 &= \frac{g}{s}, & (-d)^2 &= zws g, & x_e &= \frac{1}{\sqrt{zw}}, \\
 \tau &= \lambda \frac{z''}{w''}, & \tau_a &= a \frac{f''}{r''}, & \tau_b &= bx_e r'' f'', \\
 \tau_c &= c \frac{g''}{s''}, & \tau_d &= dx_e s'' g'', & x_m &= bdz' w' r'' s'' f'' g'',
 \end{aligned}$$

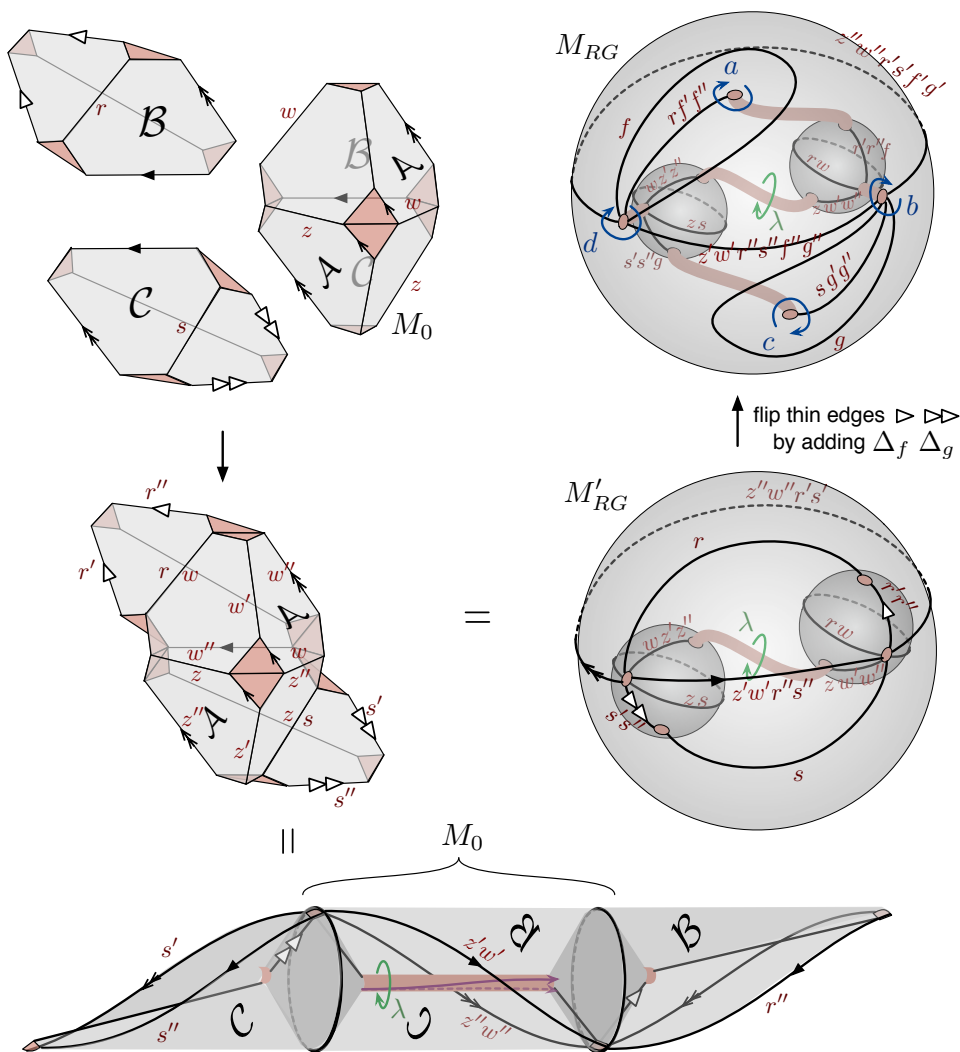


Figure 43: Construction of the $N_f = 4$ RG manifold M_{RG} from six tetrahedra, starting with a basic M_0 building block. The big-boundary edges of M'_{RG} and M_{RG} are labelled by the products of tetrahedron edge-parameters incident there.

where we emphasize that only the twists τ, τ_a, τ_c for the annuli that end in degenerate triangles (and have enhanced $SU(2)$ symmetry) are canonically defined. We have also chosen an electric coordinate x_e to associate to the IR $U(1)$ gauge charge (well defined up to flavor charges), and a conjugate magnetic coordinate x_m . We let Π_e be the polarization with positions

λ, a, b, c, d, x_e and the τ 's and x_m as momenta. The theory $T_2[M_{\text{RG}}, \mathbf{t}, \Pi_e]$ requires no gauging because all the position coordinates are made from tetrahedron positions.

The charge assignment for the six chiral fields $(\phi_z, \phi_w, \phi_r, \phi_s, \phi_f, \phi_g)$ is obtained by inverting the position relations in (7.4). It is most easily expressed by grouping the chirals into three doublets $q_\alpha = (\phi_z, \phi_w)$, $q'_\beta = (\phi_f, \phi_r)$, $q''_\gamma = (\phi_g, \phi_s)$. Then charges are

$$(7.5) \quad \begin{array}{c|ccc} & q_\alpha & q'_\beta & q''_\gamma \\ \hline SU(2)_\lambda & \square & - & - \\ SU(2)_a & - & \square & - \\ U(1)_b & 0 & 1 & 0 \\ SU(2)_c & - & - & \square \\ U(1)_d & 0 & 0 & 1 \\ U(1)_x & -1 & 1 & 1 \\ \hline U(1)_R & 0 & 1 & 1 \end{array}$$

With momenta as above, there is also a level -1 background Chern-Simons term for the combination $A_x - A_R$ of abelian gauge fields; but there are no Chern-Simons terms for any of the other symmetries.

We see manifestly the pure $SU(2)$ domain wall theory built from the first pair of chirals q_α , together with a fundamental representation $u_A = (q'_\beta, q''_\gamma)$ of a full $U(4)$ flavor group, of opposite $U(1)_x$ abelian gauge charge. Then the eight bilinear operators $q_\alpha u_A$ are associated to the three-punctured spheres and couple neatly to the UV hypermultiplets. Moreover, for the four-punctured sphere on the IR boundary, the four operators u_A are associated to the degenerate triangles, while the two long diagonal edges are associated to appropriate operators in a monopole background.

If we focus on the electric superpotential couplings, we will find a situation analogous to the $\mathcal{N} = 2^*$ case:

$$(7.6) \quad W_{\text{bdy}} = X^{\alpha A} \Big|_{\partial} q_\alpha u_A + X_e^A \Big|_{\partial} u_A$$

which enforces the notion that the quadruplet of light abelian particles X_e^A coincide with the projection of the UV hypers along the direction q_α of the IR abelian gauge field inside the UV $SU(2)$ gauge field.

7.3. The S-duality wall

We can triangulate the geometry M_F in order to reproduce 3d $N_f = 4$ SQED on the nose, with the monopole and anti-monopole superpotentials arising

from two internal edges of the triangulation. Notice that as all chirals appear in doublets of a flavor symmetry, the geometry should somehow be built out of four copies of the basic RG manifold M_0 . It takes a certain amount of geometric imagination to guess the construction. Luckily, the physics does the job for us: the gluing constraints must take the form $\prod_{i=1}^4 z'_{i,+} z''_{i,-} = 1$ and $\prod_{i=1}^4 z''_{i,+} z'_{i,-} = 1$ in order to correspond to the monopoles of SQED. This allows us to get the correct geometry, depicted in Figure 44.




Figure 44.

To build M_F , we place four copies of M_0 end-to-end in a loop, trapping four three-punctured spheres at their interfaces. The four annuli in these M_0 's carry the $SU(2)_\lambda \times SU(2)_a \times SU(2)_c \times SU(2)_{\lambda'}$ symmetries that we want to enhance. Then we fold this loop over on itself (making face identifications A, B, C, D) to obtain the complement of the tetrahedral graph in S^3 . Alternatively, the gluing of M_F can be obtained by taking two copies of the RG manifold M'_RG , before its edges were thickened, and identifying their IR boundaries, as at the top of Figure 44. We'll give the four tetrahedra in one copy of M_RG parameters z, w, r, s and in the other copy $\tilde{z}, \tilde{w}, \tilde{r}, \tilde{s}$.

With a little work, we find the relations between eigenvalue/twist coordinates and tetrahedron parameters

$$\begin{aligned}
 (7.7) \quad \lambda^2 &= \frac{\tilde{z}}{w}, & a^2 &= \frac{\tilde{r}}{r}, & (-b)^2 &= zwr\tilde{r}, \\
 c^2 &= \frac{\tilde{s}}{s}, & (-d)^2 &= zws\tilde{s}, & \lambda'^2 &= \frac{\tilde{\tilde{z}}}{\tilde{w}}; \\
 \tau &= \lambda \frac{z''}{w''}, & \tau_a &= a \frac{\tilde{r}''}{r''}, & \tau_b^2 &= \frac{z''w''s'\tilde{s}'}{z'w's''\tilde{s}''}, \\
 \tau_c &= c \frac{\tilde{s}''}{s''}, & \tau_d^2 &= \frac{z''w''r'\tilde{r}'}{z'w'r''\tilde{r}''}, & \tau' &= \lambda' \frac{\tilde{\tilde{z}}'}{\tilde{w}''}.
 \end{aligned}$$

Now all of the twists are canonically defined, because all annuli begin and end at three-punctured spheres. Moreover, the gluing constraints c_1 and c_2 , coming from internal edges labelled \blacktriangleright and \ggg in the figure, take the promised form

$$(7.8) \quad c_1 = z'\tilde{z}'w'\tilde{w}'r''\tilde{r}''s''\tilde{s}'' , \quad c_2 = z''\tilde{\tilde{z}}''w''\tilde{w}''r'\tilde{r}'s'\tilde{s}' .$$

With a polarization Π_F in which all the eigenvalues are position coordinates, we use the relations (7.7)-(7.8) to find that $T_2[M_F, \mathbf{t}, \mathbf{\Pi}]$ is a copy of $\mathcal{N}_f = 4$ SQED with eight chirals $u_\alpha = (\phi_z, \phi_w)$, $u'_\beta = (\phi_{\tilde{z}}, \phi_{\tilde{w}})$, $v_\gamma = (\phi_{\tilde{r}}, \phi_r)$, $v'_\delta = (\phi_{\tilde{s}}, \phi_s)$ of charges

$$(7.9) \quad \begin{array}{c|cccc} & u_\alpha & u'_\beta & v_\gamma & v'_\delta \\ \hline SU(2)_\lambda & \square & - & - & - \\ SU(2)_{\lambda'} & - & \square & - & - \\ SU(2)_a & - & - & \square & - \\ SU(2)_c & - & - & - & \square \\ U(1)_b & \frac{1}{2} & -\frac{1}{2} & \frac{1}{2} & -\frac{1}{2} \\ U(1)_d & \frac{1}{2} & -\frac{1}{2} & -\frac{1}{2} & \frac{1}{2} \\ \hline U(1)_{\text{gauge}} & 1 & 1 & -1 & -1 \\ \hline U(1)_R & 1 & 1 & 1 & 1 \end{array}$$

The monopole and anti-monopole in the superpotential break the axial and topological symmetries associated to c_1 and c_2 , and there are no background Chern-Simons terms of any kind.

The couplings of the bulk hypermultiplets to 3d $N_f = 4$ SQED are readily described. The natural gauge-invariant chiral bilinears, or mesons, carry an index of both $SU(4)$ fundamental representations, *i.e.* a spinor index of both $SO(6)$'s inside the $SO(12)$ symmetry that is expected to emerge

after the coupling. This is half of the $SO(12)$ spinor representation of the bulk operators, and thus we have a natural coupling, which can be reproduced from the geometry. Notice that in the tetrahedral graph each of the quartets of bilinears $u_\alpha v_\gamma$, $u_\alpha v'_\delta$, $u'_\beta v_\gamma$, $u'_\beta v'_\delta$ is associated to one of the three-punctured spheres.

7.3.1. Line operators. We may compute the Ward identities for line operators at the S-duality domain wall using a logarithmic version of relations (7.7)-(7.8) and the standard elimination procedure. The classical result, at $q = 1$, sets the Wilson loop $\lambda + \lambda^{-1}$ on one side equal to the 't Hooft loop on the other,

$$(7.10) \quad \lambda + \lambda^{-1} = \frac{1}{\mathfrak{s}(\lambda')^2} \left[-(\mathfrak{c}(c) + \mathfrak{c}(\lambda'/b))(\mathfrak{c}(a) + \mathfrak{c}(\lambda'/d)) \frac{1}{\tau'} \right. \\ \left. + \mathfrak{c}(\lambda')(\mathfrak{c}(a)\mathfrak{c}(c) + \mathfrak{c}(b)\mathfrak{c}(d)) \right. \\ \left. + 2(\mathfrak{c}(a)\mathfrak{c}(b) + \mathfrak{c}(c)\mathfrak{c}(d)) \right. \\ \left. - (\mathfrak{c}(c) + \mathfrak{c}(b\lambda'))(\mathfrak{c}(a) + \mathfrak{c}(d\lambda'))\tau' \right],$$

where $\mathfrak{s}(x) = x - x^{-1}$ and $\mathfrak{c}(x) = x + x^{-1}$. In Teichmüller theory, this encodes the action of the nontrivial generator of the mapping class group of a 4-punctured sphere in Fenchel-Nielsen-like coordinates.

The fully quantized operator identity appeared in [50, Sec 5.4], and is equivalent to (7.10) in the classical limit after flipping the three-punctured spheres to democratic triangulations. Using the dictionary from Figure 33 and being more careful about Chern-Simons prefactors, one sees that the flips have the effect of rescaling

$$(7.11) \quad \tau \mapsto \frac{\mathfrak{c}(\sqrt{d\lambda/c})}{\mathfrak{c}(\sqrt{b/(a\lambda)})} \tau, \quad \tau' \mapsto \frac{\mathfrak{c}(\sqrt{c\lambda'/b})}{\mathfrak{c}(\sqrt{a/(d\lambda')})} \tau'$$

inside classical operator relations.

7.4. The 6j symbol and volumes of non-ideal tetrahedra

The ellipsoid partition function of the S-wall theory $T_2[M_F, \mathfrak{t}, \Pi]$ can readily be computed as³⁷

³⁷The quantum dilogarithms Φ_{\hbar} used here are equivalent to the “double-sine” functions that one encounters in work on supersymmetric localization, Liouville theory, etc., as reviewed in Appendix C.

$$\begin{aligned}
(7.12) \quad & \mathcal{Z}_{\hbar}(\Lambda, A, B, C, D, \Lambda') \\
&= \int_{\mathbb{R}} du e^{\frac{1}{2\hbar}(u^2 + \Lambda^2 + \Lambda'^2 + A^2 + B^2 + C^2 + D^2 + (i\pi + \frac{\hbar}{2})^2)} \\
&\quad \times \Phi_{\hbar}(-u \pm \Lambda + \frac{1}{2}(B + D)) \Phi_{\hbar}(-u \pm \Lambda' - \frac{1}{2}(B + D)) \\
&\quad \cdot \Phi_{\hbar}(u \pm A + \frac{1}{2}(B - D)) \Phi_{\hbar}(u \pm C - \frac{1}{2}(B - D)),
\end{aligned}$$

where as usual the upper-case variables are logarithms of the eigenvalues, $\lambda = \exp(\Lambda)$, etc. Alternatively, this is an $SL(2)$ Chern-Simons wavefunction for the manifold M_F , obtained through quantization and symplectic gluing of tetrahedra. In terms of Chern-Simons theory, two observations from [58] about the function (7.12) acquire a beautiful geometric interpretation.

The first observation is that (7.12) coincides with a formula for the 6j symbols of the quantum group $U_q(sl_2)$, first calculated in [59]. Specifically, the 6j symbols in (2.17) of [58] are obtained from (7.12) by relating $(\Lambda, A, B, C, D, \Lambda') \leftrightarrow (\alpha_s, -\alpha_1, -\alpha_2, -\alpha_3, -\alpha_4, \alpha_t)$ and flipping the triangulation of 3-punctured spheres to the democratic one. Following Figure 33, the flips multiply the partition function by

$$(7.13) \quad \Phi_{\hbar}(B - \Lambda - A) \Phi_{\hbar}(C + \Lambda' - B) \Phi_{\hbar}(D + \Lambda - C) \Phi_{\hbar}(A - D - \Lambda'),$$

modulo some quadratic exponentials.

One may compare the relation to 6j symbols with a familiar situation in the compact case. The 6j symbols for $U_q(su(2))$ can be obtained as an $SU(2)$ Chern-Simons partition function of a tetrahedral network of Wilson lines “colored” by representations of $SU(2)$ that extend to representations of $U_q(su(2))$ [60–62]. This is sometimes called a quantum spin network. The tetrahedral network precisely encodes how the tensor product of representations of $U_q(su(2))$ can be decomposed in two different ways. Now we are encountering a noncompact analogue of this relation. The $SL(2)$ Chern-Simons wavefunction of a tetrahedral network is equivalent to the partition function of our manifold M_F . We are simply replacing the “colors” on edges of the tetrahedron with boundary conditions for Chern-Simons theory at the tubes of M_F , in the form of fixed holonomy eigenvalues $\lambda, a, b, c, d, \lambda'$. Correspondingly, we fully expect to recover the 6j symbols for $U_q(sl(2))$, where the holonomy eigenvalues are weights of principal-series representations.

We may also understand the Chern-Simons partition function of a tetrahedral network in terms of conformal field theory. In the compact case, the $SU(2)$ partition function would reproduce the Moore-Seiberg fusion kernel

for the $SU(2)$ WZW model. Now, in the $SL(2)$ case, with our particular Chern-Simons boundary conditions, we would expect to get the fusion kernel for Liouville theory. But, due to work of [63], this is equivalent to the 6j symbol of $U_q(sl_2)$.

The second observation of [58] is that in the classical $\hbar \rightarrow 0$ limit, the leading asymptotics of (7.12) reproduce the hyperbolic volume of a *non-ideal* tetrahedron. This can easily be understood from Chern-Simons theory by noting that the semi-classical limit of a partition function on M_F must reproduce the volume of a flat $SL(2)$ connection, *i.e.* the (complexified) hyperbolic volume of M_F . The boundary conditions at the six small annuli of M_F are such that a hyperbolic metric would have deformed cusps, with deficit angles equal to the argument of the eigenvalues $\lambda, a, b, c, d, \lambda'$. We can then cut M_F symmetrically into two halves along geodesic surfaces lying on the *faces* of the tetrahedral network, obtaining two non-ideal hyperbolic tetrahedra $\Delta, \bar{\Delta}$ of opposite orientation. The dihedral angles of each tetrahedron are half the arguments of $\lambda, a, b, c, d, \lambda'$. Then for $\hbar \in i\mathbb{R}_{>0}$ we expect

$$(7.14) \quad \begin{aligned} |\mathcal{Z}_{\hbar}(\Lambda, A, B, C, D, \Lambda')| &\stackrel{\hbar \rightarrow 0}{\sim} e^{-\frac{1}{|\hbar|} \text{Vol}(M_F)}, \\ \text{Vol}(M_F) &= \text{Vol}(\Delta) + \text{Vol}(\bar{\Delta}) = 2\text{Vol}(\Delta), \end{aligned}$$

where the absolute value isolates the real hyperbolic volume of the geometries M_F and Δ .

Acknowledgements

We would like to thank Christopher Beem, Clay Cordova, Alexander Goncharov, Sergei Gukov, Lotte Hollands, Gregory Moore, Andrew Neitzke, and Edward Witten for enlightening discussions during the course of this project. TD would like to thank the U.C. Berkeley mathematics and physics departments and the Perimeter Institute for Theoretical Physics for generously hosting his visits there in 2012–2013 to collaborate on this project. RV would similarly like to thank the Institute for Advanced study for hosting him in March 2012. The research of TD was supported in part by a William D. Loughlin membership at the Institute for Advanced Study and in part by the Friends of the Institute for Advanced Study, with additional support from DOE grant DE-FG02-90ER-40542. The research of DG was supported in part by the Roger Dashen membership at the Institute for Advanced Study, and by the Perimeter Institute for Theoretical Physics. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of

Economic Development and Innovation. The research of RV was supported by the Netherlands Organisation for Scientific Research.

Appendix A. Coordinates from hybrid triangulations

In this appendix, we discuss the construction of coordinates for spaces of framed flat connections on framed 3-manifolds and their boundaries. We follow the framework of [16] in two dimensions and [6] in three dimensions.

Recall that a framed, oriented 3-manifold M (as defined in Section 1.1 or [6]) has two types of boundary components, “small” and “big.” The big boundary consists of surfaces of any genus, at least one hole/puncture, and with negative Euler characteristic, and can be given a 2d ideal triangulation \mathbf{t} . The small boundary consists of tori, annuli, or discs, which connect to the holes on the big boundary.

Also recall that a framed flat connection on M is a standard flat connection together with a choice of flat section of an associated flag bundle at every small boundary. This means a *global* flat section at each small boundary — *i.e.* a flag that’s invariant under the boundary holonomy. Similarly, a framed flat connection on ∂M is (by convention) a standard flat connection on $\partial M^* := (\partial M \text{ with small discs removed})$ with unipotent holonomy around every excised disc, together with a choice of flat section of the flag bundle along every remaining small boundary component and every S^1 boundary of an excised disc.

Our first goal will be to consider the algebraic moduli space $\mathcal{P}_2(\partial M)$ of framed flat $PGL(2, \mathbb{C})$ connections on the boundary ∂M , and to show that Zariski-open patches $\mathcal{P}_2(\partial M, \mathbf{t})$ (which depend on a big-boundary triangulation \mathbf{t}) can be given cluster \mathbb{C}^* coordinates. In fact, there is a factorization

$$(A.1) \quad \mathcal{P}_2(\partial M, \mathbf{t}) \simeq (\mathbb{C}^*)^{2d} \times (\mathbb{C}^*)^{2a} \times (\mathbb{C}^*)^{2t}.$$

Here the $(\mathbb{C}^*)^{2d}$ factor is the standard \mathcal{X} -coordinate space of Fock and Goncharov [16] for a big boundary with an ideal triangulation \mathbf{t} , where

$$(A.2) \quad d = 3g - 3 + 2h, \quad g = \text{genus}(\text{big bdy}), \quad h = \# \text{ holes}(\text{big bdy});$$

the middle factor consists of pairs of coordinates for each of the a small annuli of ∂M , which are a novel generalization of Fenchel-Nielsen coordinates [30, 31]; and the final factor contains pairs of coordinates for each of the t small tori.

We also describe a natural holomorphic symplectic form Ω on $\mathcal{P}_2(\partial M, \mathbf{t})$, which on every patch $\mathcal{P}_2(\partial M, \mathbf{t})$ splits into a sum of three canonical pieces

$\Omega = \Omega_{\text{big}} + \Omega_a + \Omega_t$ according to the decomposition (A.1). This form admits K-theoretic and motivic “avatars” as in [16]. The form Ω projects to the standard -Atiyah-Bott-Goldman form $\Omega = \int_{\partial M} \text{Tr } \delta \mathcal{A} \wedge \delta \mathcal{A}$ [26, 27] on *smooth* parts of the standard moduli space of unframed flat connections \mathcal{A} .

Our second goal is to define a holomorphic Lagrangian submanifold $\mathcal{L}_2(M) \subset \mathcal{P}_2(\partial M)$, via its intersections with the patches

$$(A.3) \quad \mathcal{L}_2(M, \mathbf{t}) \subset \mathcal{P}_2(\partial M, \mathbf{t}).$$

The Lagrangian $\mathcal{L}_2(M)$ describes the subset of framed flat connections on ∂M that extend to framed flat connections in the bulk M and obey certain regularity conditions — in particular that the holonomy of 3d connections is irreducible, and that choices of framing flags are generic.

Given a 3d ideal triangulation \mathbf{t}_{3d} for M (compatible with the boundary triangulation \mathbf{t}), we will relate \mathbb{C}^* coordinates on $\mathcal{P}_2(\partial M, \mathbf{t})$ to edge coordinates of tetrahedra, and thus show that

$$(A.4) \quad \mathcal{P}_2(\partial M, \mathbf{t}) = \left(\prod_i \mathcal{P}_{\partial \Delta_i} \right) // (\mathbb{C}^*)^{N_I}$$

is a symplectic reduction of a product of tetrahedron phase spaces $\mathcal{P}_{\partial \Delta} \simeq \mathbb{C}^* \times \mathbb{C}^*$ by \mathbb{C}^* actions. (This statement is independent of \mathbf{t}_{3d} .) We will also see that pulling the product of tetrahedron Lagrangians $\prod_i \mathcal{L}_{\Delta_i}$ through the symplectic quotient produces an open subvariety $\mathcal{L}_2(M, \mathbf{t}, \mathbf{t}_{3d}) \subseteq \mathcal{L}_2(M, \mathbf{t})$. The closure of $\mathcal{L}_2(M, \mathbf{t}, \mathbf{t}_{3d})$ equals $\mathcal{L}_2(M, \mathbf{t})$ (and thus is independent of \mathbf{t}_{3d}) when the 3d triangulation is sufficiently refined. Note that expressing $\mathcal{L}_2(M, \mathbf{t})$ via symplectic reduction makes manifest not only that it is a Lagrangian submanifold but that the K_2 avatar of the symplectic form Ω vanishes on it — *i.e.* it is a K_2 Lagrangian.

A.1. Boundary coordinates

Let us begin to describe the components of (A.1).

A.1.1. Small tori. Since small tori are disjoint from the rest of the topological boundary of M , they can be considered separately. They contribute canonical factors to $\mathcal{P}_2(\partial M, \mathbf{t})$, independent of \mathbf{t} .

Let us choose A- and B-cycles for a small torus T^2 , oriented so that the intersection product $\langle \gamma_B, \gamma_A \rangle = 1$ (see Figure A1). Since $\pi_1(T^2) \simeq \mathbb{Z}^2$ is abelian, the holonomies of any flat $PGL(2)$ connection on T^2 can be simultaneously conjugated into Jordan normal form. The flat connection is




Figure A1: A- and B-cycles with $\langle \gamma_B, \gamma_A \rangle = 1$, if viewing the small torus boundary from *inside* M .

then uniquely determined by the two A- and B-cycle eigenvalues λ_A, λ_B , modulo any residual action of the Weyl group, which would send $\lambda \rightarrow \lambda^{-1}$. A choice of framing flag on T^2 is simply a choice of simultaneous eigenline for the two holonomies, and thus *orders* the eigenvalues, breaking the Weyl symmetry.

Suppose that upon parallel transport around γ_A and γ_B the 1-dimensional subspace in the framing flag is multiplied by $\pm\lambda_A^{-1}$ and $\pm\lambda_B^{-1}$, respectively. Another way to say this would be that any projective basis (f_1, f_2) of the two-dimensional space $V_2 = \mathbb{C}^2$ that's compatible with the flag F (*i.e.* $f_1 \in F$) is multiplied by the holonomy matrix

$$(A.5) \quad \begin{pmatrix} f_1 \\ f_2 \end{pmatrix} \mapsto \begin{pmatrix} 1 & 0 \\ * & \lambda^2 \end{pmatrix} \begin{pmatrix} f_1 \\ f_2 \end{pmatrix}$$

after being parallel-transported around a respective cycle γ . The matrix here is unnormalized, in $PGL(2)$.

The \mathbb{C}^* coordinates for framed flat $PGL(2)$ on the small torus are simply λ_A^2, λ_B^2 . The symplectic form and Poisson bracket are

$$(A.6) \quad \Omega = \frac{1}{2} \frac{d\lambda_B^2}{\lambda_B^2} \wedge \frac{d\lambda_A^2}{\lambda_A^2}, \quad \{\log \lambda_B^2, \log \lambda_A^2\} = 2.$$

The factors of 2 and 1/2 here ultimately come from the Cartan matrix of sl_2 and its inverse, *cf.* [6]. Note that only the squares of eigenvalues are well-defined in $PGL(2)$. We can lift “partway” to $SL(2)$ by taking a square root of one eigenvalue or the other, as discussed physically in Section 2.1.2. Then the Poisson bracket becomes more canonical, *e.g.*, $\{\log \lambda_B, \log \lambda_A\} = 1$.

A.1.2. Big boundary. Suppose that a component of the big boundary of M has genus g and h holes. Then every 2d ideal triangulation (like the WKB triangulations of Section 2) has exactly $2d + h = 3(g - 1 + h)$ edges. Each edge E is assigned a coordinate x_E . Namely, the sections of the flag




Figure A2: Taking a cross-ratio to define x_E .

bundle at the four vertices of the quadrilateral containing E are parallel-transported to a common point inside the quadrilateral and evaluated there to obtain four flags A, B, C, D in the two-dimensional space $V_2 \simeq \mathbb{C}^2$. (These are just lines in \mathbb{C}^2 .) The orientation convention is as in Figure A2. Then one takes a cross-ratio

$$(A.7) \quad x_E = \frac{\langle a \wedge b \rangle \langle c \wedge d \rangle}{\langle a \wedge c \rangle \langle b \wedge d \rangle},$$

where $\langle * \wedge * \rangle$ denotes an $SL(2)$ -invariant volume form on V_2 , and a, b, c, d are any vectors in the lines A, B, C, D . Note that the boundary in Figure 12 is viewed from *outside* the manifold M , and that the orientation of M determines the sign in (A.7). This sign is opposite the “positive” convention of Fock and Goncharov [16] but is more natural for relations to 3d tetrahedra.³⁸

Now consider a hole on the big boundary and a flat section F of the flag bundle there. Suppose that after a *counterclockwise* monodromy around the hole, as viewed from outside of M , the 1-dimensional subspace in the flag F is rescaled by a factor $\pm\lambda^{-1}$, as in (A.5). Then it is not hard to see that the product of edge-coordinates around the hole equals λ^2 ,

$$(A.8) \quad \prod_{E \text{ around hole}} (-x_E) = \lambda^2.$$

³⁸Asking for full 3d symmetry of the coordinates on tetrahedra — for example, that coordinates are invariant with respect to \mathbb{Z}_3 cyclic rotations — seems to necessarily break positivity. It can be restored by choosing a global “time” direction in a 3-manifold M , so that a pair of opposite edges of each tetrahedron is distinguished. This choice is automatically present when just studying framed flat connections a big boundary and building 3-manifolds from sequences of flips. It would be interesting to study the positive structure of the coordinates we discuss here more generally.




Figure A3: Central elements on the big boundary, coming from holonomies around holes. Logarithmically, $\sum_i (X_i - i\pi) = 2\Lambda$ while $\sum_i (Y_i - i\pi) = -2\Lambda$. Around a small disc, with unipotent holonomy, the relation is $\prod_i (-z_i) = 1$, which lifts to $\sum_i (Z_i - i\pi) = -2\pi i$.

as in Figure A3. For future reference, it is convenient to lift this to a logarithmic relation

$$(A.9) \quad \sum_{E \text{ around hole}} (X_E - i\pi) = 2\Lambda.$$

Aside from the h relations (A.8), the edge-coordinates x_E are completely independent. This leaves $2d$ coordinates to parameterize the space of framed flat connections on the big boundary. The symplectic and Poisson structures are given by

$$(A.10) \quad \Omega = \sum_{E, E'} \epsilon_{E, E'}^{-1} \frac{dx_E}{x_E} \wedge \frac{dx_{E'}}{x_{E'}}, \quad \{\log x_E, \log x_{E'}\} = \epsilon_{E, E'},$$

where $\epsilon_{E, E'} \in \{0, \pm 1, \pm 2\}$ counts the number of oriented triangles shared by edges E and E' . The contribution is $+1$ if E' occurs to the left of E in a triangle, as viewed from outside of M . Note that all the hole constraints (A.8) are central with respect to the Poisson brackets.

In order to calculate the $PGL(2)$ holonomy around *any* path on the big boundary, one can use 2d snakes.³⁹ The idea is that in each triangle one can parallel-transport the three framing flags from the vertices to a common point in the interior, evaluate them there, and use them to define

³⁹2d snakes were introduced by [16] to study $PGL(K)$ connections. For $PGL(2)$, however, such holonomy rules go way back to work of Thurston, Fock, and others on 2d hyperbolic geometry.




Figure A4: Projective bases associated to snakes.

six projective bases for the space V_2 . Graphically, each basis is associated to a *snake*: an arrow going from one vertex of the triangle to another, with a “fin” pointing towards the interior of the triangle. (There are six ways to draw a snake in a given triangle.) The projective basis (v_1, v_2) associated to a snake is given by any vector v_1 from the line in the flag at the tail of the snake and the unique vector v_2 from the line at the head of the snake such that either $v_2 + v_1$ or $v_2 - v_1$ (according to orientation, as in Figure A4) lies in the line at the third vertex of the triangle. Then, from this definition, one discovers rules I and II of Figure A5 for transforming the projective basis from one snake to another. In addition, if a snake is moved from one triangle to an adjacent one (rule III), the vector v_2 at the head of the snake gets rescaled by (minus) the cross-ratio coordinate x_E at the edge that the snake sits on.




Figure A5: Transformation matrices associated to fundamental snake moves: I and II within a triangle and III between neighboring triangles. The convention is such that an (unnormalized) $\text{PGL}(2)$ transformation matrix M acts on a projective basis $\begin{pmatrix} v_1 \\ v_2 \end{pmatrix}$ by left multiplication.

By moving snakes around a closed loop on the big boundary, and multiplying the transformation matrices that relate consecutive projective bases along the way, any holonomy can be calculated. Indeed, by simply starting

with a set of abstract edge-coordinates x_E and using the snake rules, one can completely reconstruct both the holonomy representation of a certain flat $PGL(2)$ connection and its choice of framing flags. This constitutes the proof that edge-coordinates really do parametrize an open subset of framed flat connections on the big boundary.

A.1.3. Small annuli. Finally, we come to the small annular parts of ∂M . We already know how to define one coordinate for framed flat $PGL(2)$ connections on a small annulus. We choose an orientation for the non-contractible A-cycle γ_A (as in Figure A3 above) and take the eigenvalue-squared λ^2 for the holonomy along that cycle (as in (A.5)). In $SL(2)$, we would use λ itself. We have just explained that this coordinate Poisson-commutes with all big-boundary coordinates x_E .

What we need is a second coordinate τ that is canonically conjugate to λ and also commutes with (almost) all big-boundary coordinates. It will be defined next. The pair (λ, τ) generalize classical Fenchel-Nielsen length and twist coordinates, respectively.

Notice that the total number of independent big-boundary coordinates x_E , small torus coordinates (λ_A, λ_B) , and small-annulus coordinates (λ, τ) exactly equals the expected dimension of $\mathcal{P}_2(\partial M)$. For example, suppose that the topological boundary ∂M has a single connected component, of genus \tilde{g} , with no holes. If this component is split into two big boundaries $\mathcal{C}_1, \mathcal{C}_2$ of genus g_1, g_2 and h_1, h_2 holes, such that a small annuli connect all the holes pairwise, then the total genus is

$$(A.11) \quad \tilde{g} = g_1 + g_2 + a - 1 = g_1 + g_2 + \frac{1}{2}(h_1 + h_2) - 1,$$

which lets us rewrite the expected complex dimension of $\mathcal{P}_2(\partial M)$ as

$$(A.12) \quad 6\tilde{g} - 6 = (6g_1 - 6 + 2h_1) + (6g_2 - 6 + 2h_2) + 2a.$$

On the RHS, we see the sum of big-boundary dimensions $d_1 + d_2$, and the small annulus contribution $2a$.

A.2. Twists for small annuli

In order to define the twist coordinate τ for a small annulus, let us first assume that the big-boundary triangulation \mathbf{t} is such that the annulus ends in two degenerate triangles t_1 and t_2 , as shown in Figure A6. We call the degenerate edges of these triangles E_1, E_2 ; and we call the second vertex of each triangle (the one disjoint from the small annulus) v_1 and v_2 . We will use




Figure A6: The setup for defining a twist coordinate.

the framing flags at v_1 and v_2 to normalize the framing flag on the annulus in two different ways, and take a ratio of these normalizations to define τ .⁴⁰

We must choose an oriented open path γ_τ along the annulus that has intersection number one with the A-cycle γ_A . We require that the endpoints of γ_τ lie anywhere along the S^1 boundaries of the annulus *except* at the points where the degenerate edges E_1, E_2 attach. We take two paths $\gamma_\tau, \gamma_\tau'$ to be equivalent if there exists a homotopy relating them *without* crossing the endpoints of E_1 and E_2 . Thus, any two inequivalent paths are related by extra twists around the A-cycle of the annulus.




Figure A7: The parallel transport of vectors in the framing flags to a common point p on the annulus. The twist can be defined by either (a) comparing normalizations given by these framing flags; or b) comparing projective snake-bases at the ends of the annulus.

Now, let us use the flat connection on ∂M to parallel-transport the framing flags F_1 and F_2 from vertices v_1 and v_2 to the endpoints of the

⁴⁰The complex generalization of Fenchel-Nielsen coordinates that we describe here are closely related to coordinates studied by Kabaya [56], and previously by [64–66] and others. The novel feature here is use of framed moduli spaces and their eventual quantization. A discussion of length/twist coordinates in the context of framed moduli spaces also appeared in the first arXiv version of [16].

curve γ_τ and then to some common point p around the middle of γ_τ . In fact, inside each triangle t_i there are two inequivalent ways to transport the flags from vertex v_i to the endpoint of γ_τ on the central hole, as shown in Figure A7a. Therefore, we obtain four different lines in the fiber $V_2 \simeq \mathbb{C}^2$ at p , denoted $F_1|_p, F'_1|_p, F_2|_p, F'_2|_p$. Notice that if we let M_λ denote the $PGL(2)$ holonomy matrix around the A-cycle starting and ending at p , then

$$(A.13) \quad F'_1|_p = M_\lambda F_1|_p, \quad F'_2|_p = M_\lambda F_2|_p.$$

Of course, in addition to the flags we have parallel-transported to the annulus, we also have the framing flag A on the annulus itself. Let a be any vector in the one-dimensional subspace of $A|_p$. Let f_1, f'_1, f_2, f'_2 be any vectors in the lines $F_1|_p$, etc. Then

$$(A.14) \quad a_1 = \sqrt{\frac{\langle f_1 \wedge f'_1 \rangle}{\langle f_1 \wedge a \rangle \langle f'_1 \wedge a \rangle}} a \quad \text{and} \quad a_2 = \sqrt{\frac{\langle f_2 \wedge f'_2 \rangle}{\langle f_2 \wedge a \rangle \langle f'_2 \wedge a \rangle}} a$$

provide two different invariant normalizations of the vector a . The square of the ratio of normalizations,

$$(A.15) \quad \tau = -\frac{\langle f_1 \wedge f'_1 \rangle}{\langle f_1 \wedge a \rangle \langle f'_1 \wedge a \rangle} \cdot \frac{\langle f_2 \wedge a \rangle \langle f'_2 \wedge a \rangle}{\langle f_2 \wedge f'_2 \rangle} = -\frac{\langle f_1 \wedge M_\lambda f_1 \rangle \langle f_2 \wedge a \rangle^2}{\langle f_1 \wedge a \rangle^2 \langle f_2 \wedge M_\lambda f_2 \rangle}$$

is invariant under rescaling a, f_1, f'_1 , etc., the $SL(2)$ action on V_2 , and the chosen volume form $\langle * \wedge * \rangle$; it defines the twist coordinate τ . (To obtain the expression on the RHS of (A.15) we used $\langle Mf \wedge a \rangle = \langle f \wedge M^{-1}a \rangle$ for any $M \in SL(2)$, and $M_\lambda^{\pm 1}a = \lambda^{\mp 1}a$.)

Alternatively, we may define τ by comparing two projective bases at the ends of the annulus. Let us form two bases associated to the snakes in Figure A7b. We still evaluate all flags at the common point p in the middle of the annulus. We can use the same vector a as the first vector in both bases, which then uniquely selects the second components w_1 and w_2 obeying

$$(A.16) \quad w_i \in F_i|_p, \quad a + w_1 \in F'_1|_p, \quad a + w_2 \in F'_2|_p.$$

Then we have

$$(A.17) \quad w_2 = -\tau w_1 + \alpha a$$

for some α ; in other words,

$$(A.18) \quad \tau = -\frac{\langle a \wedge w_2 \rangle}{\langle a \wedge w_1 \rangle}.$$

Note again that this definition of τ is independent of the choice of vector a_1 , or the precise choice of point p . To see that the definitions (A.15) and (A.18) are equivalent, it suffices to take $f_i = w_i$, $f'_i = a + w_i$ in (A.15) and simplify the RHS.

The definition of τ does depend on the equivalence class of the path γ_τ . Indeed, if γ_τ is given an extra turn around the A-cycle of the annulus, the twist coordinate transforms as

$$(A.19) \quad \tau \mapsto \lambda^2 \tau.$$

From the techniques for calculating Poisson brackets developed in [16] and [18, 21], it is straightforward to show that

$$(A.20) \quad \{\log \tau, \log \lambda^2\} = 2.$$

This is basically a consequence of the fact that both the A-cycle eigenvalue and the twist can be expressed by comparing projective snake-bases at the beginning and end of two paths γ_A and γ_τ on the annulus, and these paths have intersection number one. To get a canonical bracket, we should lift to $SL(2)$ and use λ rather than λ^2 , or take a square root of τ . Since the bracket (A.20) is nontrivial, whereas λ commutes with all edge coordinates x_E on the big boundary, it must be that τ is a new coordinate, completely independent of the x_E . Thus it fills out the required dimension of $\mathcal{P}_2(\partial M)$ as desired. The twist τ commutes with all big-boundary coordinates *except* those on the central edges of the degenerate triangles $x_{E_1} = x_{E_2}^{-1} = -\lambda^2$ and the circular outside edges of the degenerate triangles. The latter implies that the Poisson bracket on the phase space (A.1) is not completely independent between the big-boundary and small-annulus factors. However, the twist τ can be adjusted by a Laurent monomial of x_E coordinates (on a case-by-case basis) to obtain complete factorization, if desired.

Having defined all the coordinates on $\mathcal{P}_2(\partial M, \mathfrak{t})$, we see now exactly which patch of $\mathcal{P}_2(\partial M)$ it describes. It is the open subset of $\mathcal{P}_2(\partial M)$ for which all cross-ratios in the definitions of edge coordinates x_E and twists τ are well-defined in \mathbb{C}^* . This means that the configurations of framing flags entering the cross-ratios are generic. It implies in particular that in the presence of small annuli the holonomy representation of the flat $PGL(2)$ connection must not be reducible to a proper reductive subgroup (*i.e.* to $GL(1)$).

A.2.1. General boundary triangulation and canonical twists. In general, a chosen big-boundary triangulation \mathbf{t} may not have degenerate triangles at the ends of all small annuli. If so, it is still possible to define *some* twist coordinates τ the same way as above. Namely, we arbitrarily choose two triangles t_1, t_2 in the triangulation at the ends of each annulus; choose paths $\tilde{\gamma}_\tau$ that begin and end inside these triangles; and then compare normalizations of flags along the path. The resulting coordinates $\tilde{\tau}$ will always satisfy the Poisson bracket (A.20), but may in general have very complicated Poisson brackets with the various x_E on the big boundary. If desired, one can then proceed to modify τ by Laurent monomials in big-boundary coordinates (or even their roots) in order to simplify or factorize the Poisson bracket.

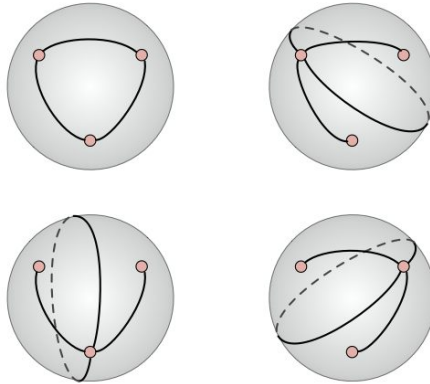


Figure A8: The four triangulations of a 3-punctured sphere.

When all big boundaries are three-punctured spheres, connected by small annuli — such as in a “pants degeneration” of some surface \mathcal{C} — much more canonical twist coordinates can be defined. Indeed, three-punctured spheres only admit four distinct triangulations (Figure A8). Correspondingly, an annulus may attach to a hole incident to one, two, or four edges of the big-boundary triangulation.

If only one (degenerate) edge touches the end of an annulus, the choice of endpoint for γ_τ there is canonical. If two edges touch the end of an annulus, one can take an average over two different paths $\gamma_\tau, \gamma'_\tau$ ending on the triangles separated by these edges, and related (locally) by a twist around the A-cycle. If four edges touch, then one can take an average over four paths, differing from each other by quarter-twists, as illustrated in Figure 38 on page 219. Globally, one must also choose an overall (possibly fractional) twisting in the bulk of each annulus.

The resulting coordinates $\{\lambda_i, \tau_i\}_{i=1}^a$ have the pleasant property that they diagonalize the Poisson bracket:

$$(A.21) \quad \{\log \tau_i, \log \lambda_j\} = \delta_{ij}.$$


Also notice that since three-punctured spheres carry no independent degrees of freedom, the entire phase space is parametrized by these pairs: $\mathcal{P}_2(\partial M, \mathbf{t}) \simeq (\mathbb{C}^*)^{2a}$. These are the most direct generalizations of Fenchel-Nielsen lengths and twists.

A.3. Tetrahedra and coordinates in the bulk

Having defined coordinates for framed flat connections on the boundary of M , we proceed to extend the flat connections into the bulk. We explain how to relate boundary coordinates to edge-parameters of tetrahedra, and how to obtain $\mathcal{L}_2(M, \mathbf{t})$ via symplectic gluing. This is largely review (covering aspects of [6, 34–36]), aside from the computation of the twist τ .

A.3.1. A single tetrahedron. For a single tetrahedron, the phase space $\mathcal{P}_2(\partial\Delta)$ is the space of framed flat connections with unipotent holonomy on a 4-punctured sphere; while the subspace $\mathcal{L}_2(\Delta)$ is simply the set of configurations of four flags at the punctures (since flat connections become trivial in the bulk).

A tetrahedron comes with a conical boundary triangulation \mathbf{t} . Then $\mathcal{P}_{\partial\Delta} := \mathcal{P}_2(\partial\Delta, \mathbf{t})$ describes the framed flat connections whose framing flags are generic — giving well-defined cross-ratios at the six edges. From Section A.1.2 we know that the product of edge-coordinates around any vertex must equal -1 (since the holonomy is unipotent). It’s then standard to call the six edge-coordinates z, z', z'' , equal on opposite edges, occurring in the same counter-clockwise order around any vertex, and satisfying $zz'z'' = -1$. We find



$$(A.22) \quad \mathcal{P}_{\partial\Delta} = \{z, z', z'' \in \mathbb{C}^* \mid zz'z'' = -1\} \simeq (\mathbb{C}^*)^2,$$

with $\{\log z, \log z'\} = \{\log z', \log z''\} = \{\log z'', \log z\} = 1$, or $\Omega_{\partial\Delta} = \frac{dz}{z} \wedge \frac{dz'}{z'}$. For future reference, we also define a logarithmic lift of phase space coordinates that satisfy

$$(A.23) \quad Z + Z' + Z'' = i\pi, \quad \{Z, Z'\} = \{Z', Z''\} = \{Z'', Z\} = 1.$$

One can compute any holonomy on the boundary of the tetrahedron by moving snakes around closed loops and using the rules of Figure A5. For example, the holonomy around any vertex is conjugate to

$$(A.24) \quad M_{\text{vertex}} = \begin{pmatrix} 1 & 0 \\ -zz'(z'' + z^{-1} - 1) & -zz'z'' \end{pmatrix}.$$

We see explicitly that $zz'z'' = -1$ is the unipotent constraint. If we want to extend a flat connection into the bulk of the tetrahedron, all holonomies must be trivial, so we also find that the Lagrangian must be given by

$$(A.25) \quad \mathcal{L}_\Delta := \mathcal{L}_2(\Delta, \mathbf{t}) := \{z'' + z^{-1} - 1 = 0\} \subset \mathcal{P}_{\partial\Delta}.$$

The Lagrangian condition is trivially satisfied, and moreover this is the canonical K_2 -Lagrangian for the form $z \wedge z'$.

Alternatively, we can obtain the Lagrangian constraint (A.25) by computing all three cross-ratio coordinates z, z', z'' from flags that have been parallel-transported to a single common point in the bulk of the tetrahedron. The three cross-ratios obey a Plücker relation, which is precisely $z'' + z^{-1} - 1 = 0$. Thus \mathcal{L}_Δ describes the configuration space of four generic flags in $V_2 \simeq \mathbb{C}^2$.

A.3.2. Triangulated 3-manifold. Now suppose that $M = \cup_{i=1}^N \Delta_i$ has a 3d triangulation \mathbf{t}_{3d} compatible with the boundary triangulation \mathbf{t} , and suppose that we fix a framed flat connection on M . We can define edge-parameters z_i, z'_i, z''_i for every tetrahedron Δ_i by parallel-transporting the framing flags at the vertices of Δ_i to a common point in its interior, and taking cross-ratios. (For each tetrahedron, we must choose which edges to label z_i, z'_i, z''_i , consistent with the cyclic order induced by the orientation of M ; but since all constructions are invariant under cyclic permutations, the choice is immaterial.) The parameters will automatically satisfy

$$(A.26) \quad z_i z'_i z''_i = -1, \quad z''_i + z_i^{-1} - 1 = 0 \quad \forall i.$$

In addition, for every internal edge E_I of the triangulation, the product of edge-parameters incident to E_I automatically equals 1:

$$(A.27) \quad c_I := \prod (z_i, z'_i, z''_i \text{ around } E_I) = 1.$$

The edge constraints $c_I = 1$ hold algebraically due to a perfect cancellation in the numerators and denominators of cross-ratios computed around the

edge E_I . They reflect the fact that the holonomy on a contractible loop in M that winds around E_I is trivial. We will lift (A.27) to the logarithmic constraints

$$(A.28) \quad C_I := \sum (Z_i, Z'_i, Z''_i \text{ around } E_I) - 2\pi i = 0.$$

The holonomy around any closed loop in M can be calculated with the help of 3d snakes. We already explained how to use snakes to define local projective bases for $V_2 \simeq \mathbb{C}^2$, and how to move them around the boundary of a single tetrahedron. To move them between adjacent tetrahedra inside M , we simply supplement the rules of Figure A5 with one more rule in Figure A.3.2.




Figure A9: Moving a snake between two tetrahedra glued along face \mathcal{A} .

Conversely, if we start with abstract edge-parameters z_i, z'_i, z''_i for every tetrahedron that satisfy (A.26) and (A.27), we can use the four snake rules to fully reconstruct the holonomy representation of a flat $PGL(2)$ connection and its framing flags. Conditions (A.26)–(A.27) are precisely the “cocycle constraints” that ensure flatness locally.

Thus, the set of tetrahedron parameters satisfying (A.26) and (A.27) parameterizes an open algebraic subset $\mathcal{L}_2(\Delta, \mathbf{t}, \mathbf{t}_{3d})$ in the space of framed flat connections on M . This subset contains those flat connections whose framing data is generic with respect to the 3d triangulation \mathbf{t}_{3d} : the parallel-transported flags inside each tetrahedron define non-degenerate cross-ratios $z_i, z'_i, z''_i \in \mathbb{C}^* \setminus \{1\}$. The dependence of $\mathcal{L}_2(\Delta, \mathbf{t}, \mathbf{t}_{3d})$ on the actual 3d triangulation is very mild; for sufficiently refined triangulations, the closures of all $\mathcal{L}_2(\Delta, \mathbf{t}, \mathbf{t}_{3d})$ ’s agree. Therefore, we will usually just write $\mathcal{L}_2(\Delta, \mathbf{t})$.⁴¹ In order to show that $\mathcal{L}_2(\Delta, \mathbf{t})$ is a K_2 -Lagrangian, we will show in Section A.3.4 that it comes from symplectic reduction.

A.3.3. Boundary and bulk coordinates. The boundary coordinates $x_E, (\lambda_A, \lambda_B), (\lambda, \tau)$ for big and small boundaries of a framed 3-manifold can *all* be written as Laurent monomials in the tetrahedron edge-parameters. These monomials provide an explicit embedding of $\mathcal{L}_2(\Delta, \mathbf{t})$ in $\mathcal{P}_2(\partial\Delta, \mathbf{t})$, and are also the key to symplectic gluing.

⁴¹A more careful way to define $\mathcal{L}_2(\Delta, \mathbf{t})$ would be to take the union of all $\mathcal{L}_2(\Delta, \mathbf{t}, \mathbf{t}_{3d})$ over all 3d triangulations \mathbf{t}_{3d} that agree with \mathbf{t} .

First, for an edge E on the big boundary of M , it is not hard to see from the definition of various cross-ratios that

$$(A.29) \quad \begin{aligned} x_E &= \prod (z_i, z'_i, z''_i \text{ incident to } E) \quad \text{or} \\ X_E &= \sum (Z_i, Z'_i, Z''_i \text{ incident to } E). \end{aligned}$$

We can then use 3d snakes to calculate holonomy eigenvalues and twists on the small boundary. In particular, given *any* closed or open path γ on the small boundary, the projective basis defined by a snake at the end of the path (meaning a snake whose tail lies at the endpoint of γ and whose head extends into M along the face of a tetrahedron, *cf.* Figure A10) will be related to the projective basis defined by a snake at the beginning of the path by a transformation matrix of the form

$$(A.30) \quad M_\gamma = \begin{pmatrix} 1 & 0 \\ * & \lambda_\gamma^2 \end{pmatrix}.$$

The matrix must be triangular, because the holonomy along a small boundary must preserve the framing flag there. The matrix element λ_γ^2 is *either* the holonomy eigenvalue or the twist coordinate associated to the path, and turns out to be a product of $z_i^{\pm 1}, z'_i{}^{\pm 1}, z''_i{}^{\pm 1}$.





Figure A10: Snake extending from a path into M .

It is very convenient to rewrite the snake rules for such a coordinate λ_γ more directly in terms of the edge-parameters encountered by the path. Recall that the small boundary is tiled by small triangles from truncated tetrahedra vertices in a 3d triangulation \mathbf{t}_{3d} . The angles of these small triangles are labelled by edge-parameters z_i, z'_i, z''_i . Let γ be an oriented path

drawn on the small boundary so that inside every triangle γ looks like



with either orientation. We assume that γ has no self-intersections and that it is either closed (as relevant for A- and B-cycle coordinates) or begins and ends at S^1 boundaries of the small boundary (as relevant for twists).

Then the parameter λ_γ^2 can be computed by multiplying and dividing by edge-parameters on angles subtended by γ , using the rules of Figure A11. These rules can consistently be lifted to calculate logarithmic parameters $2\Lambda_\gamma$ (with $\lambda_\gamma^2 = \exp(2\Lambda_\gamma)$), as shown. An example appears in Figure A12.




Figure A11: Standard factors associated to segments of an oriented path, and their logarithmic lifts.





Figure A12: Length and twist for a small annulus as combinations of tetrahedron edge-parameters. Here $\lambda^2 = z_3/(z_2 z_4'')$ and $\tau = -z_6 z_5'' (z_6' z_5')^{-1} z_3'' (z_4' z_1')^{-1}$. Or, Logarithmically, $2\Lambda = -Z_4'' + Z_3 - Z_2$, while $\mathcal{T} = Z_6 + Z_5'' - Z_6' - Z_5' + Z_3'' - Z_4' - Z_1' + i\pi - Z_1$. Here the small boundary is viewed from *outside* of M .

Notice that the parameter $2\Lambda_\gamma$ is independent of path homotopies, generated by moves of the type

(A.31)



The first move holds due to the linear logarithmic relation (A.23) among Z_i, Z'_i, Z''_i . The second move adds a sum of edge-parameters C_I around an internal edge of the triangulation, as in (A.28), to $2\Lambda_\gamma$; but C_I should be set to zero according to the gluing constraint. One can understand the logarithmic gluing constraint as a *consistency condition* that allows unambiguous parameters $2\Lambda_\gamma$, independent of path homotopy, to be defined. Also note, in general, that any contractible clockwise (resp., counterclockwise) loop on the small boundary must be assigned logarithmic parameter $2\pi i$ ($-2\pi i$), modulo gluing functions C_I .




Figure A13: The immediate neighborhood of a hole on the big boundary, connected to a small annulus. Calculating the logarithmic path parameter 2Λ shows that $2\Lambda = \sum_i (X_i - i\pi)$.

Invariance under path homotopy can be used to motivate the logarithmic relation (A.9) between the sum of big-boundary coordinates around a hole and the eigenvalue there. Consider the setup in Figure A13. On one hand, the path γ is in the homotopy class of the A-cycle of the annulus, so its total parameter is 2Λ . On the other hand, γ is composed of semi-circular segments that sum up edge-parameters around external edges — contributing external coordinates X_i — and “bounces” that subtract $i\pi$. Therefore, $2\Lambda = \sum (X_E - i\pi)$.

A.3.4. Symplectic reduction. One would like to show that the phase space $\mathcal{P}_2(\partial M, \mathbf{t})$ defined in Sections A.1–A.2 can be obtained as the symplectic reduction of a product of tetrahedron phase spaces in *any* 3d triangulation of M . This requires two basic results.

The first result is that the symplectic structure on boundaries of tetrahedra is compatible with the symplectic structure on ∂M . It can be stated as follows. Fix a triangulation \mathbf{t}_{3d} with N tetrahedra and let $\mathcal{P}_\times = \prod_{i=1}^N \mathcal{P}_{\partial\Delta_i} \simeq \{z_i, z'_i, z''_i \in \mathbb{C}^* | z_i z'_i z''_i = -1\} \simeq (\mathbb{C}^*)^{2N}$ be the product of tetrahedron phase spaces, with the product symplectic and Poisson structures, *i.e.* $\{\log z_i, \log z'_j\} = \{\log z'_i, \log z''_j\} = \{\log z''_i, \log z_j\} = \delta_{ij}$. For every internal edge E_I in the triangulation, we define a function c_I on \mathcal{P}_\times via (A.27). We also introduce a set of “boundary” functions on \mathcal{P}_\times , consisting of 1) functions x_E for every external edge, defined by (A.29); 2) functions λ_A^2, λ_B^2 associated to A- and B-cycles on every small torus boundary, defined by choosing representative paths γ_A, γ_B and using the path rules of Figure A11; and similarly 3) length and twist functions λ^2, τ for every small annulus. These are all Laurent monomials in the coordinates z_i, z'_i, z''_i . (Note that the definitions of $\gamma_A, \gamma_B; \gamma, \tau$ are unique up to multiplication by c_I 's.) Then

Theorem A.1. *The functions $x_E, \lambda_A, \lambda_B, \lambda, \tau$ on the symplectic space \mathcal{P}_\times satisfy the standard Poisson brackets described in Sections A.1–A.2. Namely,*

$$(A.32a) \quad \{\log x_E, \log x'_E\} = \epsilon_{EE'}, \quad \{x_E, \lambda_A\} = \{x_E, \lambda_B\} = \{x_E, \lambda\} = 0;$$

$$(A.32b) \quad \{\log \lambda_B^2, \log \lambda_A^2\} = \{\log \lambda^2, \log \tau\} = 2,$$

with functions λ_A, λ_B or λ, τ on different tori or annuli commuting with each other. Moreover, the internal-edge functions c_I commute among themselves and with all of the boundary functions.

The proof follows from basic combinatorics, and is a direct consequence of arguments in the Appendix of [6]. One way of phrasing this theorem in terms of path parameters (Section A.3.3) is that the commutation relation between any two path parameters $\{\log \lambda_\gamma^2, \log \lambda_{\gamma'}^2\}$ equals the intersection number $2\langle \gamma, \gamma' \rangle$ of paths on the small boundary — as long as at least one of the paths is closed. For example, the c_I are path parameters for closed contractible loops on the small boundary, so their intersection number with any other path is zero, which implies $\{c_I, *\} = 0$.

Now suppose that the big boundary has a phase space of expected dimension $2d = 6g - 6 + 2h$; and that there are a small annuli and t small tori. The

nontrivial brackets (A.32) show immediately that the $2t$ functions (λ_A, λ_B) , the $2a$ functions (λ, τ) and $2d$ of the functions x_E must be linearly independent, and independent of all the c_I 's.⁴² Moreover, it follows from calculating Euler characters that the number of internal edges is

$$(A.33) \quad \# \text{ internal edges } (E_I) = N - d - a .$$

The corresponding $N - d - a$ functions c_I *cannot* all be independent because there are too many of them. Indeed, an argument given in [6, Sec 5] (based on homotopy-invariance of path parameters) shows that there is a linear relation among the c_I 's for every small torus boundary of M . The second result we need is that there are no other relations:

Theorem A.2. *The number of linearly independent gluing functions c_I is exactly $N - d - a - t$.*

It follows immediately from the two Theorems that

$$(A.34) \quad \begin{aligned} \mathcal{P}_2(\partial M, \mathbf{t}) &= \left(\prod_{i=1}^N \mathcal{P}_{\partial \Delta_i} \right) // (\mathbb{C}^*)^{N-d-a-t} \\ &= \left(\prod_{i=1}^N \mathcal{P}_{\partial \Delta_i} \right) / (\mathbb{C}^*)^{N-d-a-t} \Big|_{(c_I = 1)} \\ &\simeq (\mathbb{C}^*)^{2d+2a+2t} . \end{aligned}$$

The $(\mathbb{C}^*)^{N-d-a-t}$ action in the symplectic quotient is generated by any independent subset of the commuting moment maps c_I . There is also an obvious logarithmic version of the symplectic reduction, using logarithmic functions $X_E, \Lambda, \mathcal{T}, Z_i$, etc., and an additive $(\mathbb{C})^{N-d-a-t}$ action generated by the C_I .

Having (A.34), we may attempt to pull the canonical product Lagrangian $\prod_i \mathcal{L}_{\Delta_i} \subset \prod_i \mathcal{P}_{\partial \Delta_i}$ through the reduction. This makes sense as long as the moment maps and their flows are transverse to the Lagrangian equations, and will automatically define a K_2 Lagrangian submanifold $\mathcal{L}_2(M, \mathbf{t}, \mathbf{t}_{3d})$. It must (tautologically) be the same space we described in Section A.3.2, but now the Lagrangian property is manifest.

A better way to phrase this second Theorem above (and incorporate it with the first) is by generalizing a beautiful combinatorial result of Neumann

⁴²We mean linearly independent with respect to multiplication, since everything is \mathbb{C}^* -valued.

[36, Thm 4.2], which in turn generalized a classic result of Neumann and Zagier [35]. This will appear elsewhere, along with a full proof.

A.3.5. Logarithms and quantization. The description of the symplectic pair $\mathcal{L}_2(M, \mathfrak{t}) \subset \mathcal{P}_2(\partial M, \mathfrak{t})$ as a symplectic quotient (A.34) enables it to be quantized using methods of [34].⁴³ By quantization here we mean promoting $\mathcal{P}_2(\partial M, \mathfrak{t})$ to a q -commutative ring (an algebra of operators) $\hat{\mathcal{P}}_2(\partial M, \mathfrak{t})$ generated by $\hat{x}_E, \hat{\lambda}_A, \hat{\lambda}_B, \hat{\lambda}, \hat{\tau}$ with commutation relations

$$(A.35) \quad \hat{x}\hat{y} = q^{\{\log x, \log y\}}\hat{y}\hat{x}$$

(for instance $\hat{\tau}\hat{\lambda}^2 = q\hat{\lambda}^2\hat{\tau}$); and promoting $\mathcal{L}_2(M, \mathfrak{t})$ to a left ideal $\hat{\mathcal{L}}_2(M, \mathfrak{t})$ in this ring. One can also consider various wavefunctions annihilated by the left ideal $\hat{\mathcal{L}}_2(M, \mathfrak{t})$, as discussed in [7, 34, 70–72] and earlier in [73, 74]. We note that the quantization of boundary phase spaces alone, along with some simple Lagrangians, intertwines beautifully with quantization in $PGL(2)$ and higher Teichmüller theory [13, 14, 55].

We emphasize that no matter what method of quantization is used, an effective lift to logarithmic coordinates such as we have described above must arise. It was explained combinatorially in [34] that the seemingly artificial counting of $2\pi i$'s in logarithms translates to keeping track of nontrivial powers of q in quantization. From the perspective of wavefunctions, the imaginary parts of coordinates $\text{Im } Z_i$, etc. are treated as dihedral angles in the geometry of M , and finiteness of wavefunctions depends crucially on the existence of “positive angle structures” [70, 71]. Indeed, it was noticed much earlier (*cf.* [36]) that logarithms are necessary for a consistent definition of “complex volumes” for flat connections on M , which appear as the semi-classical piece of any quantum wavefunction.

Physically, the importance of the $2\pi i$'s is also clear. By the 3d-3d correspondence of [5], operators and wavefunctions become quantized when a 3d gauge theory is placed in a twisted and/or compactified 3d geometry (*cf.* Section 2.2.3 here). This requires a well-defined $U(1)_R$ symmetry. Moreover, in the 3d-3d correspondence, it is precisely the imaginary parts of logarithms Z_i , etc. that are mapped to a physical R-charge assignment. Thus the logarithms tell us how to compactify a 3d gauge theory, and how to quantize.

⁴³It is strongly suspected that such quantization “via symplectic reduction” is completely equivalent to other new quantization methods, such as topological recursion [29, 67, 68], as well quantization via recursion relations for knot polynomials [32, 69]. Showing equivalence of these different approaches should be a fascinating endeavor.

The combinatorial quantization of [34] basically says to quantize tetrahedron Lagrangians and then to pull them through a quantum symplectic reduction. The basic rules are as follows.

First, given a 3d triangulation of M , one forms a product algebra $\hat{\mathcal{P}}_\times$ generated by $\{\hat{z}_i, \hat{z}'_i, \hat{z}''_i\}_{i=1}^N$ with

$$(A.36) \quad \hat{z}_i \hat{z}'_i \hat{z}''_i = -q, \quad \hat{z}_i \hat{z}'_i = q \hat{z}'_i \hat{z}_i, \quad \text{etc.}$$

The central constraint here is most naturally interpreted as an exponentiation of the logarithmic $\hat{Z}_i + \hat{Z}'_i + \hat{Z}''_i = i\pi + \hbar/2$, where $\hat{z}_i = \exp \hat{Z}_i$ (etc.), $q = e^\hbar$, and $[\hat{Z}_i, \hat{Z}'_i] = [\hat{Z}'_i, \hat{Z}''_i] = [\hat{Z}''_i, \hat{Z}_i] = \hbar$. (In general, every logarithmic $i\pi$ gets quantum-corrected to $i\pi + \hbar/2$.) Inside the q -commutative ring $\hat{\mathcal{P}}_\times$, one identifies a canonical left ideal

$$(A.37) \quad \hat{\mathcal{L}}_\times = (\hat{z}''_i + \hat{z}_i - 1)_{i=1}^N.$$

Next, one defines logarithmic gluing operators $\hat{C}_I = \sum (\hat{Z}_i, \hat{Z}'_i, \hat{Z}''_i$ surrounding $E_I) - 2\pi i - \hbar$, and exponentiates them to obtain $\hat{c}_I = \exp \hat{C}_I$. The \hat{c}_I 's are Laurent monomials in the $\hat{z}_i, \hat{z}'_i, \hat{z}''_i$, along with a well-defined q -correction. Similarly, one unambiguously quantizes the logarithmic formulas for $X_E, \Lambda_A, \Lambda_B, \Lambda, \mathcal{T}$ (correcting $i\pi \rightarrow i\pi + \hbar/2$) and exponentiates to obtain expressions for $\hat{x}_E, \hat{\lambda}_A, \hat{\lambda}_B, \hat{\lambda}, \hat{\tau}$.

Finally comes the reduction. One takes the centralizer of the product ideal $\hat{\mathcal{L}}_\times$ with respect to the \hat{c}_I . This means eliminating elements of $\hat{\mathcal{L}}_\times$ that do not commute with the \hat{c}_I . In the centralizer, one may unambiguously set \hat{c}_I , obtaining a left ideal $\hat{\mathcal{L}}_2(M, \mathbf{t}, \mathbf{t}_{3d})$. This ideal can be written entirely in terms of the boundary coordinates $\hat{x}_E, \hat{\lambda}_A, \hat{\lambda}_B, \hat{\lambda}, \hat{\tau}$ (since these generate the centralizer of \hat{c}_I in the ring $\hat{\mathcal{P}}_\times$). Indeed, $\hat{\mathcal{L}}_2(M, \mathbf{t}, \mathbf{t}_{3d}) \subset \hat{\mathcal{P}}_2(M, \mathbf{t})$ is the sought-after quantization of the Lagrangian. By construction, taking the classical limit $q \rightarrow 1$ reduces $\hat{\mathcal{L}}_2(M, \mathbf{t}, \mathbf{t}_{3d})$ to a classical ideal in $\mathcal{P}_2(M, \mathbf{t})$ (viewed as a ring), which must contain the classical Lagrangian $\mathcal{L}_2(M, \mathbf{t}, \mathbf{t}_{3d})$.

Appendix B. Triangulation of general RG and duality manifolds

In this final section, we describe explicitly the triangulation of a framed 3-manifold describing an RG wall or a duality wall for (almost) any 4d $\mathcal{N} = 2$ theory of class $\mathcal{S}[A_1]$ — obtained from compactifying the A_1 (2,0) theory on a punctured Riemann surface. This is actually not too hard to do; there are plenty of 3d triangulations of these manifolds. However, it is useful

to find triangulations that make certain properties of the corresponding gauge theories $T_2[M]$ particularly simple or manifest. Here, following the philosophy from previous sections, we will build the triangulations out of basic RG manifolds M_0 (Section 4.3), which describe free chiral theories with automatic $SU(2)$ flavor enhancements.

Mathematically, the 3-manifolds provide 1) the relation between cross-ratio coordinates and Fenchel-Nielsen-type coordinates, and their quantization, in Teichmüller theory; and 2) the mapping class group action and its quantization in (especially) Fenchel-Nielsen-type coordinates. Our constructions can be viewed as a 3d analog of the work of Teschner [15], in which he described 2d ideal triangulations adapted to pants decompositions of surfaces for this same purpose.




Figure B1: M_0 , repeated from Figure 27

It is actually sufficient to consider RG manifolds only. Suppose, for example, that we want to create a duality manifold $M_{\mathbf{p}\mathbf{p}'}$ that interpolates between two different pants decompositions \mathbf{p}, \mathbf{p}' of a punctured surface \mathcal{C} . We can first create the RG manifolds $M_{\mathbf{p}}$ and $M_{\mathbf{p}'}$ that interpolate between big triangulated boundaries \mathcal{C} and pants decompositions (a usual network of small annuli). Then we flip the big-boundary triangulations of $M_{\mathbf{p}}$ and $M_{\mathbf{p}'}$ so that they match, according to any desired identification (any element of the mapping class group), and we glue $M_{\mathbf{p}}$ and $M_{\mathbf{p}'}$ together to get $M_{\mathbf{p}\mathbf{p}'}$.

B.1. The construction

Let \mathcal{C} be a surface of genus g with h regular or irregular punctures, and let \mathbf{p} be a pants decomposition of \mathcal{C} . We require as usual that $2g - 2 + h + \sum(\text{ranks of irregular punctures}) > 0$, and in addition that $h > 0$. This is our only constraint. It is necessary, since framed 3-manifolds cannot represent RG walls if $h = 0$.⁴⁴ We want to create a triangulated RG 3-manifold $M_{\mathbf{p}}$.

We assume that if \mathcal{C} is a sphere it has at least four punctures. The case of three regular punctures is uninteresting, and $M_{\mathbf{p}}$ for three or fewer irregular punctures can be obtained by small modifications of the basic M_0 geometry.

⁴⁴Framed 3-manifolds *can* represent duality walls for surfaces with no holes ($h = 0$), *i.e.* for theories of class \mathcal{S} with no continuous flavor symmetry in the UV. But these duality manifolds cannot be obtained by gluing RG manifolds, and we will not consider them here.

Step 1: Irregular punctures. First we observe that, given a triangulated RG manifold M' for a surface \mathcal{C}' with regular punctures, it is easy to modify M' to make any of the punctures irregular.

Let us focus on a particular regular puncture that we want to make irregular of minimal rank $1/2$. The big IR boundary of M' has a hole v' on which a small annulus (carrying the regular mass parameter) ends. We choose a triangle t' in the big-boundary triangulation \mathbf{t}' that is adjacent to v' , and we glue a single tetrahedron onto this triangle, as in Figure B2. This introduces a new hole v'_0 on the big boundary, now filled in by a small disc (one of the truncated vertices of the new tetrahedron). We flip the big-boundary triangulation by gluing on additional tetrahedra (two faces at a time) to obtain a triangulation that has a degenerate triangle t_{deg} surrounding v' , with additional vertex v'_0 . Let M be the resulting triangulated 3-manifold.




Figure B2: A degenerate modification of a hole v' on a big boundary. First glue a single tetrahedron onto a triangle adjacent to v' , and then flip to surround v' by a degenerate triangle.

This is enough. We now can “cut” along the circular edge of the degenerate triangle t_{deg} , identifying the big boundary outside this triangle as the IR boundary of M , and t_{deg} itself together with the old UV boundary as the new UV boundary. The circular edge of t_{deg} (connecting the IR and UV boundaries) represents a rank- $1/2$ irregular puncture.

It is clear that we can promote any regular punctures of an RG manifold to rank- $1/2$ irregular ones in this way. To increase the rank further, we simply continue to glue more tetrahedra around the puncture on the IR boundary, adding new holes filled by small discs, and then redefine the UV/IR splitting so ∂M is “cut” by polygons that pass through the new small discs.

Note that irregular punctures often arise naturally in practice (as in the basic RG manifold M_0 itself), or can be obtained by *ungluing* some tetrahedron faces, rather than adding new tetrahedra. Otherwise, the above method can be applied systematically.

For future reference, we will refer to the procedure of Figure B2 — adding a tetrahedron next to a puncture v' and flipping to surround v' by a degenerate triangle — as a *degenerate modification of v'* .

Step 2: Choosing $SU(2)$ enhancements. We may now assume that all punctures are regular.

Let us draw a trivalent skeleton of \mathcal{C} according to the pants decomposition \mathbf{p} (Figure B3), and choose a maximal set of internal edges to promote to non-abelian flavor symmetry $SU(2)$. Recall from Section 5.3 what this means. Each trivalent vertex of this graph will become a 3-punctured sphere in the UV boundary of $M_{\mathbf{p}}$. The edges of the graph become small annuli connecting the spheres. We want to choose un-democratic triangulations for each 3-punctured sphere ($\mathbf{t}_1, \mathbf{t}_2$, or \mathbf{t}_3 of Figure 31) so that as many small annuli as possible begin and end at holes surrounded by degenerate triangles. In the 3d theory $T_2[M, \dots]$ these annuli will correspond to enhanced symmetries $SU(2)_{\lambda_i}$ — which in turn get gauged in coupling to a 4d theory of class \mathcal{S} .




Figure B3: Drawing the skeleton of \mathcal{C} according to a fixed pants decomposition, then choosing a locally maximal set of internal edges to enhance (highlighted in red). The enhanced edges form open or closed paths γ_i .

Combinatorially, we choose a subset of internal edges of the skeleton of \mathcal{C} to “enhance” so that at most two enhanced edges end at any vertex. Some examples appear in Figure B3. We want this subset to be (locally) maximal, meaning that no additional internal edges can be added to it.

We will not worry about external edges of the skeleton — corresponding to small annuli in M that connect the UV and IR boundaries, and whose flavor symmetries in $T_2[M]$ are identified with bulk flavor symmetries of $T_2[\mathcal{C}]$. If these small annuli happen to end in degenerate triangles of 3-punctured spheres on the UV boundary, the triangulation of the IR boundary can always be flipped to create degenerate triangles (and hence $SU(2)$ enhancement) there as well.

Notice that the enhanced edges of the skeleton form either 1) closed loops; or 2) straight open paths. In particular, enhanced edges cannot branch. Moreover, the complement of the enhanced edge paths γ_i consists of:




Figure B4: The three options for unenhanced edges: extremal, isolated, or trinion.

- a) external edges of the skeleton (connecting to punctures of \mathcal{C});
- b) isolated internal edges that connect two enhanced paths; or
- c) “trinions” that have one, two, or three legs connecting to enhanced paths.

See Figure B4. We may (and will) require that there are no un-enhanced trinions (c) with just one leg connecting to an enhanced path (and two legs connecting to punctures). This can be arranged by slightly modifying the set of enhanced edges.

Step 3: Surround enhanced edges with M_0 . For every closed loop γ containing n enhanced edges in the skeleton of \mathcal{C} , build a 3-manifold M_γ by gluing n copies of M_0 end-to-end. An example is shown in Figure B5. Recall that M_0 is the basic RG manifold from Section 4.3, built from two tetrahedra. Each time two copies of M_0 are connected, the circular edges on their ends are identified, trapping a 3-punctured sphere on the inside. These newly identified edges are “thin” in the sense of Section 6.1: they don’t quite fully separate the UV and IR boundaries of M_γ ; but they can be thickened out with extra flips at the end.

For every open path γ of n enhanced edges in the skeleton, we can first build a manifold M'_γ in a similar way, by gluing n copies of M_0 end-to-end (but not closing the loop). Now we need to decorate the ends of this manifold the same way was done for the $N_f = 4$ RG manifold in Section 7.2. We add a self-glued tetrahedron (half of M_0) to each end of M'_γ , trapping 3-punctured spheres; then we fold these tetrahedra back over the big IR boundary of M'_γ , identifying a face of each with a face on the big IR boundary. See Figure B5. Call the resulting manifold M_γ .

Note that the manifolds M_γ associated to open or closed paths γ are themselves RG manifolds corresponding to some surfaces C_γ . Here C_γ is an n -punctured torus if γ is closed, and an $(n + 3)$ -punctured sphere if γ is open.

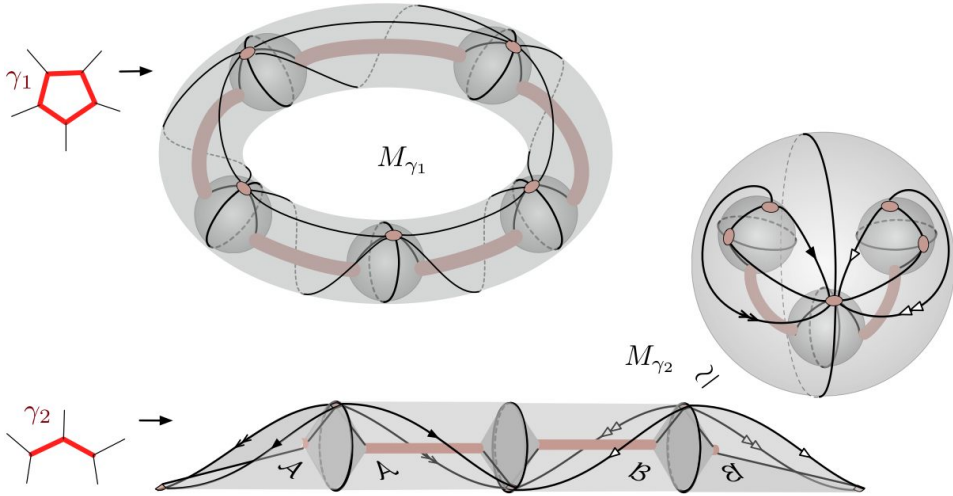


Figure B5: Constructing RG manifolds for the skeletons of a five-punctured torus (top) and a 5-punctured sphere (bottom), with $SU(2)$ enhancement along the internal paths indicated.

Step 4: Connect the enhanced RG manifolds. Finally we glue together our manifolds M_γ into an RG manifold $M_{\mathcal{P}}$.

Due to the remarks at the end of Step 2, any two paths γ, γ' in the skeleton that are to be connected must be connected by either (b) a single un-enhanced edge, or (c) an internal un-enhanced trinion. We will deal with the trinions first. Note that γ and γ' need not be distinct.

Select a trinion, and two of its legs that connect to enhanced paths. (If a third leg also connects to an enhanced path, we leave it alone for now.) Let v, v' be the corresponding holes on the IR boundaries of M_γ and $M_{\gamma'}$ that are to be connected. M_γ may coincide with $M_{\gamma'}$. Perform a degenerate modification of each hole v, v' as in Figure B6, introducing two new degenerate triangles $t_{\text{deg}}, t'_{\text{deg}}$ on the IR boundaries, along with two new holes v_0, v''_0 filled by small discs. Then identify the circular edge of t_{deg} with the circular edge of t'_{deg} . This traps a 3-punctured sphere inside a new manifold $M_\gamma \cup M_{\gamma'}$; two of its holes are identified with v, v' , and the third hole, an amalgamation of v_0 and v''_0 , connects to the new IR boundary. We have reproduced the trinion.

We repeat this procedure for all trinions, attaching two of their three legs. If in the skeleton of \mathcal{C} there were no single un-enhanced edges (b) or trinions (c) with all three legs on enhanced paths, we are done. Otherwise, we arrive a new collection of manifolds \widetilde{M}_i , which, in terms of the skeleton, must




Figure B6: Gluing pieces of RG manifolds together to produce a trinion in the skeleton.

be connected to each other by single un-enhanced internal edges. (These might be either isolated un-enhanced edges, or the third legs of trinions that we neglected to connect above.) Since \mathcal{C} has $h > 0$ punctures, at least one of the \widetilde{M}_i has a hole v_0 that doesn't get connected to anything. Call this manifold $\widetilde{M}^{(0)}$.

We now choose pairs of holes v, v' on the \widetilde{M}_i and glue them together. Sometimes this is obvious and easy to do. Otherwise, we proceed systematically as follows. We require that at least one of v, v' (say v) lies on $\widetilde{M}^{(0)}$, or, after subsequent gluing, on the connected 3-manifold that contains $\widetilde{M}^{(0)}$. To glue, we flip the big-boundary triangulation of $\widetilde{M}^{(0)}$ to create a degenerate triangle t_{deg} surrounding v (with additional vertex at v_0). We also perform a degenerate modification of v' (Figure B2), creating a degenerate triangle t'_{deg} surrounding v' . Then we identify the faces of t_{deg} and t'_{deg} , as in Figure B7. The connection is made, and we repeat.

Step 5: Touching up. We arrive at a triangulated framed 3-manifold $M_{\mathbf{p}}$ representing the RG wall. It may have two technical problems.

The separation between the UV and IR boundaries may be “thin” in places, where only the edges of tetrahedra were identified to create $M_{\mathbf{p}}$. We fix this as usual either by flipping the IR boundary or by flipping 3-punctured spheres on the UV boundary.

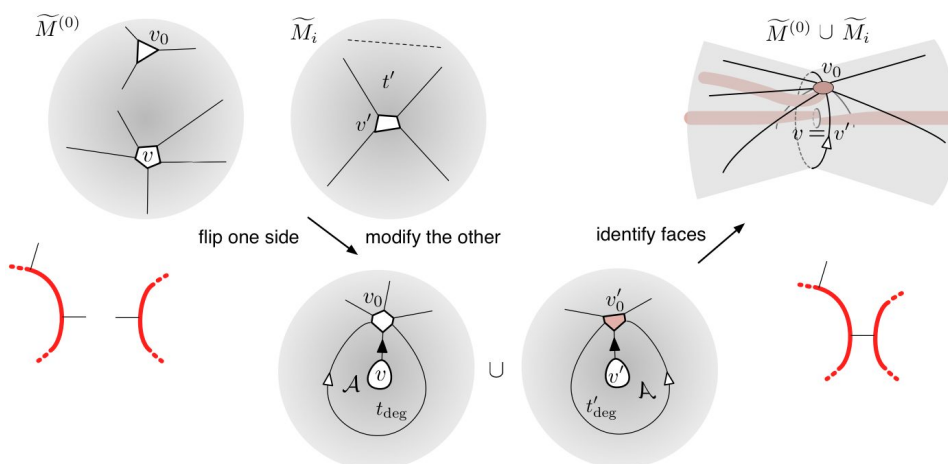


Figure B7: Gluing pieces of RG manifolds together to produce a direct connection in the skeleton. This must be done in the presence of an external puncture.

The gluing of holes corresponding to single un-enhanced edges of the skeleton of \mathcal{C} may also create “hard” internal edges in the 3d triangulation of $M_{\mathbf{p}}$, in the sense of [5, Sec. 4.1]. These are not a problem for mathematical applications or computations of partition functions, line operators, etc. However, in order to define the full 3d gauge theory $T_2[M_{\mathbf{p}}]$, the “hard” edges must be removed by changing the 3d triangulation, *i.e.* by performing a sequence of 2–3 moves.

Appendix C. Quantum dilogarithms

We briefly recall the definition of the quantum dilogarithm function $\Phi_{\hbar}(Z)$. It is used to build $SL(K, \mathbb{R})$ -like Chern-Simons wavefunctions for framed 3-manifolds M , and ellipsoid partition functions for theories $T_K[M]$. (See [7] for the related quantum dilogarithm used to build 3d indices.)

For $Z \in \mathbb{C}$ and $\text{Re } \hbar < 0$ we set

$$(C.1) \quad \Phi_{\hbar}(Z) = \prod_{r \geq 0} \frac{1 - q^{r + \frac{1}{2}} e^Z}{1 - \tilde{q}^{-r - \frac{1}{2}} e^{\tilde{Z}}}, \quad q = e^{\hbar}, \quad \tilde{q} = e^{-\frac{4\pi^2}{\hbar}}, \quad \tilde{Z} = \frac{2\pi i}{\hbar} Z.$$

This is basically the double-gamma function of Barnes [75], rediscovered by Faddeev when studying nonperturbative completions of Weyl algebras

[76]. It has many remarkable properties, most of which we will not discuss here (see *e.g.* [55, 74] for reviews). We observe, though, that $\Phi_{\hbar}(Z)$ can be extended to a meromorphic function of $Z \in \mathbb{C}$ for all \hbar in the cut plane $\mathbb{C} \setminus \{i\mathbb{R}_{\leq 0}\}$, and in particular for $\hbar \in i\mathbb{R}_{>0}$. The poles and zeroes of (C.1) can easily be found from the product (C.1).

The function $\Phi_{\hbar}(i\pi + \frac{\hbar}{2} - Z)$ is the $SL(2)$ Chern-Simons wavefunction of a tetrahedron, in a certain natural polarization Π_z [34]. It is also the ellipsoid (S_b^3) partition function of the tetrahedron theory $T_2[\Delta, \mathbf{t}, \Pi_Z]$, the theory of a free 3d chiral multiplet ϕ with $U(1)_R$ charge zero, charge +1 under a $U(1)_F$ flavor symmetry, and a (supersymmetrized) background Chern-Simons coupling at level -1/2 for the combination $A_F - A_R$ of $U(1)_F$ and $U(1)_R$ gauge fields. To interpret $\Phi_{\hbar}(i\pi + \frac{\hbar}{2} - Z)$ as an S_b^3 partition function, one identifies $\hbar = 2\pi i b^2$ ($b \in \mathbb{R}$) and sets the effective complexified mass of ϕ on S_b^3 to be $m_Z = 2\pi b Z$.

One finds several other forms of $\Phi_{\hbar}(Z)$ in the math and physics literature. In particular, one encounters a “double sine” function s_b in the study of Liouville theory and representation theory of $U_q(sl_2)$, as well as S_b^3 partition functions. It is related to $\Phi_{\hbar}(Z)$ as

$$(C.2) \quad s_b(m) = \exp \left[-\frac{i\pi}{24}(b^2 + b^{-2}) - \frac{i\pi}{2}m^2 \right] \Phi_{2\pi i b^2}(2\pi b m).$$

The function $s_b(m)$ has a balanced reflection property: $s_b(-m) = s_b(m)^{-1}$. In [58] (which is used for some comparisons in Section 7) the authors use a non-standard convention, replacing $s_b(m)$ with $s_b(m)^{-1}$. The easiest way to distinguish definitions of the quantum dilogarithm is to compare lattices of poles and zeroes.

References

- [1] D. Gaiotto, *N=2 dualities*, [arXiv:0904.2715](#).
- [2] D. Gaiotto, G. W. Moore and A. Neitzke, *Four-dimensional wall-crossing via three-dimensional field theory*. *Comm. Math. Phys.* **299** (2010), no. 1 163–224, [[arXiv:0807.4723](#)].
- [3] L. F. Alday, D. Gaiotto and Y. Tachikawa, *Liouville Correlation Functions from Four-Dimensional Gauge Theories*. *Lett. Math. Phys.* **91** (2010), no. 2 167–197, [[arXiv:0906.3219](#)].
- [4] A. Gadde, E. Pomoni, L. Rastelli and S. S. Razamat, *S-duality and 2d Topological QFT*. *JHEP* **1003** (2010) 032, [[arXiv:0910.2225](#)].

- [5] T. Dimofte, D. Gaiotto and S. Gukov, *Gauge Theories Labelled by Three-Manifolds*. [arXiv:1108.4389](#).
- [6] T. Dimofte, M. Gabella and A. B. Goncharov, *K-Decompositions and 3d Gauge Theories*. [arXiv:1301.0192](#).
- [7] T. Dimofte, D. Gaiotto and S. Gukov, *3-Manifolds and 3d Indices*. [arXiv:1112.5179](#).
- [8] S. Cecotti, C. Cordova and C. Vafa, *Braids, Walls, and Mirrors*. [arXiv:1110.2115](#).
- [9] D. Gaiotto and E. Witten, *S-Duality of Boundary Conditions In $N=4$ Super Yang-Mills Theory*. *Adv. Theor. Math. Phys.* **13** (2009), no. 2 721–896, [[arXiv:0807.3720](#)].
- [10] N. Drukker, D. Gaiotto and J. Gomis, *The Virtue of Defects in 4D Gauge Theories and 2D CFTs*. [arXiv:1003.1112](#).
- [11] D. Gaiotto, *Domain Walls for Two-Dimensional Renormalization Group Flows*. *JHEP* **1212** (2012) 103, [[arXiv:1201.0767](#)].
- [12] R. Kashaev, *The quantum dilogarithm and Dehn twists in quantum Teichmüller theory*. Integrable structures of exactly solvable two-dimensional models of quantum field theory (Kiev, 2000), NATO Sci. Ser. II Math. Phys. Chem. **35** (2001) 211–221.
- [13] R. M. Kashaev, *Quantization of Teichmüller Spaces and the Quantum Dilogarithm*. *Lett. Math. Phys.* **43** (1998), no. 2 105–115.
- [14] L. Chekhov and V. V. Fock, *Quantum Teichmüller Space*. *Theoret. and Math. Phys.* **120** (1999), no. 3 1245–1259, [[arXiv:math/9908165v2](#)].
- [15] J. Teschner, *An Analog of a Modular Functor from Quantized Teichmüller Theory*. *Handbook of Teichmüller Theory*, Vol. I (2007) 685–760, [[arXiv:math/0510174v4](#)].
- [16] V. V. Fock and A. B. Goncharov, *Moduli spaces of local systems and higher Teichmüller theory*. *Publ. Math. Inst. Hautes Etudes Sci.* **103** (2006) 1–211, [[arXiv:math/0311149v4](#)].
- [17] T. Dimofte and D. Gaiotto, *An $E7$ Surprise*. [arXiv:1209.1404](#).
- [18] D. Gaiotto, G. W. Moore and A. Neitzke, *Wall-crossing, Hitchin Systems, and the WKB Approximation*. [arXiv:0907.3987](#).

- [19] P. C. Argyres and M. R. Douglas, *New Phenomena in $SU(3)$ Supersymmetric Gauge Theory*. Nucl. Phys. **B448** (1995) 93–126, [[arXiv:hep-th/9505062v1](#)].
- [20] T. Dimofte and S. Gukov, *Chern-Simons Theory and S-duality*. [arXiv:1106.4550](#).
- [21] D. Gaiotto, G. W. Moore and A. Neitzke, *Framed BPS States*. [arXiv:1006.0146](#).
- [22] N. Drukker, D. R. Morrison and T. Okuda, *Loop operators and S-duality from curves on Riemann surfaces*. JHEP **0909** (2009) 031, [[arXiv:0907.2593](#)].
- [23] N. Seiberg and E. Witten, *Electric-Magnetic Duality, Monopole Condensation, and Confinement In $N=2$ Supersymmetric Yang-Mills Theory*. Nucl. Phys. **B426** (1994) 19–52, [[arXiv:hep-th/9407087v1](#)].
- [24] N. Seiberg and E. Witten, *Monopoles, Duality and Chiral Symmetry Breaking in $N=2$ Supersymmetric QCD*. Nucl. Phys. **B431** (1994) 484–550, [[arXiv:hep-th/9408099v1](#)].
- [25] E. Witten, *Solutions of Four-Dimensional Field Theories Via M Theory*. Nucl. Phys. **B500** (Jan, 1997) [[arXiv:hep-th/9703166v1](#)].
- [26] M. F. Atiyah and R. Bott, *The Yang-Mills equations over Riemann surfaces*. Philos. Trans. Roy. Soc. London Ser. A **308** (1983), no. 1505 523–615.
- [27] W. M. Goldman, *The symplectic nature of fundamental groups of surfaces*. Adv. Math. **54** (1984), no. 2 200–225.
- [28] S. Gukov and E. Witten, *Gauge Theory, Ramification, and the Geometric Langlands Program*. Curr. Dev. Math. **2006** (Dec, 2008) 35–180, [[arXiv:hep-th/0612073v2](#)].
- [29] S. Gukov and P. Sulkowski, *A-polynomial, B-model, and Quantization*. [arXiv:1108.0002](#).
- [30] S. Wolpert, *The Fenchel-Nielsen deformation*. Ann. of Math. **115** (1982), no. 3 501–528.
- [31] S. Wolpert, *On the symplectic geometry of deformations of a hyperbolic surface*. Ann. of Math. **117** (1983), no. 2 207–234.
- [32] S. Gukov, *Three-Dimensional Quantum Gravity, Chern-Simons Theory, and the A-Polynomial*. Commun. Math. Phys. **255** (2005), no. 3 577–627, [[arXiv:hep-th/0306165v1](#)].

- [33] D. Cooper, M. Culler, H. Gillet, D. Long and P. Shalen, *Plane Curves Associated to Character Varieties of 3-Manifolds*. *Invent. Math.* **118** (1994), no. 1 47–84.
- [34] T. Dimofte, *Quantum Riemann Surfaces in Chern-Simons Theory*. [arXiv:1102.4847](#).
- [35] W. D. Neumann and D. Zagier, *Volumes of hyperbolic three-manifolds*. *Topology* **24** (1985), no. 3 307–332.
- [36] W. Neumann, *Combinatorics of Triangulations and the Chern-Simons Invariant for Hyperbolic 3-Manifolds*. *Topology '90*, Ohio State Univ. Math. Res. Inst. Publ. **1** (1992).
- [37] W. Thurston, *The Geometry and Topology of Three-Manifolds*. Lecture notes at Princeton University (1980).
- [38] S. S. Pufu and S. Sachdev, *Monopoles in 2+1-dimensional conformal field theories with global $U(1)$ symmetry*. [arXiv:1303.3006](#).
- [39] E. Witten, *$SL(2, Z)$ Action On Three-Dimensional Conformal Field Theories With Abelian Symmetry*. [arXiv:hep-th/0307041v3](#).
- [40] O. DeWolfe, D. Z. Freedman and H. Ooguri, *Holography and Defect Conformal Field Theories*. *Phys. Rev.* **D66** (2002) 025009, [[arXiv:hep-th/0111135v3](#)].
- [41] D. Gaiotto and E. Witten, *Supersymmetric Boundary Conditions in $N=4$ Super Yang-Mills Theory*. *J. Stat. Phys.* **135** (Dec, 2009) 789–855, [[arXiv:0804.2902](#)].
- [42] P. C. Argyres, M. R. Plesser, N. Seiberg and E. Witten, *New $N=2$ Superconformal Field Theories in Four Dimensions*. *Nucl. Phys.* **B461** (1996) 71–84, [[arXiv:hep-th/9511154v1](#)].
- [43] M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi and C. Vafa, *BPS Quivers and Spectra of Complete $N=2$ Quantum Field Theories*. [arXiv:1109.4941](#).
- [44] M. Alim, S. Cecotti, C. Cordova, S. Espahbodi, A. Rastogi and C. Vafa, *$N=2$ Quantum Field Theories and Their BPS Quivers*. [arXiv:1112.3984](#).
- [45] M. Kontsevich, *Lecture presented at the meeting “New Mathematical Structures in Gauge Theory?”* Perimeter Institute, March 2013.
- [46] D. Gaiotto, G. Moore and E. Witten, *In preparation*, .

- [47] O. Aharony, A. Hanany, K. Intriligator, N. Seiberg and M. J. Strassler, *Aspects of $N=2$ Supersymmetric Gauge Theories in Three Dimensions*. Nucl. Phys. **B499** (1997), no. 1-2 67–99, [arXiv:hep-th/9703110v1].
- [48] N. Hama, K. Hosomichi and S. Lee, *SUSY Gauge Theories on Squashed Three-Spheres*. arXiv:1102.4716.
- [49] L. F. Alday, D. Gaiotto, S. Gukov, Y. Tachikawa and H. Verlinde, *Loop and Surface Operators in $N=2$ Gauge Theory and Liouville Modular Geometry*. JHEP **1001** (2010) 113, [arXiv:0909.0945].
- [50] N. Drukker, J. Gomis, T. Okuda and J. Teschner, *Gauge Theory Loop Operators and Liouville Theory*. JHEP **1002** (2010) 057, [arXiv:0909.1105].
- [51] V. Pestun, *Localization of gauge theory on a four-sphere and supersymmetric Wilson loops*. arXiv:0712.2824.
- [52] N. Hama and K. Hosomichi, *Seiberg-Witten Theories on Ellipsoids*. arXiv:1206.6359.
- [53] D. Green, Z. Komargodski, N. Seiberg, Y. Tachikawa and B. Wecht, *Exactly Marginal Deformations and Global Symmetries*. JHEP **1006** (2010) 106, [arXiv:1005.3546].
- [54] V. V. Fock, *Dual Teichmüller spaces*. arXiv:dg-ga/9702018v3.
- [55] V. V. Fock and A. B. Goncharov, *The quantum dilogarithm and representations quantum cluster varieties*. Invent. Math. **175** (2009), no. 2 223–286, [arXiv:math/0702397v6].
- [56] Y. Kabaya, *Parametrization of $PSL(2, C)$ -representations of surface groups*. arXiv:1110.6674.
- [57] K. Hosomichi, S. Lee and J. Park, *AGT on the S-duality Wall*. JHEP **1012** (2010) 079, [arXiv:1009.0340].
- [58] J. Teschner and G. S. Vartanov, *$6j$ symbols for the modular double, quantum hyperbolic geometry, and supersymmetric gauge theories*. arXiv:1202.4698.
- [59] B. Ponsot and J. Teschner, *Clebsch-Gordan and Racah-Wigner coefficients for a continuous series of representations of $Uq(sl(2, R))$* . Commun. Math. Phys. **224** (2001) 613–655, [arXiv:math/0007097v2].
- [60] E. Witten, *Gauge theories and integrable lattice models*. Nucl. Phys. B **322** (1989), no. 3 629–697.

- [61] E. Witten, *Gauge theories, vertex models, and quantum groups*. Nucl. Phys. **B330** (1990), no. 2-3 285–346.
- [62] V. G. Turaev, *Quantum invariants of knots and 3-manifolds*. de Gruyter Studies in Mathematics 18, Walter de Gruyter and co., Berlin **18** (1994) x+588.
- [63] J. Teschner, *Liouville Theory Revisited*. Class. Quant. Grav. **18** (2001), no. 23 R153–R222, [[arXiv:hep-th/0104158v3](#)].
- [64] S. P. Tan, *Complex Fenchel-Nielsen coordinates for quasi-Fuchsian structures*. Internat. J. Math. **5** (1994), no. 2 239–251.
- [65] C. Kourouniotis, *Complex length coordinates for quasi-Fuchsian groups*. Mathematika **41** (1994), no. 1 173–188.
- [66] W. M. Goldman, *The complex-symplectic geometry of $SL(2, C)$ -characters over surfaces*. Algebraic groups and arithmetic, Tata Inst. Fund. Res., Mumbai (2004) 375–407, [[arXiv:math/0304307v2](#)].
- [67] B. Eynard and N. Orantin, *Invariants of Algebraic Curves and Topological Expansion*. Commun. Number Theory Phys. **1** (2007), no. 2 347–452, [[arXiv:math-ph/0702045v4](#)].
- [68] G. Borot and B. Eynard, *All-order asymptotics of hyperbolic knot invariants from non-perturbative topological recursion of A -polynomials*. [arXiv:1205.2261](#).
- [69] S. Garoufalidis, *On the Characteristic and Deformation Varieties of a Knot*. Geom. Topol. Monogr. **7** (2004) 291–304, [[arXiv:math/0306230v4](#)].
- [70] J. E. Andersen and R. Kashaev, *A TQFT from quantum Teichmüller theory*. [arXiv:1109.6295](#).
- [71] S. Garoufalidis, *The 3D index of an ideal triangulation and angle structures*. [arXiv:1208.1663](#).
- [72] C. Beem, T. Dimofte and S. Pasquetti, *Holomorphic Blocks in Three Dimensions*. [arXiv:1211.1986](#).
- [73] K. Hikami, *Generalized Volume Conjecture and the A -Polynomials — the Neumann-Zagier Potential Function as a Classical Limit of Quantum Invariant*. J. Geom. Phys. **57** (2007), no. 9 1895–1940, [[arXiv:math/0604094v1](#)].

- [74] T. Dimofte, S. Gukov, J. Lenells and D. Zagier, *Exact Results for Perturbative Chern-Simons Theory with Complex Gauge Group*. *Comm. Num. Thy. and Phys.* **3** (2009), no. 2 363–443, [[arXiv:0903.2472](https://arxiv.org/abs/0903.2472)].
- [75] E. Barnes, *The Genesis of the Double Gamma Functions*. *Proc. London Math. Soc.* **31** (1899) 358–381.
- [76] L. D. Faddeev, *Discrete Heisenberg-Weyl Group and Modular Group*. *Lett. Math. Phys.* **34** (1995), no. 3 249–254.

INSTITUTE FOR ADVANCED STUDYI, SCHOOL OF NATURAL SCIENCES
EINSTEIN DRIVE, PRINCETON, NJ 08540, USA
E-mail address: tdimofte@gmail.com

PERIMETER INSTITUTE FOR THEORETICAL PHYSICS
WATERLOO, ONTARIO, CANADA N2L 2Y5
E-mail address: dgaiotto@gmail.com

KORTEWEG-DE VRIES INSTITUTE FOR MATHEMATICS, UNIV. OF AMSTERDAM
P.O. BOX 94248, 1090 GE AMSTERDAM, THE NETHERLANDS
E-mail address: R.I.vanderVeen@uva.nl