

K-decompositions and 3d gauge theories

Tudor Dimofte,^{a,b} Maxime Gabella^{a,c} and Alexander B. Goncharov^d

^a*Institute for Advanced Study, Einstein Dr., Princeton, NJ 08540, U.S.A.*

^b*U.C. Davis, Dept. of Mathematics and Center for Quantum Mathematics and Physics, Davis, CA 95616, U.S.A.*

^c*Institut de Physique Théorique, CEA/Saclay, 91191 Gif-sur-Yvette, France*

^d*Yale University Mathematics Dept., New Haven, CT 06520, U.S.A.*

E-mail: tudor@math.ucdavis.edu, gabella@ias.edu,
alexander.goncharov@yale.edu

ABSTRACT: This paper combines several new constructions in mathematics and physics. Mathematically, we study *framed* flat $\mathrm{PGL}(K, \mathbb{C})$ -connections on a large class of 3-manifolds M with boundary. We introduce a moduli space $\mathcal{L}_K(M)$ of framed flat connections on the boundary ∂M that extend to M . Our goal is to understand an open part of $\mathcal{L}_K(M)$ as a Lagrangian subvariety in the symplectic moduli space $\mathcal{X}_K^{\mathrm{un}}(\partial M)$ of framed flat connections on the boundary — and more so, as a “ K_2 -Lagrangian,” meaning that the K_2 -avatar of the symplectic form restricts to zero. We construct an open part of $\mathcal{L}_K(M)$ from elementary data associated with the *hypersimplicial K -decomposition* of an ideal triangulation of M , in a way that generalizes (and combines) both Thurston’s gluing equations in 3d hyperbolic geometry and the cluster coordinates for framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connections on surfaces. By using a canonical map from the complex of configurations of decorated flags to the Bloch complex, we prove that any generic component of $\mathcal{L}_K(M)$ is K_2 -isotropic as long as ∂M satisfies certain topological constraints (theorem 4.2). In some cases this easily implies that $\mathcal{L}_K(M)$ is K_2 -Lagrangian. For general M , we extend a classic result of Neumann and Zagier on symplectic properties of $\mathrm{PGL}(2)$ gluing equations to reduce the K_2 -Lagrangian property to a combinatorial statement.

Physically, we translate the K -decomposition of an ideal triangulation of M and its symplectic properties to produce an explicit construction of 3d $\mathcal{N} = 2$ superconformal field theories $T_K[M]$ resulting (conjecturally) from the compactification of K M5-branes on M . This extends known constructions for $K = 2$. Just as for $K = 2$, the theories $T_K[M]$ are described as IR fixed points of abelian Chern-Simons-matter theories. Changes of triangulation (2–3 moves) lead to abelian mirror symmetries that are all generated by the elementary duality between $N_f = 1$ SQED and the XYZ model. In the large K limit, we find evidence that the degrees of freedom of $T_K[M]$ grow cubically in K .

KEYWORDS: Differential and Algebraic Geometry, Supersymmetric gauge theory, Supersymmetry and Duality

ARXIV EPRINT: [1301.0192](https://arxiv.org/abs/1301.0192)

Contents

1	Introduction	1
1.1	The 3d-3d correspondence	5
1.2	A mathematical perspective	9
1.3	Organization	19
2	Basic tools and definitions	20
2.1	Gluing admissible 3-manifolds from truncated tetrahedra	20
2.2	Configurations of flags	22
2.3	The basic moduli spaces	23
2.4	From framed flat bundles to configurations of flags	26
2.5	Hypersimplicial K -decomposition	27
2.6	Configuration of flags and hypersimplices	29
3	Localization of framed flat connections	31
3.1	Boundary phase spaces	31
3.1.1	Coordinates from the K -triangulation	32
3.1.2	Symplectic structure	35
3.2	The Lagrangian pair assigned to a 3d triangulation of M	35
3.2.1	The Lagrangian pair for the $\mathrm{PGL}(2)$ -tetrahedron	36
3.2.2	The Lagrangian pair for the $\mathrm{PGL}(K)$ -tetrahedron	37
3.2.3	The coordinate phase space for a tetrahedron via symplectic reduction	41
3.2.4	Gluing constraints for 3-manifolds	42
3.3	Holonomy representation	44
3.3.1	3d snakes in a tetrahedron	44
3.3.2	3d snakes in a 3-manifold	49
3.4	Small torus boundaries	52
3.5	2–3 moves via 4d cobordism	54
3.5.1	Coordinates	56
4	K_2-Lagrangians and the Bloch complex	58
4.1	Perspective	59
4.2	From decorated flags to the Bloch complex	61
4.2.1	Decorated flags	61
4.2.2	BiGrassmannian complex	62
4.2.3	Bloch complex	66
4.2.4	Conclusion: from decorated flags complex to the Bloch complex	67
4.3	$\mathcal{L}_K(M)$ is a K_2 -isotropic subvariety	68
4.4	Differential of the volume of a framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connection	71

5	Combinatorics of K-decompositions	73
5.1	Octahedra and slices	74
5.2	Coordinates on the product phase space	75
5.2.1	Paths on slices	76
5.2.2	Gluing functions	77
5.2.3	Big boundary coordinates	79
5.2.4	Small torus coordinates	80
5.3	Poisson brackets and the quotient	81
5.3.1	Unipotent conditions	83
5.3.2	Symplectic data	84
5.4	Remarks on quantization	85
5.5	2–3 moves and path coordinates	86
5.5.1	Elementary moves	86
5.5.2	Composing elementary moves	88
6	Tetrahedron and polyhedron theories	89
6.1	Review: the octahedron theory	90
6.2	Warmup: elementary 2–3 move theories	92
6.3	Tetrahedra at higher K	95
6.3.1	$K = 3$	95
6.3.2	$K = 4$	96
6.3.3	$K = 5$	98
6.4	The $K = 3$ bipyramid	100
6.4.1	Two tetrahedra	101
6.4.2	Three tetrahedra	103
6.4.3	Mirror symmetry	105
7	Knot complement theories	106
7.1	Flavor symmetry and marginal operators of $T_K[M]$	108
7.1.1	Real masses and boundary holonomy	110
7.2	K^3 scaling in $T_K[M]$	111
7.2.1	$\mathrm{PGL}(K)$ volume from octahedra	111
7.2.2	From free energy to volumes	113
7.3	Trefoil knot invariants	114
7.3.1	$\mathrm{SL}(2)$ invariants	114
7.3.2	The $\mathrm{SL}(3)$ character variety	115
7.3.3	Symplectic gluing data	116
7.4	Figure-eight knot gluing	118
A	Symplectic data and class \mathcal{R}	120
A.1	The symplectic data	121
A.2	Constraints on the data	122
A.3	The gauge theory	122

A.4	Equivalences	125
A.5	The existence of operators \mathcal{O}_k	126
A.6	Associated invariants	127
B	The Poisson bracket for eigenvalues	129
B.1	$\{C, U\} = \{\mathcal{X}, U\} = 0$	130
B.2	$\{U^\gamma, U^\lambda\} = \kappa \cdot \langle \gamma, \lambda \rangle$	132
B.2.1	Parallel slices	133
B.2.2	Slices in different families	136

1 Introduction

This paper presents a combination of mathematical and physical results. Its main goal is a physical one: to algorithmically define three-dimensional supersymmetric $\mathcal{N} = 2$ quantum field theories $T_K[M]$ labeled by an oriented topological 3-manifold M and an integer $K \geq 2$. The theories $T_K[M]$ are meant to coincide with the compactification of the six-dimensional $\mathcal{N} = (2, 0)$ superconformal theory with symmetry algebra A_{K-1} on 3-manifolds M . This implies that the theories $T_K[M]$ should possess several important properties, relating their observables to the topology and geometry of M — in particular, to the geometry of the moduli space of flat $\mathrm{SL}(K, \mathbb{C})$ -connections on M and its quantization. These properties are summarized in table 1 below (page 6). The combinatorial definition we give of $T_K[M]$ makes many of its expected properties manifest.

For $K = 2$, and a particular class of 3-manifolds M with boundary, theories $T_2[M]$ were defined in [1]. The key idea of [1] was to decompose M into topological ideal tetrahedra (i.e. to choose a *triangulation* of M). Then, after assigning a canonical “tetrahedron theory” $T_2[\Delta]$ to each tetrahedron, $T_2[M]$ was constructed by “gluing” together the tetrahedron theories

$$M = \bigcup_{i=1}^N \Delta_i \quad \Rightarrow \quad T_2[M] = (T_2[\Delta_1] \otimes \cdots \otimes T_2[\Delta_N]) / \sim . \quad (1.1)$$

Physically, each $T_2[\Delta_i]$ contained a 3d chiral multiplet, and the gluing operation ‘ \sim ’ added superpotential interactions and gauged some global symmetries, producing an abelian Chern-Simons-matter theory with an explicit Lagrangian. The final step in the gluing was to flow to the infrared, defining $T_2[M]$ as the infrared limit of the Chern-Simons-matter theory.

The gluing of (1.1) was done in such a way that, generically, the definition of $T_2[M]$ would be independent of the choice of triangulation. Topologically, any two ideal triangulations can be related by local 2–3 Pachner moves as depicted in figure 1 [2, 3]. The basic theory $T_2[\text{bipyramid}]$ associated with a triangular bipyramid and constructed from the triangulation on the l.h.s. of figure 1 was 3d $\mathcal{N} = 2$ quantum electrodynamics (SQED) with two chiral multiplets. The theory constructed from the triangulation on the r.h.s. was the “XYZ model,” consisting of three chiral multiplets and a cubic interaction. These




Figure 1. The 2–3 move: decomposing a bipyramid into two or three tetrahedra.

two 3d theories are equivalent (dual) in the infrared [4], ensuring a local triangulation-independence. Lifting the local independence to a global independence of any theory $T_2[M]$ turned out to be subtle for two reasons: 1) it required interchanging infrared limits for different pieces of a Chern-Simons-matter Lagrangian; and 2) not all triangulations could be used to produce sensible Chern-Simons-matter Lagrangians, and the “good” triangulations that work are not all known to be related by sequences of 2–3 moves. Thus, strictly speaking, the triangulation-independence of $T_2[M]$ was only conjectural.

The observables of theories $T_2[M]$ were related to classical and quantum hyperbolic geometry, because flat $\mathrm{SL}(2, \mathbb{C})$ -connections on M are (roughly speaking) hyperbolic metrics.¹ The properties that allowed $T_2[M]$ to be defined as in (1.1) were a direct generalization of the symplectic properties that Neumann and Zagier [8] observed in Thurston’s gluing equations [5] for ideal hyperbolic tetrahedra. These same symplectic properties allowed Thurston’s gluing equations to be quantized in [9] (following [10, 11]). The lift of Thurston’s gluing equations to theories $T_2[M]$ amounts to a categorification of hyperbolic geometry — though many details of this categorification remain to be worked out.

In this paper, we extend the triangulation methods of [1] to general $K \geq 2$. This requires developing some new mathematics. At the classical level, we need to describe flat $\mathrm{PGL}(K, \mathbb{C})$ -connections on triangulated 3-manifolds in a manner analogous to Thurston’s description of hyperbolic metrics. To achieve this, we enhance moduli spaces of flat connections on a 3-manifold M with additional *framing data*, much as in the 2d constructions of [12]. Namely, we consider flat connections together with a choice of invariant flags along certain loci on the boundary ∂M . This requires the introduction of extra topological data on ∂M .² Then, for admissible 3-manifolds M , we construct two algebraic varieties

$$\begin{aligned} \tilde{\mathcal{L}}_K(M) &\simeq \{\text{framed flat } \mathrm{PGL}(K, \mathbb{C})\text{-connections on } M\}, \\ \mathcal{X}_K^{\mathrm{un}}(\partial M) &\simeq \{\text{framed unipotent flat } \mathrm{PGL}(K, \mathbb{C})\text{-connections on } \partial M\}. \end{aligned} \tag{1.2}$$

Although we do not give precise definitions here, let us present a basic example.

¹Precisely, hyperbolic metrics on M are in 1-1 correspondence with flat $\mathrm{PGL}(2, \mathbb{C})$ -connections whose holonomy representation $\rho : \pi_1(M) \rightarrow \mathrm{PGL}(2, \mathbb{C})$ is discrete, faithful, and torsion-free, cf. [5]. The ρ is the holonomy representation of the hyperbolic metric. The relation between hyperbolic geometry and $\mathrm{PGL}(2, \mathbb{C})$ -connections was heavily exploited in [6] and subsequent works to study quantum gravity in three dimensions, and in [7] and subsequent works to understand aspects of quantum Chern-Simons theory.

²Examples of this “extra structure” include the laminations discussed in early versions of [13], and expanded on in [14]. The slightly more general notion of framing data that we present in section 2 arose in collaboration with the authors of [15].

Let M be a polyhedron. Then the moduli space $\mathcal{X}_K^{\text{un}}(\partial M)$ parametrizes flat $\text{PGL}(K, \mathbb{C})$ -connections on the sphere ∂M punctured at the vertices of the polyhedron, with unipotent monodromies around the vertices, plus a choice of an invariant flag near each of the vertices. Turning to $\tilde{\mathcal{L}}_K(M)$, it makes little sense to talk about a moduli space of just flat connections on a polyhedron, since a polyhedron is simply connected, and so any flat connection is trivial. However, the *framed* moduli space of a polyhedron parameterizes configurations of flags in \mathbb{C}^K , labelled by the vertices of the polyhedron. For example, in the most fundamental example when M is a tetrahedron we get a configuration space of four flags in \mathbb{C}^K .

The space $\mathcal{X}_K^{\text{un}}(\partial M)$ is a singular complex symplectic space, which carries a canonical symplectic form on its non-singular part, and more so, a K_2 -avatar of the symplectic form [12]. We consider the image of the natural projection

$$\mathcal{L}_K(M) := \text{im}(\tilde{\mathcal{L}}_K(M) \rightarrow \mathcal{X}_K^{\text{un}}(\partial M)). \tag{1.3}$$

When M is a knot complement and $K = 2$, the subvariety $\mathcal{L}_K(M)$ is a curve defined by the A-polynomial of the knot [16]. The spaces $\mathcal{L}_K(M)$ and $\tilde{\mathcal{L}}_K(M)$ may have several irreducible components. We introduce a notion of a generic component, and conjecture that any generic component of $\mathcal{L}_K(M)$ is a K_2 -Lagrangian subvariety of $\mathcal{X}_K^{\text{un}}(\partial M)$. This implies that it is Lagrangian for the symplectic form on $\mathcal{X}_K^{\text{un}}(\partial M)$. We prove that, assuming that M has no toric components on the boundary, any generic component of $\mathcal{L}_K(M)$ is K_2 -isotropic.³

The moduli space $\mathcal{X}_K^{\text{un}}(\partial M)$, as well as a larger moduli space $\mathcal{X}_K(\partial M)$ where the unipotence condition is dropped, were introduced and studied in [12], generalizing the classical Teichmüller theory. The moduli space $\mathcal{X}_K^{\text{un}}(\partial M)$ is birationally isomorphic to Hitchin’s moduli space \mathcal{M}_H for ∂M in any fixed complex structure. Hitchin’s moduli space, as a hyperkähler manifold, has played a major role in the study of 4d supersymmetric gauge theories, cf. [21, 22]. In applications it is often important to consider a hyperkähler resolution of singularities in \mathcal{M}_H . The framing data of $\mathcal{X}_K^{\text{un}}(\partial M)$ in (1.2) partially resolves the space of flat $\text{PGL}(K, \mathbb{C})$ -connections on ∂M and can be thought of as an algebraic counterpart for the hyperkähler resolution of \mathcal{M}_H .

For $K = 3$, the variety $\tilde{\mathcal{L}}_K(M)$ was described in [23]. For general $K > 2$, the space $\tilde{\mathcal{L}}_K(M)$ was studied in [24], following [25–27]. We will comment on the relation between these works and the current paper at the end of section 1.2.

In [12] it was shown that an open part of $\mathcal{X}_K(\partial M)$ is covered by affine charts $\mathcal{X}_K(\partial M, \mathbf{t})$ of specific type, the cluster Poisson coordinate charts, labeled by 2d ideal triangulations \mathbf{t} of ∂M . Here we similarly glue together $\tilde{\mathcal{L}}_K(M)$ from affine varieties $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$, labeled by 3d ideal triangulations \mathbf{t}_{3d} . To define $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$, we use an important tool: the

³The graph of any cluster transformation is K_2 -isotropic, and hence K_2 -Lagrangian: this follows immediately from the basic relation of cluster transformations to the Bloch complex [17, Sec 6], and provided first general examples of K_2 -Lagrangians. In particular, when M is a cobordism of triangulated 2d surfaces generated by flips of the triangulation, this implies that $\mathcal{L}_K(M)$ is a K_2 -Lagrangian. The K_2 -Lagrangian property of the cluster transformations is crucial for their quantization. The K_2 -Lagrangian property of A-polynomials was discussed in [18, 19], and argued to be necessary for quantization in [20].




Figure 2. Stacking $\frac{1}{6}K(K^2 - 1)$ octahedra to form an A_{K-1} tetrahedron (for $K = 4$).

hypersimplicial K -decomposition of an ideal triangulation, or *K -decomposition* for short. Given an ideal triangulation $\mathfrak{t}_{3d} = \{\Delta_i\}_{i=1}^N$ of M , we further divide each tetrahedron Δ_i into $\frac{1}{6}K(K^2 - 1)$ octahedra \diamond_j (figure 2). To each octahedron we assign a triple of “cross-ratio parameters” z, z', z'' with $zz'z'' = -1$, and the variety $\tilde{\mathcal{L}}_K[M, \mathfrak{t}_{3d}]$ is cut out from $(\mathbb{C}^*)^{2 \cdot \#\text{octahedra}} = (\mathbb{C}^*)^{\frac{1}{3}NK(K^2-1)}$ by *gluing equations* that include

$$z'' + z^{-1} = 1 . \tag{1.4}$$

These gluing equations are analogues of Thurston’s gluing equations from hyperbolic geometry. Symplectic properties of the gluing equations (which we prove in some cases) provide one of the ways to prove that $\mathcal{L}_K(M) \subset \mathcal{X}_K^{\text{un}}(\partial M)$ is K_2 -isotropic.

Some features of the cluster coordinate charts for $\mathcal{X}_K(\partial M)$ were recently generalized by the spectral networks of [28, 29]. The generalization uses the geometry of K -fold spectral covers of the surface ∂M . It would be very interesting to describe the variety $\tilde{\mathcal{L}}_K(M)$ in a similar way, using K -fold covers of the 3-manifold M , perhaps along the lines of [30, 31].

Coming back to 3d gauge theories: we use the structure of the space $\tilde{\mathcal{L}}_K(M)$, and in particular the symplectic properties of gluing equations, to formulate a definition of $T_K[M]$. Namely, given a triangulation $M = \bigcup_{i=1}^N \Delta_i$ and a further K -decomposition into $\frac{1}{6}NK(K^2 - 1)$ octahedra \diamond_j , we propose that

$$T_K[M] = \left(\bigotimes_{i=1}^N T_K[\Delta_i] \right) / \sim = \left(\bigotimes_{j=1}^{\frac{1}{6}NK(K^2-1)} T[\diamond_j] \right) / \sim , \tag{1.5}$$

in direct analogy to (1.1). Here $T[\diamond_j]$ is the canonical theory of a single chiral multiplet (identical to $T_2[\Delta]$), and gluing is again implemented by superpotentials and gauging. Thus $T_K[M]$ acquires a description as the infrared limit of an abelian Chern-Simons-matter theory. The invariance of $T_K[M]$ under 2–3 Pachner moves follows (again, conjecturally) from basic 2–3 moves for the octahedra.

The theory $T_K[\Delta]$ associated with a single tetrahedron at higher K is no longer so simple or canonical as for $K = 2$. The definition (1.5) implies that its UV Lagrangian contains $\frac{1}{6}K(K^2 - 1)$ chiral multiplets. For the first few K , we find in section 6:

K	$T_K[\Delta]$	flavor sym.
2	1 free chiral	$U(1)$
3	4 free chirals	$U(1)^4$
4	10 chirals + degree-six superpotential	$U(1)^9$
5	20 chirals, $U(1)^2$ gauge group, $W = \sum(\text{four monopole ops})$	$U(1)^{16}$

(1.6)

In general, for all $K \geq 5$, the gluing rules force $T_K[\Delta]$ to have both a nontrivial superpotential and a nontrivial gauge group.

Constructing theories $T_K[M]$ for all $K \geq 2$ allows us to study some features of the $K \rightarrow \infty$ limit. In particular, by relating degrees of freedom of theories $T_K[M, \Pi]$ to the volume of $\text{PGL}(K, \mathbb{C})$ -connections on M , we will find evidence of the large- K scaling

$$\# \text{ degrees of freedom of } T_K[M] \sim \# \text{ octahedra} \sim K(K^2 - 1). \quad (1.7)$$

This agrees beautifully with predictions from M-theory and holography [32, 33]. We note, however, that K does not appear as a continuously tunable parameter in $T_K[M]$, such as the rank of a gauge group. (This is obvious, for example, in (1.6).) Recently a number of 3d theories that do allow analytic continuation in K were studied in [34]; it would be very interesting to relate these to the present constructions.

As we review momentarily, part of the 3d-3d correspondence relates partition functions of $T_K[M]$ on spheres (and more generally lens spaces) to partition functions of Chern-Simons theory on M itself, with complex gauge group $\text{PGL}(K, \mathbb{C})$. Thus, the combinatorial construction of UV Lagrangians for theories $T_K[M]$ implies a combinatorial construction of partition functions for $\text{PGL}(K, \mathbb{C})$ Chern-Simons theory on M . Alternatively, one may simply say that the symplectic properties of $\text{PGL}(K, \mathbb{C})$ gluing equations allow a systematic quantization of the pair of spaces $\mathcal{L}_K(M) \subset \mathcal{X}_K^{\text{un}}(\partial M)$, generalizing [9]. The result is a construction of a subsector of $\text{PGL}(K, \mathbb{C})$ Chern-Simons theory that generalizes a circle of ideas initiated in [7, 11] for $K = 2$.⁴

In the remainder of this introduction, we review some basic features of the 3d-3d correspondence between observables of $T_K[M]$ and the geometry of flat $\text{PGL}(K, \mathbb{C})$ -connections on M ; then we provide a more detailed summary of our main mathematical results.

1.1 The 3d-3d correspondence

The 3d-3d correspondence was first conjectured in [38], and has since been studied in a multitude of papers. Some of its fundamental ideas were developed in [1, 30, 39], and the first physical proofs of the correspondence appeared in [40, 41]. (The recent review [42] contains further discussion and references.) New and interesting aspects of the correspondence are still being developed, cf. [35–37, 43], which appeared after the first version of this paper.

The basic idea behind the 3d-3d correspondence is that the compactification of the 6d (2,0) theory of type A_{K-1} on $\mathbb{R}^3 \times M$, with a topological twist along M that preserves four supercharges, should lead to a 3d $\mathcal{N} = 2$ theory $T_K[M]$ in \mathbb{R}^3 . Alternatively, $T_K[M]$ may be defined as the effective field theory on a stack of K M5-branes wrapping $M \times \mathbb{R}^3$ in the M-theory background $\mathbb{R}^5 \times T^*M$. More generally, compactification of the 6d (2,0) theory (or of 5-branes) on a d -manifold X^d should produce an effective theory $T_K[X^d]$ in

⁴The $\text{PGL}(K, \mathbb{C})$ Chern-Simons partitions constructed with the current methods have limited TQFT-like properties under cutting-and-gluing operations that preserve a triangulation and some additional structures. It is nevertheless an open problem to define $\text{PGL}(K, \mathbb{C})$ Chern-Simons theory (for any $K \geq 2$) as a full TQFT, with complete cutting-and-gluing rules. This has recently been emphasized in [35–37], along with new proposals for resolving the problem.

$T_K[M, \Pi]$		$M, G_{\mathbb{C}}$ connections
coupling to $T_K[\partial M]$	\longleftrightarrow	polarization Π for $\mathcal{X}_K^{\text{un}}(\partial M)$
rank of flavor symmetry group	\longleftrightarrow	$\frac{1}{2} \dim_{\mathbb{C}} \mathcal{X}_K^{\text{un}}(\partial M)$
twisted superpotential on $\mathbb{C} \times S^1$	\longleftrightarrow	classical $G_{\mathbb{C}}$ Chern-Simons functional
SUSY parameter space on $\mathbb{C} \times S^1$	\longleftrightarrow	$\mathcal{L}_K(M)$
flavor line operators	\longleftrightarrow	quantized algebra of functions on $\mathcal{X}_K^{\text{un}}(\partial M)$
identities for line operators	\longleftrightarrow	quantization of $\mathcal{L}_K(M)$
partition function on lens space $L(k, 1)$	\longleftrightarrow	$G_{\mathbb{C}}$ Chern-Simons theory on M at level k
index on $S^2 \times_q S^1$	\longleftrightarrow	$G_{\mathbb{C}}$ Chern-Simons theory on M at level $k = 0$
holomorphic blocks on $\mathbb{C} \times_q S^1$	\longleftrightarrow	analytically cont'd CS wavefunctions on M

Table 1. A sampling of relations in the 3d-3d correspondence.

\mathbb{R}^{6-d} . Notable examples include the 2d-4d correspondence developed in [22, 44–46] and many other works, and the 4d-2d correspondence introduced in [47].

It is often expected, but so far not strictly proven, that the theory $T_K[M]$ only depends on the topology of M , along with a bit of boundary data discussed below. From a field-theory perspective, this statement crucially depends on defining $T_K[M]$ as the *infrared limit* of the compactification of the 6d (2, 0) theory on M . Away from the infrared, all bets are off. Indeed, one may view a particular triangulation (or K -decomposition) of M as a choice of discretized metric, and it is clear even in the simplest examples (such as SQED \simeq XYZ associated to the Pachner move of figure 1) that two triangulations only define the same theory after flowing to the infrared. In the infrared, the expectation⁵ is that $T_K[M]$ only depends on a choice of uniformized, constant-curvature metric on M . If M happens to be hyperbolic, then this choice is unique, whence metric invariants are actually topological invariants. However, if M is not hyperbolic and (more importantly) the constant-curvature metric is not rigid, then $T_K[M]$ should admit additional parameters corresponding to metric deformations. We will not encounter this latter scenario in this paper.

Various properties of the parent (2, 0) theory imply relations between observables in the theory $T_K[M]$ and the geometry of flat $G_{\mathbb{C}}$ connections on M , where $G_{\mathbb{C}}$ is a complex group with Lie algebra A_{K-1} .⁶ Some of these relations are summarized in table 1.

An important aspect to mention is that when M has a boundary (which will be true throughout this paper), the theory $T_K[M]$ is not truly an isolated 3d theory, but rather a boundary condition for a 4d theory $T_K[\partial M]$. By the 2d-4d correspondence, the rank of the gauge group of $T_K[\partial M]$ is half the dimension of the space of flat connections on the boundary $\mathcal{X}_K^{\text{un}}(\partial M)$. To obtain an isolated 3d theory, we must choose a weak-coupling

⁵See some relevant discussion in the Introduction of [42]. The statement about uniformization comes in part from holography. It was analyzed carefully by [48] for the analogous 2d-4d correspondence, though it has yet to be fully developed for the 3d-3d correspondence.

⁶Some subtle discrete choices can modify the theory $T_K[M]$ and the appropriate form of the group $G_{\mathbb{C}}$. See, for example, [49].

limit⁷ for the theory $T_K[\partial M]$ [15]. This choice amounts to a polarization Π of one of affine charts of the complex symplectic space $\mathcal{X}_K(\partial M)$. The resulting 3d theory should be denoted $T_K[M, \Pi]$. It acquires a flavor symmetry of rank $\frac{1}{2} \dim \mathcal{X}_K^{\text{un}}(\partial M)$.

By the 2d-4d correspondence, putting the 4d theory $T_K[\partial M]$ in a background with angular momentum quantizes its algebra of line operators (see [50] and similar ideas in [51]). An example of such a background is $(\mathbb{C} \times_q S^1) \times \mathbb{R}_+$, where $(\mathbb{C} \times_q S^1) := \{(w, t) \in \mathbb{C} \times [0, 1]\} / (w, 1) \sim (qw, 0)$. The resulting quantum algebra $\hat{\mathcal{X}}_K^{\text{un}}(\partial M)$ is a quantization of $\mathcal{X}_K^{\text{un}}(\partial M)$ — it is the same quantization presented in [52, 53] for $K = 2$ and [12, 13, 17] for $K > 2$. When acting on a 3d boundary theory $T_K[M, \Pi]$, line operators of $T_K[\partial M]$ satisfy additional relations. These relations define an $\hat{\mathcal{X}}_K^{\text{un}}(\partial M)$ -module $\hat{\mathcal{L}}_K(M)$ whose characteristic variety (or “classical limit”) should be $\mathcal{L}_K(M)$. In other words, the module $\hat{\mathcal{L}}_K(M)$ is a quantization of the Lagrangian $\mathcal{L}_K(M)$.

Some of the more interesting observables of a 3d $\mathcal{N} = 2$ theory $T_K[M, \Pi]$ are partition functions on squashed lens spaces $L(k, 1)_b$. For $k = 0$, $k = 1$, and $k > 1$, respectively, it was conjectured in [1, 54], and (after the first version of this paper) in [55] that these partition functions are equivalent to partition functions of $G_{\mathbb{C}}$ Chern-Simons theory at level k on M itself. This was proven for $k = 0, 1$ in [40, 41]. The quantum algebra of line operators (in fact, two commuting copies of it) acts on the $L(k, 1)_b$ partition function of $T_K[M, \Pi]$ in such a way that the relations of the module $\hat{\mathcal{L}}_K(M)$ are satisfied. Correspondingly, a quantization of the Lagrangian variety $\mathcal{L}_K(M)$ should annihilate a $G_{\mathbb{C}}$ Chern-Simons partition function on M [7].

A somewhat simpler observable of $T_K[M, \Pi]$ is the effective twisted superpotential of the theory on $\mathbb{C} \times S^1$ (at $q = 1$). It is the generating function for the supersymmetric parameter space that also appears on the l.h.s. of table 1. As a concrete example, let’s review how this works for a standard $K = 2$ tetrahedron and for the 2–3 Pachner move.

The tetrahedron theory $T_2[\Delta] = T[\diamond]$ contains a single free chiral multiplet ϕ , with a $U(1)$ flavor symmetry (see section 6.1). Its effective twisted superpotential on $\mathbb{C} \times S^1$ is

$$\widetilde{W}(z) = \text{Li}_2(z^{-1}) . \tag{1.8}$$

The dilogarithm here is a standard 1-loop contribution to the twisted superpotential coming from Kaluza-Klein modes of the 3d chiral ϕ [56]. The complex parameter z is the twisted mass of ϕ , associated with its $U(1)$ symmetry. The supersymmetric parameter space is then defined by

$$\mathcal{L}_2[\Delta] := \left\{ \exp \frac{d\widetilde{W}}{d \log z} = z'' \right\} = \{z'' + z^{-1} = 1\} . \tag{1.9}$$

By the 3d-3d correspondence, this should correspond to the Lagrangian $\mathcal{L}_2[\Delta]$. Indeed, (1.9) is identical to (1.4), and we recognize it as the canonical K_2 -Lagrangian subvariety of $\mathcal{X}_2[\partial \Delta] \simeq (\mathbb{C}^*)^2$, with symplectic form $d \log z'' \wedge d \log z$ and K_2 form $z'' \wedge z$. The superpotential $\widetilde{W}(z)$ itself is interpreted as the complexified hyperbolic volume of an ideal hyperbolic tetrahedron.

⁷More precisely, we mean here a weak-coupling limit for the abelian Seiberg-Witten theory that describes $T_K[\partial M]$ on its Coulomb branch. Such a weak-coupling limit always exists.

Now consider a bipyramid as in figure 1. Triangulating by two tetrahedra (with a suitable choice of boundary polarization) leads to a description of $T_2[\text{bipyramid}]$ as 3d $\mathcal{N} = 2$ SQED, i.e. a $U(1)$ gauge theory with two chiral multiplets. Its twisted superpotential is

$$\widetilde{W}_{(2)}(x, y) = \left[\text{Li}_2(s^{-1}) + \text{Li}_2(s/x) + \frac{1}{2} \log(s)^2 + \log s(\log y - i\pi) \right]_{d/ds=0}. \quad (1.10)$$

Each dilogarithm comes from a chiral multiplet of SQED (geometrically: from a tetrahedron), and the effect of the $U(1)$ gauge group is to extremize the superpotential with respect to a dynamical variable s . In contrast, triangulation by three tetrahedra gives the XYZ model, whose twisted superpotential is

$$\widetilde{W}_{(3)}(x, y) = \text{Li}_2(x^{-1}) + \text{Li}_2(y^{-1}) + \text{Li}_2(xy) + \frac{1}{2}(\log(xy) - i\pi)^2. \quad (1.11)$$

This time there is no minimization (since there is no dynamical gauge group), but the arguments of the dilogarithms are constrained by the superpotential of the XYZ model, so that their product is one. Geometrically, the constraint is associated with the internal edge of the triangulation.

Twisted superpotentials of a 3d $\mathcal{N} = 2$ theory are insensitive to renormalization-group flow. (Supersymmetry protects them from quantum corrections beyond one-loop.) Then the infrared equivalence of SQED and the XYZ model implies that we should have

$$\widetilde{W}_{(2)}(x, y) = \widetilde{W}_{(3)}(x, y), \quad (1.12)$$

with appropriate choices of branch cuts. Mathematically (1.12) is the 5-term identity for the dilogarithm. The superpotentials $\widetilde{W}_{(2)}$ and $\widetilde{W}_{(3)}$ simply compute the complexified hyperbolic volume of the bipyramid in two different ways. The SUSY parameter space of $T_2[\text{bipyramid}]$ is $\mathcal{L}_2[\text{bipyramid}] = \{d\widetilde{W}(x, y)/d \log x = p_x, d\widetilde{W}(x, y)/d \log y = p_y\}$; it is an algebraic variety that may equivalently be computed from either $\widetilde{W}_{(2)}$ or $\widetilde{W}_{(3)}$.

Heuristically, the theory $T_K[M]$ should be thought of as a categorical analogue of the motivic $\text{PSL}(K, \mathbb{C})$ -volume of M , or (equivalently) the $\text{PSL}(K, \mathbb{C})$ -class of M in the Bloch group (see pages 15–17). For $K = 2$, this perspective was advocated in the introduction to [1]. In brief, the Bloch group $B_2(F)$ of a field F is the abelian group generated by elements $\{z\}_2$ with $z \in F^* - \{1\} = F - \{0, 1\}$, modulo five-term relations

$$[x] - [y] + \left[\frac{y}{x} \right] - \left[\frac{1-x^{-1}}{1-y^{-1}} \right] + \left[\frac{1-x}{1-y} \right] = 0, \quad x, y \in F^* - \{1\}. \quad (1.13)$$

For a three-manifold M , we take $F = \mathbb{C}(\widetilde{\mathcal{L}}_K(\partial M))$ to be the field of functions on (a regular component of) the moduli space of framed flat $\text{PGL}(K, \mathbb{C})$ -connections on M . Then the motivic $\text{PGL}(K, \mathbb{C})$ -volume of M is an element $\{M\}_2 = \sum_{i=1}^{\#\text{octahedra}} \{z_i\}_2 \in B_2(F)$, obtained by summing the parameters z_i associated to each octahedron in a K -decomposition of M , viewing these parameters as functions on $\widetilde{\mathcal{L}}_K(\partial M)$. The five-term relations ensure that the element $\{M\}_2$ is invariant under (generic) 2–3 moves. The real $\text{PGL}(K, \mathbb{C})$ -volume of M is obtained by applying the Bloch-Wigner dilogarithm map (4.57)

to the motivic volume. We may compare this to the theory $T_K[M, \Pi]$. It is defined by taking a tensor product of octahedron theories and adding some interactions to enforce the gluing — similar to the formal sum of octahedron parameters defining $\{M\}_2$. Just like $\{M\}_2$, the theory $T_K[M, \Pi]$ is invariant under (generic) 2–3 moves. The real $\mathrm{PGL}(K, \mathbb{C})$ -volume of M may be recovered (roughly) by evaluating the real part of the twisted superpotential of the theory on $\mathbb{C} \times S^1$; but by instead evaluating partition functions of $T_K[M, \Pi]$ one also recovers many quantum invariants, as in table 1.

Finally, we note that while in this paper we construct theories $T_K[M]$ as *abelian* Chern-Simons-matter theories, one might generally expect them to have dual non-abelian descriptions as well, related to the $\mathfrak{su}(K)$ symmetry of the underlying 6d (2,0) theory. In principle, a non-abelian construction of $T_K[M]$ would arise from decomposing M into pieces (e.g. handlebodies) glued along *smooth* surfaces, as apposed to the triangulations and K -decompositions of this paper. Some isolated examples of non-abelian $T_K[M]$ have recently appeared (after the initial version of this paper) in [36, 37, 43], but a systematic construction remains an open problem, even for $K = 2$.

1.2 A mathematical perspective

Given a 3-manifold M with boundary, consider the moduli space $\mathrm{Loc}_K(\partial M)$ of flat $\mathrm{PGL}(K, \mathbb{C})$ -connections on the boundary ∂M . This space is symplectic, with the symplectic form at a generic point given by the Atiyah-Bott/Goldman construction. The moduli space of flat connections on the boundary ∂M that can be extended to M is expected to be a Lagrangian subvariety $\mathcal{L}_K(M) \subset \mathrm{Loc}_K(\partial M)$.

One can quantize the symplectic space $\mathrm{Loc}_K(\partial M)$ by defining a non-commutative q -deformation of the $*$ -algebra of regular functions on $\mathrm{Loc}_K(\partial M)$, and constructing its $*$ -representation in an infinite-dimensional Hilbert space $\mathcal{H}_K(\partial M)$. This has nothing to do with a 3-manifold: the problem makes sense for any oriented 2d surface \mathcal{C} . Assuming that the surface is hyperbolic, that is $\chi(\mathcal{C}) < 0$, and has at least one hole, the quantization was done in [12, 13, 17]. It generalizes the quantization of Teichmüller spaces [52, 53], related to the $\mathrm{SL}(2)$ case.

The next goal is to quantize the Lagrangian subvariety $\mathcal{L}_K(M)$. By this we mean defining a line in $\mathcal{H}_K(\partial M)$ that must be annihilated by the q -deformations of the equations defining the subvariety $\mathcal{L}_K(M)$. In the $\mathrm{SL}(2)$ case, this was discussed in a series of physics papers [7, 9–11, 20, 54, 57, 58], and in the closely related mathematical works [59–62]. (An alternative mathematical approach to the quantization of $\mathcal{L}_K(M)$ has been proposed using the theory of skein modules, cf. [63, 64].) Subsequent to the initial version of this paper, further constraints on a consistent quantization were uncovered in [55, 65].

The following problem motivated this project. We would like to have a local procedure for the quantization of the moduli space of flat $\mathrm{PGL}(K, \mathbb{C})$ -connections on 3-manifolds, by decomposing the manifold into tetrahedra, quantizing the tetrahedra, and then gluing the quantized tetrahedra. This approach for $\mathrm{SL}(2)$ was implemented in [9] by introducing the phase space and the Lagrangian subvariety related to a hyperbolic tetrahedron. However, already for $\mathrm{SL}(2)$, an attempt to understand even the Lagrangian subvariety itself as a

moduli space of flat connections immediately faces a serious problem:

Any flat connection on a tetrahedron is trivial, so the corresponding moduli space is just a point, and thus cannot produce a Lagrangian subvariety in the phase space assigned to the boundary, which in this case is a four-holed sphere.

The problem is that the tetrahedron has trivial topology, while the moduli space of flat connections is a topological invariant, and hence also becomes trivial.

We suggest a solution to this problem based on the following idea. We consider moduli spaces of flat connections on 3-manifolds with *framings*. A framing amounts to introducing invariant flags on each of the so-called *small boundary components*, which we define below, invariant under the holonomy around the component. This, remarkably, allows one to produce the missing Lagrangian subvariety for the tetrahedron. The corresponding moduli spaces are defined for arbitrary admissible manifolds, and can be “symplectically” glued from the ones assigned to the tetrahedra. So we use the invariant flags to *localize* flat connections to tetrahedra.

This notion of framing generalizes the key idea used in [12] to introduce cluster coordinates on the moduli space $\mathcal{X}_K(\mathcal{C})$ of framed flat $\text{PGL}(K)$ -connections on a surface \mathcal{C} with punctures: invariant flags were invoked to localize flat connections on ideal triangles. In the three-dimensional case, a related notion of “decorations” was used by [25–27] to study representations of $\pi_1(M)$.

In section 2.1 we start with a careful discussion of a class of 3-manifolds with boundary, which we call *admissible 3-manifolds*, which are glued from truncated tetrahedra. Since the boundary faces of truncated tetrahedra are of two different types, triangles and hexagons, the boundary of the manifold obtained by gluing them along the hexagonal faces also has two kinds of boundary components, *big* (formed by the unglued hexagons) and *small* (formed by the triangles). We say that such a manifold is admissible if 1) the fundamental groups of the small boundary components are abelian; 2) the small boundary components are not spheres; and 3) the big boundary components have negative Euler character.

The second and third conditions are technical, and simply allow us to avoid stacks in our constructions. The first condition, however, is dictated by the notion of framing: *every* vector bundle with flat connection on an admissible 3-manifold admits at least one choice of framing. Indeed, if the fundamental group of the small boundary is abelian, then the holonomies of a flat connection on the small boundary all commute with each other, and a family of commuting operators in a vector space V_K always has an invariant flag. (Typically there are just $K!$ invariant flags; in a basis in which all operators are diagonal with different eigenvalues, choosing a flag amounts to ordering the basis.) So by adding a framing to a flat vector bundle we enlarge the moduli space by taking its cover and partially resolving its singularities, rather than cutting it down.

In section 2.3 we discuss the moduli spaces presented in (1.2): the space $\tilde{\mathcal{L}}_K(M)$ of framed flat $\text{PGL}(K, \mathbb{C})$ -connections on M , the space $\mathcal{X}_K^{\text{un}}(\partial M)$ of framed unipotent flat $\text{PGL}(K, \mathbb{C})$ -connections on the boundary, and the prospective Lagrangian $\mathcal{L}_K(M) = \text{im}(\tilde{\mathcal{L}}_K(M) \rightarrow \mathcal{X}_K^{\text{un}}(\partial M))$.

The simplest example of an admissible 3-manifold is the tetrahedron Δ , whose boundary is understood as a sphere with four punctures. We associate to the boundary the moduli space $\mathcal{X}_2^{\text{un}}(\partial\Delta)$ of framed flat $\text{PGL}(2, \mathbb{C})$ -connections on the four-punctured sphere with unipotent monodromy around the punctures. It is a two-dimensional symplectic space. Its Lagrangian subspace $\mathcal{L}_2(\Delta)$ consists of the connections that can be extended to the bulk with framings at the four vertices at the boundary. Since any connection on the ball is trivial, the only data left is the four flags, which in this case amounts to a configuration of four lines in a two-dimensional space V_2 . The resulting Lagrangian pair

$$\mathcal{L}_2(\Delta) \subset \mathcal{X}_2^{\text{un}}(\partial\Delta) \tag{1.14}$$

is our main building block. It has a natural Zariski open part, which deserves a special notation $\mathcal{L}_\diamond \subset \mathcal{P}_{\partial\diamond}$, described in coordinates as follows:

$$\mathcal{P}_{\partial\diamond} := \{z, z', z'' \in \mathbb{C}^* \mid zz'z'' = -1\}, \quad \mathcal{L}_\diamond := \{z, z', z'' \in \mathcal{P}_{\partial\diamond} \mid z'' + z^{-1} - 1 = 0\}. \tag{1.15}$$

The natural compactification of the symplectic space $\mathcal{P}_{\partial\diamond} = \mathbb{C}^* \times \mathbb{C}^*$ given as a moduli space (i.e. a stack) $\mathcal{X}_2^{\text{un}}(\partial\Delta)$ should help to deal with non-generic framings.

Framed flat connections from octahedra. Our first major goal is to build the moduli space $\tilde{\mathcal{L}}_K(M)$ of framed flat $\text{PGL}(K, \mathbb{C})$ -connections on an admissible 3-manifold M out of these building blocks. To achieve this, we choose an ideal triangulation \mathbf{t}_{3d} of M , and a further *hypersimplicial K -decomposition* of each tetrahedron in \mathbf{t}_{3d} (figure 2). We show in section 3 that this K -decomposition has precisely the combinatorial and geometric data that we need to describe $\tilde{\mathcal{L}}_K(M)$ and its projection $\mathcal{L}_K(M)$ to the boundary moduli space.

We use the framing data on a flat connection to assign to each tetrahedron in the triangulation \mathbf{t}_{3d} a *configuration* of four flags at its vertices. One can think of these as four flags in the K -dimensional space V_K , considered modulo the action of $\text{PGL}(K, \mathbb{C})$.⁸ The K -decomposition of tetrahedra is used to construct various generalized cross-ratios that determine the configuration of four flags. These generalized cross-ratios correspond precisely to the parameters z_i, z'_i, z''_i assigned to the vertices of each octahedron \diamond_i (as on the left of figure 2). Once the parameters are identified with cross-ratios, they naturally satisfy certain monomial relations of the form $c_v = 1$, stating that the product of all octahedron parameters sitting at any vertex v of the K -decomposition is trivial. These gluing constraints generalize Thurston’s gluing equations for an ideal hyperbolic triangulation.

These projective geometry constructions can be restated as follows. Let \mathbf{t}_{3d} be an ideal triangulation of M , inducing an ideal triangulation \mathbf{t} of the big boundary. We construct a Zariski-open subset

$$\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}) \simeq (\mathbb{C}^*)^{2d} \subset \mathcal{X}_K^{\text{un}}(\partial M), \quad 2d = \dim_{\mathbb{C}} \mathcal{X}_K^{\text{un}}(\partial M) \tag{1.16}$$

of the space of framed unipotent flat connections on ∂M , generalizing the cluster coordinate charts defined by [12]. We write the \mathbb{C}^* -coordinates of $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ as monomials of the

⁸It is important to notice that the sets of configurations of objects of any kind associated with a vector space depend only on the dimension of the space and not on the choice of vector space itself; thus configurations assigned to isomorphic vector spaces are *canonically* isomorphic.

octahedron parameters z_i, z'_i (where i ranges over the octahedra). This just means that we get a projection

$$\tilde{\pi} : \prod_{i=1}^N \mathcal{P}_{\partial\diamond_i} \simeq (\mathbb{C}^*)^{2N} \longrightarrow \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}) \simeq (\mathbb{C}^*)^{2d}, \quad (1.17)$$

where $i = 1, \dots, N$ indexes the octahedra in the K -decomposition of \mathbf{t}_{3d} . The projection is not quite canonical in the presence of small-torus boundary components, but canonical if they are absent.

We define an open subset of the space $\tilde{\mathcal{L}}(M)$ by intersecting the product of octahedron Lagrangians \mathcal{L}_{\diamond_i} with the gluing constraints

$$\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d}) = \left(\prod_{i=1}^N \mathcal{L}_{\diamond_i} \right) \cap \{c_v = 1\}_{\text{vertices } v}. \quad (1.18)$$

By showing that a set of octahedron parameters that satisfies the gluing equations can be used to uniquely reconstruct a framed flat connection on M (section 3.3), we prove that

Theorem 1.1 (Theorem 3.1 (page 44)). *The intersection (1.18) parameterizes a Zariski-open subset of the space of framed flat connections on M .*

The restriction of the map (1.17) to the subvariety $\{c_v = 1\}$ is a canonical map

$$\pi : \prod_{i=1}^N \mathcal{P}_{\partial\diamond_i} \cap \{c_v = 1\} \longrightarrow \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}). \quad (1.19)$$

Let $\mathcal{L}_K(M, \mathbf{t}_{3d}) := \tilde{\pi}(\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})) = \pi(\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d}))$ be the image of $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ under the projections $\tilde{\pi}$ or π :

$$\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}). \quad (1.20)$$

It is an open subset of the space $\mathcal{L}_K(M)$.

Theorem 3.1 implies that changing the bulk (\mathbf{t}_{3d}) and boundary (\mathbf{t}) triangulations amounts to birational transformations of the spaces $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$, $\mathcal{L}_K(M, \mathbf{t}_{3d})$ and $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$.

Changes of bulk triangulation are generated by 2–3 moves, which can be decomposed into elementary 2–3 moves acting on octahedra in a K -decomposition, as described in section 3.5. Changes of boundary triangulation correspond to cluster transformations on $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ [12]. Taking the union of the spaces assigned to a given bulk triangulation \mathbf{t}_{3d} over the set of the bulk triangulations compatible with a fixed boundary triangulation \mathbf{t} , we obtain a triple

$$\tilde{\mathcal{L}}_K(M, \mathbf{t}) \rightarrow \mathcal{L}_K(M, \mathbf{t}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}).$$

It depends only on \mathbf{t} . One may then vary \mathbf{t} to eliminate the dependence on boundary triangulation. We emphasize that even then we do not cover the whole moduli spaces $\tilde{\mathcal{L}}_K(M)$ and $\mathcal{X}_K^{\text{un}}(\partial M)$. In particular, only components of $\tilde{\mathcal{L}}_K(M)$ corresponding to irreducible flat connections on M will be detected.

Symplectic gluing. Our next major goal is to understand the symplectic properties of $\mathcal{L}_K(M)$ and $\mathcal{X}_K^{\text{un}}(\partial M)$. They are summarized in the Symplectic Gluing Conjecture (Conj. 4.1, page 59): for any admissible 3-manifold M with bulk triangulation \mathbf{t}_{3d} and corresponding big-boundary triangulation \mathbf{t} , we expect that

- The moduli space $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$, equipped with the canonical complex symplectic form, is isomorphic to a holomorphic symplectic quotient of the product of octahedron spaces $\prod_{i=1}^N \mathcal{P}_{\partial\phi_i}$, equipped with the product symplectic structure $\sum_{i=1}^N d \log z_i \wedge d \log z'_i$. Precisely, it is the symplectic quotient for the Hamiltonian $(\mathbb{C}^*)^{N-d}$ action whose Hamiltonians are given by the gluing monomials c_v . Thus

$$\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}) \simeq \left(\prod_{i=1}^N \mathcal{P}_{\partial\phi_i} \cap \{c_v = 1\} \right) / (\mathbb{C}^*)^{N-d}. \tag{1.21}$$

- $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ is a Lagrangian subvariety; it coincides with the image of the product Lagrangian $\prod_{i=1}^N \mathcal{L}_{\phi_i} \subset \prod_{i=1}^N \mathcal{P}_{\partial\phi_i}$ under the symplectic quotient (1.21).

Recall the projection $\tilde{\pi} : \prod_{i=1}^N \mathcal{P}_{\partial\phi_i} \rightarrow \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ from (1.17). We prove that the gluing monomials c_v Poisson commute, and that they Poisson commute with the pullbacks by the map $\tilde{\pi}^*$ of the cluster coordinates on $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$, given by the monomials in the octahedron parameters. Thus the first claim (1.21) of the Symplectic Gluing Conjecture reduces to the claim that exactly $N - d$ of the equations $c_v = 1$ are independent. An easy Euler-characteristic count shows that in the presence of t small-torus boundaries, the total number of gluing monomials c_v equals $N - d + (K - 1)t$; therefore, it remains to show that there are exactly $(K - 1)t$ relations among the gluing monomials.

Also recall the map π from (1.19). The claim that $(\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}), \Omega)$, as a holomorphic symplectic space, is the reduction of the product of octahedron spaces means that

$$\pi^* \Omega = \sum_{i=1}^N d \log z_i \wedge d \log z'_i \Big|_{c_v=1}. \tag{1.22}$$

Since the number of monomials c_v is $N - d + (K - 1)t$, the dimension of $\tilde{\mathcal{L}}(M; \mathbf{t}_{3d})$ is at least $d - (K - 1)t$, see (1.18). The second claim of the Symplectic Gluing Conjecture is that the image of $\tilde{\mathcal{L}}(M; \mathbf{t}_{3d})$ under the projection by the Hamiltonian flows of the Hamiltonians c_v is Lagrangian. When the only small-boundary components of ∂M are discs, theorem 4.2 guarantees that the image of $\tilde{\mathcal{L}}(M; \mathbf{t}_{3d})$ is isotropic; thus it would suffice to show that its dimension is $\geq d$.

We refer to the relationship between the pair $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ and the elementary octahedron pairs $\mathcal{L}_{\phi_i} \subset \mathcal{P}_{\partial\phi_i}$ as *symplectic gluing*. Since our parameters are assigned to the octahedra of the K -decomposition of M , the Symplectic Gluing Conjecture says that

The Lagrangian pair $\mathcal{L}_K(M) \subset \mathcal{X}_K^{\text{un}}(\partial M)$ is obtained by symplectic gluing of the elementary Lagrangian pairs (1.14) parametrized by the octahedra of the K -decomposition of M corresponding to an ideal triangulation \mathbf{t}_{3d} of M .

The symplectic gluing is studied in sections 4 and 5 from two different perspectives. In both settings, it is very natural to promote the notion of a Lagrangian subvariety to the much stronger notion of a K_2 -Lagrangian. Let us explain what this means.

K₂-Lagrangians. Let $F^* := F - \{0\}$ be the multiplicative group of a field F . Recall that $F^* \wedge F^* = \Lambda^2 F^*$ is the abelian group generated by elements of the form $a \wedge b$, $a, b \in F^*$, with $a \wedge b = -b \wedge a$ and $(a_1 a_2) \wedge b = a_1 \wedge b + a_2 \wedge b$. The group $K_2(F)$ is the quotient of $F^* \wedge F^*$ by the subgroup generated by Steinberg relations $(1 - z) \wedge z$, where $z \in F^* - \{1\}$:

$$K_2(F) = F^* \wedge F^* / \{(1 - z) \wedge z\}. \tag{1.23}$$

Next, let X be a complex algebraic variety. Denote by $\mathbb{C}(X)$ the field of rational functions on X . Then there is a homomorphism to the space $\Omega_{\log}^2(X)$ of holomorphic 2-forms with logarithmic singularities on X :

$$d \log : \mathbb{C}(X)^* \wedge \mathbb{C}(X)^* \longrightarrow \Omega_{\log}^2(X), \quad f \wedge g \longmapsto d \log f \wedge d \log g. \tag{1.24}$$

The map kills elements $(1 - f) \wedge f$, and so the image of an element $W \in \mathbb{C}(X)^* \wedge \mathbb{C}(X)^*$ depends only on its class in $K_2(\mathbb{C}(X))$.

It was proved in [12] that the symplectic form on the moduli space $\mathcal{X}_K^{\text{un}}(\mathcal{C})$ of unipotent flat connections on a 2d surface \mathcal{C} can be upgraded to its motivic avatar, a class \mathcal{W} in K_2 of $\mathcal{X}_K^{\text{un}}(\mathcal{C})$. The symplectic form is recovered as $d \log(\mathcal{W})$. From our current perspective, the construction of [12] applies directly to the big boundaries of admissible 3-manifolds, and we explain the simple generalization to small boundaries in sections 4.3. (An even wider class of examples is provided by cluster \mathcal{A} -varieties in [17].) This motivates the following definition.

Definition 1.1. *Let X be a complex variety with a class \mathcal{W} in $K_2(\mathbb{C}(X))$ such that $d \log(\mathcal{W})$ is a symplectic form at the generic part of X . A subvariety $L \subset X$ is called a K_2 -Lagrangian subvariety if \mathcal{W} restricts to zero in $K_2(\mathbb{C}(L))$, and $2 \dim L = \dim X$.*

Examples: 1. The space $\mathcal{P}_{\partial \diamond}$ has the symplectic form $d \log z \wedge d \log z'$. It lifts to a symbol $z \wedge z'$, which is invariant, up to 2-torsion, under the cyclic shift $z, z', z'' \mapsto z', z'', z$. The curve \mathcal{L}_{\diamond} is a K_2 -Lagrangian, since $z \wedge z'$ restricts to $(1 - z) \wedge z$ on \mathcal{L}_{\diamond} .

2. The graph of any cluster transformation $\mathcal{A} \rightarrow \mathcal{A}$ of a cluster \mathcal{A} -variety \mathcal{A} is a K_2 -Lagrangian subvariety of the product $\mathcal{A} \times \mathcal{A}$, see section 6 of [17].

Theorem 1.2 (Theorem 4.2).

(i) *If ∂M has only big boundary and small discs, any generic component of $\mathcal{L}_K(M)$ is a K_2 -isotropic subspace of $\mathcal{X}_K^{\text{un}}(\partial M)$, i.e. the restriction of the K_2 -class \mathcal{W} is zero.*

(ii) *If M is a convex polyhedron, then $\mathcal{L}_K(M)$ is a K_2 -Lagrangian subvariety of $\mathcal{X}_K^{\text{un}}(\partial M)$.*

Part (ii) of theorem 4.2 follows from part (i) and an easy dimension count. Indeed, $\mathcal{L}_K(M)$ is just the configuration space of flags parametrized by vertices of the polyhedron M :

$$\mathcal{L}_K(M) = G \backslash \mathcal{B}^{\{\text{vertices of } M\}}, \quad \dim \mathcal{L}_K(M) = v(M) \cdot \dim \mathcal{B} - \dim G.$$

Here $G = PGL_K$ and $v(M)$ is the number of vertices of M . On the other hand,

$$\dim \mathcal{X}_K^{\text{un}}(M) = \dim(G)(v(M) - 2) - \text{rk}(G)v(M) = 2 \dim \mathcal{L}_K(M).$$

Since the cluster coordinates on $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ and the gluing constraints c_v are monomials of the octahedron parameters z_i, z'_i , a proof of Conjecture 4.1 implies a K_2 -analog of (1.22):

$$\pi^* \mathcal{W} = \sum_{i=1}^N z_i \wedge z'_i \Big|_{c_v=1} \pmod{\text{torsion}}. \tag{1.25}$$

It would immediately imply that $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ is K_2 -isotropic.

In section 4 we prove theorem 4.2 by using the canonical map of complexes [66]:

$$\alpha_\bullet : \text{Complex of generic configurations of decorated flags in } V_K \longrightarrow \text{Bloch complex}. \tag{1.26}$$

To define it, we define first a closely related homomorphism of complexes, where the notation will be explained momentarily:

$$\begin{array}{ccccccc} \dots & \xrightarrow{d} & A_4^{(K)} & \xrightarrow{d} & A_3^{(K)} & \xrightarrow{d} & A_2^{(K)} \\ & & \downarrow \gamma_5 & & \downarrow \gamma_4 & & \downarrow \gamma_3 \\ \dots & \longrightarrow & R_2(F) & \xrightarrow{i} & \mathbb{Z}[F^* - \{1\}] & \xrightarrow{\delta} & \wedge^2(F^*) \end{array} \tag{1.27}$$

Here the abelian group $A_n^{(K)}$ is given by formal integral linear combinations of the configurations of $n + 1$ generic decorated flags in V_K . The differential d is the standard simplicial differential. This way we get the complex of generic configurations of decorated flags in V_K .

Let us define the bottom complex. We denote by $\mathbb{Z}[F^* - \{1\}]$ the free abelian group with a basis $\{z\}$, where z runs through the elements of $F^* - \{1\}$. The homomorphism δ in (1.27) is defined by setting $\delta\{z\} := (1 - z) \wedge z$. So its cokernel is the group $K_2(F)$ in (1.23).

Let $R_2(F) \subset \mathbb{Z}[F^* - \{1\}]$ be the subgroup generated by the five-term relations

$$\sum_{i=0}^4 (-1)^i \{r(z_0, \dots, \widehat{z}_i, \dots, z_4)\} \subset \mathbb{Z}[F^* - \{1\}], \tag{1.28}$$

where (z_0, \dots, z_4) run through generic configurations of 5 points on $\mathbb{P}^1(F)$, and $r(-)$ is the cross-ratio. One shows that the restriction of the map δ to the subgroup $R_2(F)$ is zero. So we get the bottom complex, where i is the natural embedding.

The Bloch group $B_2(F)$ is the quotient of the group $\mathbb{Z}[F^* - \{1\}]$ by the five-term relations:

$$B_2(F) := \frac{\mathbb{Z}[F^* - \{1\}]}{R_2(F)}.$$

Since $\delta(R_2(F)) = 0$, the map δ descends to the Bloch group, and we get the Bloch complex:

$$B_2(F) \xrightarrow{\delta} \wedge^2(F^*). \tag{1.29}$$

So the map of complexes (1.27) induces a map of complexes (1.26).

Let us assume now that $F = \mathbb{C}$. Recall the Bloch-Wigner version of the dilogarithm

$$\mathcal{L}i_2(z) := \text{Im}(\text{Li}_2(z) + \log(1 - z) \log |z|). \tag{1.30}$$

It is a single-valued function, well defined for all $z \in \mathbb{C}$. So it defines a group homomorphism

$$\mathcal{L}i_2 : \mathbb{Z}[(\mathbb{C})] \longrightarrow \mathbb{R}, \quad \{z\}_2 \rightarrow \mathcal{L}i_2(z).$$

It satisfies the Abel five-term relation, i.e. its restriction to the subgroup $R_2(\mathbb{C})$ is zero. Therefore it gives rise to a group homomorphism

$$\mathcal{L}i_2 : B_2(\mathbb{C}) \longrightarrow \mathbb{R}, \quad \{z\}_2 \rightarrow \mathcal{L}i_2(z). \tag{1.31}$$

The Bloch group is famously related to 3d hyperbolic geometry, in the following way. Consider the scissor congruence group \mathcal{P}_2 of ideal hyperbolic polyhedra. It is an abelian group with the generators assigned to ideal oriented hyperbolic tetrahedra $I(z_1, z_2, z_3, z_4)$ with the vertices z_1, \dots, z_4 . The generators $[I(z_1, z_2, z_3, z_4)]$ satisfy two kinds of relations. First, the cutting and gluing relation: cutting an ideal hyperbolic bipyramid in two different ways into 2 or 3 ideal tetrahedra amounts to the same element of \mathcal{P}_2 . Second, changing the orientation of a tetrahedron amounts to a sign change: $[I(0, 1, \infty, z)] = -[I(0, 1, \infty, \bar{z})]$. Denote by $B_2(\mathbb{C})^-$ the aninvariants of the action of complex conjugation on the group $B_2(\mathbb{C})$. Then there is a canonical group isomorphism

$$B_2(\mathbb{C})^- \longrightarrow \mathcal{P}_2, \quad \{z\}_2 \rightarrow [I(0, 1, \infty, z)].$$

Finally, the map δ shows up in the formula for the differential of the dilogarithm:

$$d\mathcal{L}i_2(z) = \log|1-z|d\arg(z) - \log|z|d\arg(1-z). \tag{1.32}$$

Let us return finally to the homomorphism of complexes (1.27). To define it, one uses a key construction of [66] relating configurations of flags to configurations of vectors:

$$\text{Complex of generic configurations of decorated flags in } V_K \longrightarrow \text{biGrassmannian complex}. \tag{1.33}$$

Given a single generic configuration of $(n+1)$ decorated flags in V_K , one assigns to it a collection of points in the Grassmannians \mathbf{G}_p^q where $p+q+2 = n+1$. The K -decomposition of simplices is crucial here: viewing the $(n+1)$ flags as assigned to the vertices of an n -dimensional simplex Δ , each hypersimplex in the K -decomposition of Δ gives rise to a single point in a Grassmannian that matches the type of the hypersimplex. Combining (1.33) with the homomorphism of complexes, defined in [67], which we review in section 4.2.3:

$$\text{biGrassmannian complex} \longrightarrow \text{Bloch complex}, \tag{1.34}$$

we arrive at the homomorphism of complexes (1.27), and hence (1.26).

The homomorphism of complexes (1.26) controls a number of features of the geometry of framed flat connections on an admissible three-manifold M and its boundary. Let us elaborate on this.

First, choosing compatible triangulations \mathbf{t}_{3d} and \mathbf{t} of M and its big boundary, we assign to points in the moduli spaces $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ and $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ elements in the complex of generic configurations of decorated flags in V_K . For example, we may start from a generic

point of $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$, representing a framed flat connection on a triangulated manifold M . We pick a decorated flag representative for each flag of the framing: we can do this thanks to the unipotence condition. Then we assign to each tetrahedron of the triangulation the configuration of four decorated flags obtained by the restriction of the framed connection to the tetrahedron. The formal sum of these configurations over all tetrahedra of the triangulation is an element of the group $A_3^{(K)}$ which we assign to the framed flat connection. Similarly, we assign to a generic point of $\mathcal{X}_K^{\text{un}}(M, \mathbf{t})$ an element of $A_2^{(K)}$. Then we find:

- 1) The component γ_3 of the map (1.27) was used in [12, section 15] to define the K_2 -class \mathcal{W} on the space $\mathcal{X}_K^{\text{un}}(\mathcal{C})$ for a 2d surface \mathcal{C} . It provides a K_2 -class on the space $\mathcal{X}_K^{\text{un}}(\partial M)$.
- 2) Theorem 4.2 tells that the restriction of the K_2 -class \mathcal{W} on $\mathcal{X}_K^{\text{un}}(\partial M)$ to $\mathcal{L}_K(M)$ is zero. The component γ_4 of map (1.27) tells how exactly it becomes zero. Precisely, given an ideal triangulation \mathbf{t} of ∂M , the class \mathcal{W} has a natural lift to $\Lambda^2 \mathbb{C}^*$ [12, section 15]. The map γ_4 presents its restriction to $\mathcal{L}_K(M)$ as a sum of Steinberg relations $(1 - z_i) \wedge z_i$. These z_i 's are just our octahedron parameters. They are precisely the terms of the map γ_4 in (1.27).

The very existence of the map of complexes (1.27) implies that the image $\sum \{z_i\}_2$ of the map γ_4 in the Bloch group $B_2(\mathbb{C})$ is independent of the choice of the balk triangulation \mathbf{t}_{3d} . We call it the *motivic volume*. Applying the dilogarithm homomorphism (1.31) to it we define in section 4.4 the volume of a generic framed flat $\text{PGL}(K, \mathbb{C})$ -connection on M . We stress that although the motivic volume of a framed flat connection does not depend on the balk triangulation \mathbf{t}_{3d} , it does depend on the boundary triangulation \mathbf{t} .

Using formula (1.32) and the commutativity of the last square in (1.27), we get a formula for the variation of the volume of generic framed flat connections on M . It generalizes the Neumann-Zagier formula for variation of volumes of hyperbolic 3-manifolds with toric boundary [8], and the work of Bonahon [68] on hyperbolic 3-manifolds with geodesic boundary. The $\text{PGL}(3, \mathbb{C})$ -analog of the variation formula was beautifully established by Bergeron-Falbel-Guilloux in [23].

If the big boundary is absent, the motivic volume lies in the kernel of the Bloch complex map (1.29), and thus defines an element of $K_3^{\text{ind}}(\mathbb{C}) \otimes \mathbb{Q}$ due to a theorem of Suslin [69]. This follows immediately from the construction and the fact that the map (1.26) is a map of complexes. The value of the regulator on it is the Chern-Simons invariant of the connection.

- 3) What exactly happens under the 2-3 moves changing the triangulation \mathbf{t}_{3d} ? In section 3.5 we prove that the $\text{PGL}(K)$ pentagon relation can be reduced to a sequence of 2-3 Pachner moves in the $K = 2$ case, which we call elementary pentagon relations. The elementary pentagon relations match the terms of the component α_5 of the map (1.26).

The 2–3 move on a 3-manifold can be viewed as a cobordism that amounts to attaching a 4-simplex to the 3-manifold. We believe that the K -decomposition of this 4-simplex should play an important role in construction of invariants of 4-manifolds. The fact that the elementary pentagons which appear in the decomposition of the 2–3 $\mathrm{PGL}(K)$ -move match the hypersimplices of type $\Delta^{1,2}$ and $\Delta^{2,1}$ of the K -decomposition of the 4-simplex, which are the 4d analogs of the octahedron, agrees nicely with this idea.

Combinatorics. In section 5, we generalize to $K > 2$ a combinatorial analysis of ideal triangulations (and their symplectic properties) pioneered by Neumann-Zagier [8] and extended by Neumann [70] for $K = 2$. We introduce an algebra of paths on slices in the K -decomposition of a triangulation of an admissible 3-manifold M . These paths provide a simple, geometric interpretation of the monomial map relating coordinates on $\mathcal{X}_K(\partial M, \mathbf{t})$, gluing constraints c_v , and octahedron parameters z_i, z'_i, z''_i .

We use the geometry of paths to prove a weak version of the Symplectic Gluing Conjecture, essentially saying that the symplectic quotient on the r.h.s. of (1.21) makes sense and contains $\mathcal{L}_K(\partial M, \mathbf{t})$ as a subspace (Proposition 5.1).

The geometry of paths in section 5 have another application: they allow us to consistently define logarithmic lifts of octahedron parameters and coordinates on $\mathcal{X}_K^{\mathrm{un}}(\partial M, \mathbf{t})$ — thus lifting the symplectic quotient (1.21) to universal covers. As initially discussed in [9], this is a necessary requirement for quantization of the pair $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\mathrm{un}}(\partial M, \mathbf{t})$. It is also a necessary requirement for the definition of gauge theories $T_K[M]$ in later sections of the paper. This lift was introduced by Neumann in [70] in order to calculate the Chern-Simons invariant of a flat $\mathrm{PGL}(2, \mathbb{C})$ -connection.

Related mathematical works. There are several recent papers closely related to our work. Bergeron-Falbel-Guilloux [23] were the first to study parametrizations of spaces of flat connections on 3-manifolds for $K > 2$, considering $\mathrm{PGL}(3, \mathbb{C})$ -connections. The authors localized flat connections to tetrahedra, defined the equivalent of our octahedron parameters, and generalized the classic formulas for variations of the volume. While this paper was in preparation, Garoufalidis, Goerner, and Zickert [24] (following [25–27]) explained how to localize flat $\mathrm{PGL}(K, \mathbb{C})$ -connections to tetrahedra by considering “decorated representations” of $\pi_1(M)$. The “decorations” of [24] are identical to our notion of “framing,” and lead to the same parameters z_i, z'_i, z''_i and gluing equations for the space $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ of framed flat connections on M that we discuss. A major difference between [24] and the present work includes our simultaneous study of both bulk and boundary moduli spaces $\mathcal{L}_K(M) \subset \mathcal{X}_K^{\mathrm{un}}(\partial M)$ as K_2 -Lagrangian pairs obtained from symplectic gluing (which necessitates the introduction of admissible manifolds, with big and small boundaries). This allows us to easily obtain (e.g.) the formula for variation of the volume. Another major difference is our use of geometric K -decompositions.

Since the first version of this paper, a proof of an important result for manifolds with boundary consisting entirely of small tori, such as knot complements, generalizing the work of Neuman [70] for $K = 2$ has been proposed by [71]. It implies the Symplectic

Gluing Conjecture. The proof is extremely technical. It would be very satisfying to find a simpler, more fundamental proof. For the case $K = 2$, a simple topological perspective on symplectic gluing was given in [72].

Methods for quantizing the pair $\mathcal{L}_K(M) \subset \mathcal{X}_K^{\text{un}}(\partial M)$ have been significantly improved since the first version of this paper. In particular, it was realized that in order to quantize $\mathcal{L}_K(M, \mathfrak{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathfrak{t})$ for a fixed triangulation \mathfrak{t}_{3d} — to obtain both a Chern-Simons wavefunction and a set of difference operators that annihilate it — it is necessary for the triangulation \mathfrak{t}_{3d} to admit a *positive angle structure* [55, 59–62, 65]. This condition is also necessary physically in order for the gauge theory $T_K[M]$ to flow to a well-defined conformal field theory in the infrared. The positive angle structure seems to provide an intrinsically three-dimensional substitute for the notion of positivity in the study of cluster varieties such as $\mathcal{X}_K(\partial M, \mathfrak{t})$. This relationship will be explored elsewhere.

1.3 Organization

To recap, the paper is organized as follows.

In section 2 we introduce the notion of a triangulated admissible 3-manifold (section 2.1), different moduli spaces with framed flat connections (section 2.3), and explain how the framing allows a localization of framed flat connections to configurations of flags (section 2.4). We then define the hypersimplicial K -decomposition (section 2.5).

In section 3 we analyze how octahedron parameters are used to describe framed flat connections on an admissible M and its boundary, and begin to discuss the symplectic properties of the moduli spaces $\mathcal{L}_K(M) \subset \mathcal{X}_K^{\text{un}}(\partial M)$. We review how the K -decompositions of the 2d simplices, called K -triangulations, were used in [12] to define cluster coordinates on $\mathcal{X}_K^{\text{un}}(\partial M)$. Then we revisit the octahedron parameters related to a K -decomposition of M , and show (theorem 3.1) how the elementary symplectic pairs $\mathcal{L}_\diamond \subset \mathcal{P}_{\partial \diamond}$ for octahedra are glued together to construct a Zariski open part for the pair $\mathcal{L}_K(M) \subset \mathcal{X}_K^{\text{un}}(\partial M)$. We also study 2–3 moves, and decompose them into sequences of elementary ones.

In section 4, we formulate the Symplectic Gluing Conjecture (section 4.1). Then we introduce all the ingredients necessary to understand the homomorphism of complexes (1.26). This allows us to prove the main new result, theorem 4.2, telling that $\mathcal{L}_K(\mathcal{M})$ is K_2 -isotropic in the case that ∂M consists entirely of big boundary, with holes filled in by small discs. We discuss how to generalize this to arbitrary admissible M .

In section 5 we follow the combinatorial approach to symplectic gluing. We discuss the combinatorics of octahedron parameters, the Poisson brackets that they induce on $\mathcal{X}_K(\partial M)$, and the abstract data needed for quantization. We defer to appendix B the proof of the most nontrivial result about Poisson brackets for eigenvalue coordinates on $\mathcal{X}_K(\partial M)$.

In sections 6 and 7 we return to the physical motivations of this paper. We use the combinatorial data of section 5 to construct simple theories $T_K[M]$ associated with polyhedra and show how 2–3 moves encode mirror symmetries. We then discuss general properties of theories $T_K[M]$ when M is a knot or link complement, including the large- K scaling behavior. (At the end of section 7 we provide a few examples of the moduli spaces




Figure 3. (a) truncated tetrahedra, glued together to form 3-manifolds that have either (b) big boundaries tiled by hexagonal faces with holes filled in by small vertex triangles; or (c) small torus boundaries tiled exclusively by vertex triangles.

associated with simple knot complements, and how their coordinates are computed from paths in the K -decomposition.)

The complete dictionary between the symplectic data of section 5, class- \mathcal{R} gauge theories, and other quantum invariants is reviewed in appendix A.

2 Basic tools and definitions

We carefully introduce the basic objects studied throughout the paper.

2.1 Gluing admissible 3-manifolds from truncated tetrahedra

Let us truncate a tetrahedron by cutting off its vertices. The resulting *truncated tetrahedron* has two kinds of faces: (small) triangles replacing the original vertices, and (big) hexagons. Let us glue truncated tetrahedra into an oriented manifold with boundary by allowing pairs of hexagonal faces to be glued, but not triangular ones (figure 3). The boundary of the resulting manifold M is tiled by triangles and unglued hexagons. We call the part tiled by triangles the *small boundary*, and the part tiled by hexagons the *big boundary*.

A small boundary component could have topology of any oriented surface. However we impose two additional constraints. We say that a gluing is *admissible* if it satisfies the following topological conditions:

- The small boundary components have abelian fundamental groups.
- The small boundary components are not spheres.
- The big boundary components have negative Euler character.

The reason for the first condition is to ensure that any flat connection admits a framing (section 2.3). The second condition is ultimately necessary for the existence of logarithmic path coordinates in section 5 and subsequent quantization and definition of $T_K[M]$, though it could be relaxed classically. Mathematically, the second and third conditions eliminate stacks that should be assigned to spheres with less than three punctures.

A 3-manifold with boundary is *admissible* if it is homeomorphic to a manifold with boundary obtained by an admissible gluing of truncated tetrahedra. We consider only

admissible manifolds M . An ideal triangulation of an admissible manifold is an admissible tiling of the manifold by truncated tetrahedra.

A small boundary component must be a *closed torus*, an *annulus*, or a *disc*. A big boundary component is a surface with $h \geq 1$ holes (see figure 3b), provided by the edges of the truncated tetrahedra shared by triangles and hexagons. The hexagonal tiling is equivalent to a 2d ideal triangulation of this surface — shrinking small discs to points we get an ideal triangulation. The small and big boundary components are connected as follows.

- The small tori are closed, disjoint boundary components.
- Boundary circles of small annuli must be glued to boundary circles of big surfaces.
- The boundary circle of a small disc is glued to a boundary circle of a big boundary surface. So each disc fills a hole of a big boundary.

Here are some useful examples of admissible 3-manifolds to keep in mind:

1. The most basic example is given by a tetrahedron. Its big boundary is a four-holed sphere, and each hole is filled by a small boundary disc. More generally, any convex polyhedron gives rise to an admissible 3-manifold. Indeed, consider a triangulation of the polyhedron into tetrahedra. Its big boundary is a sphere with the holes matching the vertices of the polyhedron.
2. Another fundamental example is a hyperbolic 3-manifold, triangulated into ideal hyperbolic tetrahedra. The big boundary is then the geodesic boundary. It is a union of geodesic surfaces. The small torus boundaries are cusps or deformed cusps. The cusps can be regularized by taking horoball neighborhoods of all the ideal hyperbolic tetrahedra vertices ending on them. This truncates the vertices of the ideal hyperbolic tetrahedra. In each tetrahedron, the truncated-vertex triangles are Euclidean. Moving from tetrahedron to tetrahedron going around the cusp, we sweep out a Euclidean surface. Metrically, the surface closes up into a Euclidean torus if the cusp holonomy is parabolic. Otherwise it keeps spiraling.
3. Our last example is given by a link complement. It has only small torus boundaries, one for every excised link component.

Topologically, we can describe the boundary ∂M of an admissible 3-manifold in terms of laminations. A *simple lamination* on a surface \mathcal{C} with marked points $\{p_1, p_2, \dots\}$ is a collection of non-intersecting simple unoriented non-isotopic loops on $\mathcal{C} \setminus \{p_1, p_2, \dots\}$ modulo isotopy. The boundary ∂M is a disjoint union of tori and closed oriented surfaces \mathcal{C} with simple laminations and marked points. The small annuli in ∂M are neighborhoods of the loops of the laminations, and the small discs are neighborhoods of the marked points (figure 4). The remainder of ∂M is big boundary.

There is also a nice physical interpretation of the admissibility conditions. We ultimately want to study 3-manifolds M on which the 6d $(2, 0)$ theory can be compactified. The big boundaries we have just described are asymptotic regions of M in six dimensions (i.e.




Figure 4. A topological surface with lamination and marked points is equivalent to a big boundary with small annuli and small discs connecting its holes. Here, two loops of a lamination are replaced by two annuli (depicted as long tubes); and two marked points are replaced by discs. The remaining big boundary consists of two 3-punctured spheres.

of $M \times \mathbb{R}^3$), of the form $\mathcal{C} \times \mathbb{R}_+$. The small boundaries come from regularizing codimension-two defects of the 6d theory. The defects may either be closed loops in M (hence small torus boundaries), infinite lines connecting two asymptotic regions (hence small annuli), or in special cases *half*-infinite lines attached to one asymptotic region (hence small discs). This is explained further in [15]. Note that it is impossible to obtain small spheres — there are no codimension-three defects in the 6d theory, which could appear point-like on M .

2.2 Configurations of flags

To define the correct phase space assigned to the boundary ∂M of an admissible 3-manifold M , and especially to define its Lagrangian subvariety corresponding to M , we start with a small digression on flags and configurations of flags.

A flag F_\bullet in a K -dimensional complex vector space V_K is a collection of nested subspaces

$$F_\bullet : \quad \emptyset = F_0 \subset F_1 \subset F_2 \subset \dots \subset F_K = V_K, \quad \dim F_i = i. \quad (2.1)$$

We also use the notation $F^i := F_{K-i}$ for the codimension i subspace of the flag. Then a flag is denoted

$$F^\bullet : \quad \emptyset = F^K \subset F^{K-1} \subset F^{K-2} \subset \dots \subset F^0 = V_K, \quad \text{codim } F^i = i. \quad (2.2)$$

In fact later on we will mostly use the second convention. We say that a basis $\{f_1, f_2, \dots, f_K\}$ in V is adjusted for a flag F_\bullet if F_i is spanned by $\{f_1, f_2, \dots, f_i\}$.

A flag F_\bullet in V_K induces a canonical dual flag F_\bullet^* in the dual vector space V_K^* :

$$F_\bullet^* : \quad \emptyset = F_0^* \subset F_1^* \subset \dots \subset F_K^* = V_K^*, \quad F_i^* := (V_K/F^i)^*. \quad (2.3)$$

Indeed, the projection $V_K/F^{i+1} \rightarrow V_K/F^i$ induces a dual injection $(V_K/F^i)^* \hookrightarrow (V_K/F^{i+1})^*$.

A collection of m flags $(F_{1,\bullet}, F_{2,\bullet}, \dots, F_{m,\bullet})$ in V_K is said to be *generic* if for any collection of integers a_1, \dots, a_m that sum to K one has an isomorphism

$$F_{1,a_1} \oplus \dots \oplus F_{m,a_m} = V_K.$$

The space of flags is denoted by \mathcal{B} , the moduli space of configurations of m flags is denoted by $\text{Conf}_m(\mathcal{B})$, and the variety of generic configurations of m flags is denoted by

$\text{Conf}_m^*(\mathcal{B})$. Recall that if we have a set X with an action of a group G , the *configurations* of n elements of X are the orbits of G acting on X^n . In this case, a configuration of flags consists of the $\text{PGL}(V_K)$ -orbit of a tuple of flags. For $\text{PGL}(2, \mathbb{C})$, the space of flags is \mathbb{CP}^1 . The cross-ratio provides an isomorphism

$$\text{Conf}_4^*(\mathbb{CP}^1) \xrightarrow{=} \mathbb{CP}^1 - \{0, 1, \infty\}. \tag{2.4}$$

2.3 The basic moduli spaces

A vector bundle \mathcal{V} with a flat connection gives rise to a new bundle with a flat connection on the same space, the flag bundle $\mathcal{V}_{\mathcal{B}}$, whose fiber at a point x is the space of all flags at the fiber of \mathcal{V} at x .

Definition 2.1. *Let M be an admissible 3-manifold. A framing on a vector bundle \mathcal{V} with a flat connection on M is a choice of a flat section of the flag bundle $\mathcal{V}_{\mathcal{B}}$ on each of the small components of ∂M .*

Let ∂M^ be the total boundary ∂M with all small discs removed. A framing on a vector bundle \mathcal{V} with flat connection on ∂M^* is a choice of flat section of $\mathcal{V}_{\mathcal{B}}$ on each remaining small component of ∂M^* and on each S^1 boundary of ∂M^* .*

To define a framing we need to define it for each of the small boundary components. A framing at a boundary component can be thought of as an invariant flag, that is a flag in the fiber at a point of the component, preserved by the holonomies around the component.

A collection of commuting linear transformations in a finite-dimensional complex vector space always admits an invariant flag, i.e. a flag preserved by all the transformations. Therefore, since the fundamental group of each small boundary components is abelian, a framing is an additional structure on a vector bundle with connection — it always exists. This is why we consider only small boundary components with the abelian fundamental group. Otherwise the very existence of an invariant flag is a severe restriction on a flat connection.

If the holonomy around a puncture is unipotent, and consists of a single Jordan block, then there is a unique invariant flag there. Otherwise additional freedom arises. For example, if the holonomy is trivial, then the choice of invariant flag is completely arbitrary.

Using the notion of framing, we can define now the moduli spaces needed in the paper.

Definition 2.2.

- (i) *The moduli space $\mathcal{X}_K(\partial M)$ parametrizes framed $\text{PGL}(K, \mathbb{C})$ -bundles with flat connections on ∂M^* .*
- (ii) *The subspace $\mathcal{X}_K^{\text{un}}(\partial M) \subset \mathcal{X}_K(\partial M)$ parametrizes the framed flat connections on ∂M^* with unipotent holonomy around the holes (S^1 boundaries) where small discs were removed.*
- (iii) *The moduli space $\mathcal{L}_K(M)$ parametrizes framed flat $\text{PGL}(K, \mathbb{C})$ -connections on ∂M^* that can be extended to framed flat connections on M .*
- (iv) *The moduli space $\tilde{\mathcal{L}}_K(M)$ parametrizes framed flat $\text{PGL}(K, \mathbb{C})$ -connections on M .*

Any loop on ∂M^* around a hole is contractible in M . So the holonomy of a flat connection in M around any such loop in ∂M is trivial, and thus the invariant flags near every puncture are completely unrestricted. So $\mathcal{L}_K(M)$ embeds into the moduli space of connections with trivial holonomies around the holes. The moduli spaces are related as follows:

$$\mathcal{L}_K(M) \subset \mathcal{X}_K^{\text{un}}(\partial M) \subset \mathcal{X}_K(\partial M). \tag{2.5}$$

Moreover, $\mathcal{L}_K(M)$ is the image of the projection $\tilde{\mathcal{L}}_K(M) \rightarrow \mathcal{X}_K(\partial M)$.

The moduli space $\mathcal{X}_K(\partial M)$ has a canonical Poisson structure. The moduli space $\mathcal{X}_K^{\text{un}}(\partial M)$ is symplectic. It is realized as a closure of a symplectic leaf of $\mathcal{X}_K(\partial M)$. It serves the role of the phase space. The subspace $\mathcal{L}_K(M)$ is supposed to be a Lagrangian subspace.

Zariski open parts of these moduli spaces can be understood by introducing coordinates. However the coordinates are not everywhere defined. The moduli spaces themselves are of fundamental importance.

The moduli space $\mathcal{X}_K(\partial M)$ is naturally decomposed into a product of the moduli spaces assigned to the components of ∂M . To state this precisely, let us discuss the moduli spaces assigned to the small boundary components.

1. The phase space for a surface with a simple lamination. Let \mathcal{C} be an oriented surface with $h > 0$ holes and $\chi(\mathcal{C}) < 0$. Let γ be a simple lamination on \mathcal{C} . (Thus, we have in mind that \mathcal{C} is the union of the big boundary of M and the small annular boundaries that connect some pairs of holes; the small annuli are collapsed to lamination curves.) A framing on a flat vector bundle with connection on (\mathcal{C}, γ) is a choice of a flat section of the flag bundle over each component of the lamination γ , and near each of the holes. Then

$$\mathcal{X}_K(\mathcal{C}, \gamma) := \left\{ \text{framed flat } \text{PGL}(K, \mathbb{C})\text{-connections on } (\mathcal{C}, \gamma) \right\} / \text{isomorphisms}. \tag{2.6}$$

This is a finite-dimensional complex space. If g is the genus of \mathcal{C} , its dimension is $-\chi(\mathcal{C}) \dim(\text{PGL}(K))$. If we further fix the eigenvalues of the $\text{PGL}(K, \mathbb{C})$ holonomies around each of the holes, the dimension is cut down to

$$d_K(g, h) = -\chi(\mathcal{C}) \dim(\text{PGL}(K)) - h \text{rk}(\text{PGL}(K)) = 2(K^2 - 1)(g - 1) + K(K - 1)h. \tag{2.7}$$

Let $\text{Loc}_K(\mathcal{C})$ be the traditional moduli space of vector bundles with flat connections on \mathcal{C} . Forgetting the framing, we get a projection, which is a finite cover over the generic point:

$$\pi : \mathcal{X}_K(\mathcal{C}) \longrightarrow \text{Loc}_K(\mathcal{C}). \tag{2.8}$$

Over the locus of connections with unipotent holonomies around the punctures it is generically one to one, and partially resolves the singularities of the traditional unipotent moduli space.

2. The phase space and the coordinate phase space for a surface. The moduli space $\mathcal{X}_K(\mathcal{C})$ was introduced in [12], for the case with no lamination. It was proved there that any 2d ideal triangulation \mathbf{t} gives rise to a rational coordinate system $\{x_i^{\mathbf{t}}\}$ on the moduli space $\mathcal{X}_K(\mathcal{C})$. Moreover, it was shown that there is a Zariski open subset

$$\mathcal{X}_K(\mathcal{C}; \mathbf{t}) \subset \mathcal{X}_K(\mathcal{C}) \tag{2.9}$$

on which the coordinates $\{x_i^{\mathbf{t}}\}$ are well defined, and which is identified with a complex torus:

$$\mathcal{X}_K(\mathcal{C}; \mathbf{t}) \xrightarrow{=} (\mathbb{C}^*)^d, \quad p \longmapsto x_i^{\mathbf{t}}(p). \quad (2.10)$$

We call the complex torus $\mathcal{X}_K(\mathcal{C}; \mathbf{t})$ the *coordinate phase space associated with an ideal triangulation*, or just the *coordinate phase space*.

There is a holomorphic Poisson structure on $\mathcal{X}_K(\mathcal{C})$, which in any of the coordinate systems $\{x_i^{\mathbf{t}}\}$ is given by

$$\{x_i^{\mathbf{t}}, x_j^{\mathbf{t}}\} = \varepsilon_{ij}^{\mathbf{t}} x_i^{\mathbf{t}} x_j^{\mathbf{t}}, \quad \varepsilon_{ij}^{\mathbf{t}} \in \mathbb{Z}. \quad (2.11)$$

The Poisson tensor $\varepsilon_{ij}^{\mathbf{t}}$ depends on the choice of the triangulation. The eigenvalues of holonomies around holes in \mathbb{C} generate the center of the Poisson algebra; after fixing them, the Poisson tensor can be inverted to define a holomorphic symplectic form.

With non-empty lamination, the corresponding moduli space $\mathcal{X}_K(\mathcal{C}, \gamma)$ is studied and coordinatized in [14]; see also [15, 73] for the case $K = 2$.

3. The phase space for a torus T^2 . The phase space $\mathcal{X}_K(T^2)$ is the moduli space of flat $\mathrm{PGL}(K, \mathbb{C})$ -connections on T^2 together with a framing, given by a flat section of the associated flag bundle on T^2 . Alternatively, a framing is a choice of a flag in a fiber invariant under both A- and B-cycle holonomies, or a reduction of the structure group to a Borel subgroup. There is a birational equivalence, that is an isomorphism at the generic point

$$\mathcal{X}_K(T^2) \simeq (\mathbb{C}^* \times \mathbb{C}^*)^{K-1}. \quad (2.12)$$

It assigns to a framed flat connection on T^2 the *ordered* collections $\{\ell_a\}$ and $\{m_b\}$, where

$$\mathrm{diag}(1, \ell_1, \ell_1 \ell_2, \dots, \ell_1 \dots \ell_{K-1}) \quad \text{and} \quad \mathrm{diag}(1, m_1, m_1 m_2, \dots, m_1 \dots m_{K-1}) \quad (2.13)$$

are the diagonal parts of the $\mathrm{PGL}(K, \mathbb{C})$ -holonomies $\mathrm{Hol}_A, \mathrm{Hol}_B$ around the A- and B-cycles of the torus. Let κ_{ab} be the Cartan matrix. Then the holomorphic symplectic form is

$$\Omega_{T^2} = \mathrm{Tr} \, d \log \left(\frac{\mathrm{Hol}_A}{\det \mathrm{Hol}_A^{\frac{1}{K}}} \right) \wedge d \log \left(\frac{\mathrm{Hol}_B}{\det \mathrm{Hol}_B^{\frac{1}{K}}} \right) = \sum_{a,b=1}^{K-1} (\kappa^{-1})_{ab} d \log \ell_a \wedge d \log m_b, \quad (2.14)$$

corresponding to Poisson brackets $\{\ell_a, m_b\} = \kappa_{ab} \ell_a m_b$. The traditional moduli space of flat connections on the torus is birationally isomorphic to the quotient of (2.12) by the symmetric group S_K . It makes it more difficult to introduce the canonical coordinates, since they must be invariants of the symmetric group. The choice of invariant flag orders the eigenvalues. Notice that for the torus there is no extra unipotent condition.

Since ∂M is a disjoint union of tori T_j^2 and punctured surfaces with simple laminations $(\mathcal{C}_i, \gamma_i)$, the moduli space $\mathcal{X}_K^{\mathrm{un}}(\partial M)$ is a product:

$$\mathcal{X}_K^{\mathrm{un}}(\partial M) = \prod_i \mathcal{X}_K(T_i^2) \times \prod_j \mathcal{X}_K(\mathcal{C}_j, \gamma_j).$$

4. The moduli space $\mathcal{L}_K(M, \mathbf{t})$. Let us consider now an admissible 3-manifold M whose boundary ∂M does not have small annuli. Choose a triangulation \mathbf{t} of the big boundary of ∂M . We define

$$\mathcal{L}_K(M, \mathbf{t}) := \mathcal{L}_K(M) \cap \left(\prod_i \mathcal{X}_K(T_i^2) \times \prod_j \mathcal{X}_K(\mathcal{C}_j; \mathbf{t}_j) \right). \quad (2.15)$$

Let us see now what we get in our two running examples.

Examples: 1. Let $M = B_m$ be a 3d ball with $m \geq 3$ small discs on the boundary. It can be approached combinatorially as a convex polyhedron with m vertices. Then $\mathcal{X}_K^{\text{un}}(\partial M)$ is the moduli space of framed flat $\text{PGL}(K)$ -bundles with connections on S^2 minus m discs, and unipotent monodromies around the discs. Now let us figure out the Lagrangian subspace. Since any flat connection on a ball is trivial, the invariant flags are unrestricted, and are the only non-trivial part of the data. So $\mathcal{L}_K(B_m)$ is the configuration space of m flags, realized as a framing data in a trivialized vector bundle with connection:

$$\mathcal{L}_K(B_m) = \text{Conf}_m(\mathcal{B}) \subset \mathcal{X}_K^{\text{un}}(S^2 - m \text{ discs}). \quad (2.16)$$

The statement that this is a Lagrangian subvariety is non-trivial even in this case. It follows from the result of [12] that a flip of a 2d triangulation is a birational transformation that preserves the Poisson/symplectic structure.

2. Let $M = S^3 \setminus \mathcal{K}$ be the complement of a knot \mathcal{K} in S^3 . Then the small boundary is a torus T^2 . The $\text{SL}(2, \mathbb{C})$ version of $\mathcal{L}_2(M)$ is a complex curve.⁹ It is usually called the A-polynomial curve of the knot [16]. Therefore for $\text{SL}(K, \mathbb{C})$, we arrive at a natural generalization of the A-polynomial as a Lagrangian subvariety in the phase space (2.12). Upon quantization, the operators $\hat{\mathcal{L}}_K(M)$ quantizing the equations defining the Lagrangian subvariety $\mathcal{L}_K(M)$ are expected to provide recursion relations for the colored HOMFLY polynomials of a knot, just as the quantized A-polynomial (conjecturally) provides a recursion relation for the colored Jones polynomial [7, 74, 75]. (See e.g. [76–78] for quantizations of $\mathcal{L}_K(M)$ in several special cases.)

Conclusion: at first glance, the moduli space $\mathcal{X}_K^{\text{un}}(\partial M)$ of framed unipotent flat connections on the boundary ∂M of an admissible 3-manifold M is just a modification of the traditional moduli space $\text{Loc}_K^{\text{un}}(\partial M)$ of unipotent flat connections on ∂M , making it more accessible. It turns out however that the moduli space $\mathcal{X}_K^{\text{un}}(\partial M)$ is really indispensable in construction of the Lagrangian subvariety assigned to M :

Quite often, there is simply no room in $\text{Loc}_K^{\text{un}}(\partial M)$ for a Lagrangian assigned to M !

Indeed, the canonical projection $\pi : \mathcal{X}_K^{\text{un}}(\partial M) \longrightarrow \text{Loc}_K^{\text{un}}(\partial M)$ can map $\mathcal{L}_K(M)$ to a subvariety whose dimension is too small to be Lagrangian.

2.4 From framed flat bundles to configurations of flags

Given a framed $\text{PGL}(K, \mathbb{C})$ -vector bundle \mathcal{V} with a flat connection on a 2d surface \mathcal{C} with punctures, and given an ideal triangulation of \mathcal{C} , i.e. a triangulation with vertices at the

⁹More precisely, the top-dimensional components of $\mathcal{L}_2(M)$ are a complex curve. Exceptional zero-dimensional components can also arise, but they are usually excluded by additional stability conditions.




Figure 5. Parallel transport of flags to the common point p .

punctures, any ideal *triangle* t of the triangulation gives rise to a configuration of three flags (A, B, C) in a K -dimensional complex vector space V_K . Namely, consider the restriction of the flat bundle \mathcal{V} to the triangle t . Then the flat sections defining the framing provide flat sections of the associated flag bundle $\mathcal{V}_{\mathcal{B}}$ at each punctured disc on \mathcal{C} in the vicinity of the vertices of the triangle. Using the connection, we parallel transport each of the invariant flags to a point p inside of the triangle t . Since the connection is flat, and the triangle is contractible, the resulting collection of flags in the fiber \mathcal{V}_p of the bundle at p does not depend on the paths from the small triangles to p . The resulting configuration of three flags (A, B, C) does not depend on any choices involved. More generally, given any ideal polygon on \mathcal{C} we can run the same construction and construct a configuration of flags in V_K corresponding to the flat sections near the vertices of the polygon. This construction goes back to [12], and leads to a construction of the coordinates associated with the triangulation in *loc. cit.* We will recall it later on, in section 3.1.

Literally the same construction with minimal adjustments can be applied to framed flat vector bundles on an admissible 3-manifold with a given ideal triangulation. Here is how it goes. Consider a framed $\text{PGL}(K, \mathbb{C})$ -bundle with a flat connection \mathcal{V} on an admissible 3-manifold M , and an ideal triangulation of M . Then an ideal tetrahedron T of the triangulation gives rise to a configuration of four flags (A, B, C, D) in a K -dimensional complex vector space V_K . Indeed, let us restrict \mathcal{V} to the truncated tetrahedron T . The flat sections defining the framing provide flat sections of the flag bundle $\mathcal{V}_{\mathcal{B}}$ at each of the four small triangles assigned to the vertices of the original tetrahedra. Using the connection, we parallel transport the corresponding invariant flags to a point p inside of the tetrahedra T . Since the connection is flat, and the tetrahedron is contractible, the resulting collection of flags in the fiber \mathcal{V}_p does not depend on the paths from the small triangles to p . So the resulting configuration of four flags (A, B, C, D) is well defined. Similarly, any ideal polyhedron in M , and in particular an ideal triangle or an ideal edge of the triangulation, leads to a configuration of flags labeled by the vertices of the polyhedron.

The very notions of a *framed* bundle with flat connection and an *admissible* manifold were designed in such a way that we can run the above constructions. They deliver a collection of configurations of flags assigned to the ideal tetrahedra (in 3d) or triangles/rectangles (in 2d). So the next important question is how to deal with configurations of flags.

2.5 Hypersimplicial K -decomposition

We now describe a construction that plays a crucial role in this paper, namely the hypersimplicial K -decomposition of an ideal triangulation in dimension n .




Figure 6. Forming $\Delta^{1,1}$.

Let p and q be two non-negative integers satisfying $p + q = n - 1$. An n -dimensional hypersimplex $\Delta^{p,q}$ is defined [79] as the intersection of the $n+1$ -dimensional cube $0 \leq x_i \leq 1$ with the hyperplane $\sum_{i=0}^n x_i = q + 1$. Combinatorially, it is isomorphic to the convex hull of the centers of p -dimensional faces of an n -dimensional simplex.

The hypersimplices $\Delta^{p,0}$ and $\Delta^{0,q}$ are just simplices. The simplest hypersimplex different from a simplex is the octahedron $\Delta^{1,1}$. It is the convex hull of the centers of the edges of a tetrahedron (figure 6).

The boundary of a hypersimplex $\Delta^{p,q}$ is a union of $n+1$ hypersimplices $\Delta^{p-1,q}$ and $n+1$ hypersimplices $\Delta^{p,q-1}$. Indeed, they are given by the intersections with the hyperplanes $x_i = 0$ and $x_i = 1$ of the unit cube. For example, the boundary of the octahedron $\Delta^{1,1}$ consists of four $\Delta^{1,0}$ -triangles and four $\Delta^{0,1}$ -triangles. The boundary of $\Delta^{2,1}$ is given by five octahedra $\Delta^{1,1}$ and five $\Delta^{2,0}$ -tetrahedra, etc.

Consider the standard coordinate space \mathbb{R}^{n+1} . It contains the integral lattice $\mathbb{Z}^{n+1} \subset \mathbb{R}^{n+1}$. The hyperplanes $x_i = s$, where $s \in \mathbb{Z}$ and $i = 0, \dots, n$, cut the space into unit cubes with vertices at the integral points. Given a positive integer K , consider the n -dimensional simplex Δ_K^n given by the intersection of the hyperplane $\sum x_i = K$ with the positive orthant $x_i \geq 0$:

$$\Delta_K^n = \left\{ (x_0, \dots, x_n) \mid \sum_{i=0}^n x_i = K \right\}. \tag{2.17}$$

The hyperplanes $x_i = s$, where $s \in \mathbb{Z}$, cut the simplex Δ_K^n into a union of hypersimplices. Indeed, the hyperplane $\sum x_i = K$ intersect each of the standard unit lattice cubes either by an empty set, or by a hypersimplex.

Definition 2.3. [12] A hypersimplicial K -decomposition (or K -decomposition, for short) of an n -dimensional simplex is a decomposition of the simplex Δ_K^n into hypersimplices provided by the hyperplanes $x_i = s$, $s \in \mathbb{Z}$.

The polyhedra of a hypersimplicial K -decomposition have vertices at the lattice points

$$p_{b_0, \dots, b_n}, \quad b_0 + \dots + b_n = K, \quad b_i \in \mathbb{Z}. \tag{2.18}$$

The hypersimplices of K -decomposition are the closures of connected components of the complement to the hyperplanes $x_i = s$. The hypersimplices $\Delta^{p,q}$ of the K -decomposition match the $\binom{K+q}{n}$ solutions of the equation

$$a_0 + \dots + a_n = K - (p + 1), \quad a_i \in \mathbb{Z}_{\geq 0}. \tag{2.19}$$




Figure 7. 2-simplex with lattice points (a, b, c) , $a + b + c = 3$.

Finally, a K -decomposition of a simplex induces K -decompositions of any faces of the simplex.

Examples:

1. A K -decomposition of a segment is its decomposition into K equal little segments.
2. When $n = 2$ we get a K -triangulation of a triangle. It is given by three families of parallel lines, each consisting of K lines including one of the sides, with triple point intersections, as shown on the figure 7. They induce decompositions of the sides of the triangle into K little segments.

A K -triangulated triangle carries $\binom{K+2}{2}$ lattice points p_{abc} indexed by triples of non-negative integers (a, b, c) with $a + b + c = K$. The original triangle is decomposed into K^2 small triangles. They are of two types: $\binom{K+1}{2}$ of them are “upright” $\Delta^{0,1}$ -triangles, shaded white, and labeled by the triples (a, b, c) with $a + b + c = K - 1$; $\binom{K}{2}$ of them are “upside-down” $\Delta^{1,0}$ -triangles, shaded black and labeled by the triples (a, b, c) with $a + b + c = K - 2$.

3. When $n = 3$ we get a K -decomposition of a tetrahedron. It consists of $\binom{K+2}{3}$ upright $\Delta^{0,2}$ -tetrahedra, $\binom{K+1}{3}$ octahedra $\Delta^{1,1}$, and $\binom{K}{3}$ upside-down $\Delta^{2,0}$ -tetrahedra. It induces a K -triangulation of each of the four faces of the tetrahedron. The lattice points p_{abcd} are labeled by quadruples of non-negative integers that sum to K , see figure 8.
4. When $n = 4$ we get a K -decomposition of a four-dimensional simplex into hypersimplices of type $\Delta^{0,3}$, $\Delta^{1,2}$, $\Delta^{2,1}$, $\Delta^{3,0}$.

2.6 Configuration of flags and hypersimplices

Consider a generic configuration of $n + 1$ flags $(F_0^\bullet, \dots, F_n^\bullet)$ in a K -dimensional complex vector space V_K . We assign them to the vertices of an n -dimensional simplex Δ_K^n .

These flags define a collection of $(p + 1)$ -dimensional linear subspaces in V_K , assigned to the $\Delta^{p,q}$ -hypersimplices of the hypersimplicial K -decomposition of the simplex Δ_K^n . Namely, let $\Delta_{\mathbf{a}}^{p,q}$ be the $\Delta^{p,q}$ -hypersimplex corresponding to a given partition

$$\mathbf{a} = (a_0, \dots, a_n), \quad a_0 + \dots + a_n = K - (p + 1), \quad a_i \geq 0. \quad (2.20)$$

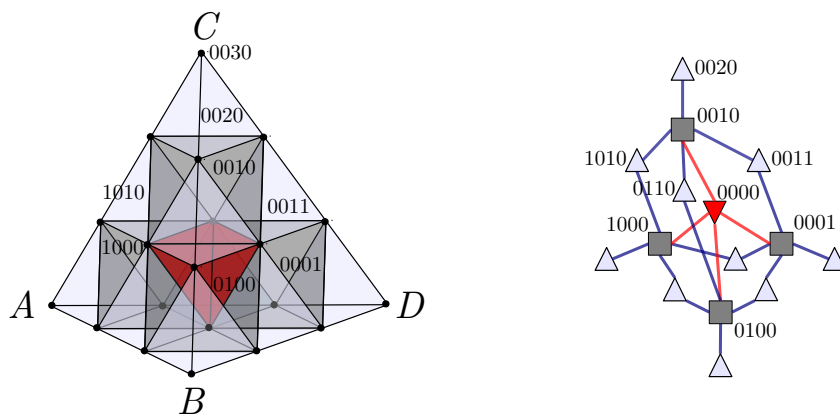


Figure 8. The 3-decomposition of the tetrahedron. The upright $\Delta^{0,2}$ -tetrahedra (white) correspond to lines, the octahedra (black) to planes, and the upside-down $\Delta^{2,0}$ -tetrahedron (red) to a 3-space. Adjacent hypersimplices obey incidence relations, which can be represented as a hypergraph, shown on the right.

Definition 2.4. Given a generic configuration of $n + 1$ flags $(F_0^\bullet, \dots, F_n^\bullet)$ in V_K , we assign to $\Delta_{\mathbf{a}}^{p,q}$ a $(p + 1)$ -dimensional subspace $\mathfrak{F}_{\mathbf{a}}$ of V_K given by intersection of the flag subspaces $F_i^{a_i}$:

$$\Delta_{\mathbf{a}}^{p,q} \rightsquigarrow \mathfrak{F}_{\mathbf{a}} = F_0^{a_0} \cap F_1^{a_1} \cap \dots \cap F_n^{a_n}, \quad \mathbf{a} = (a_0, \dots, a_n). \quad (2.21)$$

Each hypersimplex $\Delta_{\mathbf{a}}^{p,q}$ is surrounded by $n + 1$ hypersimplices $\Delta_{\mathbf{b}}^{p-1,q+1}$, where \mathbf{b} is obtained from \mathbf{a} by adding 1 to one of the coordinates (a_0, \dots, a_n) . So the collection of \mathbf{b} 's is

$$(a_0 + 1, a_1, \dots, a_n), (a_0, a_1 + 1, \dots, a_n), \dots, (a_0, a_1, \dots, a_n + 1).$$

Therefore we get a collection of $n + 1$ codimension-one subspaces $\mathfrak{F}_{\mathbf{b}} \subset \mathfrak{F}_{\mathbf{a}}$.

One way to organize this data is to assign to the hypersimplex $\Delta_{\mathbf{a}}^{p,q}$ the configuration of $n + 1$ hyperplanes $\mathfrak{F}_{\mathbf{b}}$ in the space $\mathfrak{F}_{\mathbf{a}}$, which we denote by $\mathcal{PC}_{\mathbf{a}}(F_0^{a_0}, F_1^{a_1}, \dots, F_n^{a_n})$.

Definition 2.5. Given a generic configuration of $n + 1$ flags $(F_0^\bullet, \dots, F_n^\bullet)$ in V_K , we assign to $\Delta_{\mathbf{a}}^{p,q}$ a configuration of $n + 1$ hyperplanes in a p -dimensional projective space:

$$\Delta_{\mathbf{a}}^{p,q} \longmapsto \mathcal{PC}_{\mathbf{a}}(F_0^{a_0}, F_1^{a_1}, \dots, F_n^{a_n}) \in \text{Conf}_{n+1}(\mathbb{P}^p), \quad \mathbf{a} = (a_0, \dots, a_n). \quad (2.22)$$

Equivalently, it is a configuration of $n + 1$ points in the p -dimensional projective space of hyperplanes in $\mathfrak{F}_{\mathbf{a}}$. Here are two basic examples (explained in much greater detail in section 3):

- Let $n = 3$. Then the octahedron $\Delta_{\mathbf{a}}^{1,1}$ gives rise to a configuration of four lines $\mathfrak{F}_{\mathbf{b}}$ in a 2-plane $\mathfrak{F}_{\mathbf{a}}$. The cross-ratio of these four lines is the coordinate $z_{\mathbf{a}}$ assigned to the octahedron $\Delta_{\mathbf{a}}^{1,1}$.
- Let $n = 4$. Then the hypersimplex $\Delta_{\mathbf{a}}^{1,2}$ gives rise to a configuration of five lines $\mathfrak{F}_{\mathbf{b}}$ in a 2-plane $\mathfrak{F}_{\mathbf{a}}$. The hypersimplex $\Delta_{\mathbf{a}}^{2,1}$ gives rise to a configuration of five planes $\mathfrak{F}_{\mathbf{b}}$ in a 3-space $\mathfrak{F}_{\mathbf{a}}$.

Notice that generic configurations of five planes in a 3-space can be identified with generic configurations of five lines in a 2-plane. A nice way to see this is to pass to the corresponding projective configuration of five points in the projective plane, and notice that given five generic points in the projective plane, there is a unique conic containing them. This conic is isomorphic to a projective line, thus providing an isomorphism. (Another way to see it is to combine the two canonical isomorphisms (4.17) and (4.18) from section 4.2.2.)

This way we get two kinds of “pentagons” related to the octahedral coordinates. We will return to these two pentagons in section 3.5, and explore the relevant geometry.

Remark. The constructions in definitions 2.4, 2.5 are dual to the one used in [66]. Indeed, a flag in a vector space V determines the dual flag in the dual space V^* . A configuration of flags in V gives rise to configuration of the dual flags; the intersection of flag subspaces corresponds to the quotients by direct sums of flag subspaces in [66].

3 Localization of framed flat connections

Given a 2d ideal triangulation \mathbf{t} of the big boundary of an admissible 3-manifold M , we know from [12] how to build coordinates on a Zariski open subset $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ of the space of framed flat connections on ∂M — the coordinate phase space associated with ∂M and \mathbf{t} . We will review the construction in section 3.1. (We also explained how to extend the definition of $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ to small torus boundary components in section 2.3.) Now let \mathbf{t}_{3d} be a 3d ideal triangulation of M compatible with a triangulation \mathbf{t} of the big boundary. We will construct coordinates associated with the octahedra of the K -decomposition of M . We will prove in sections 3.2–3.3 that octahedron coordinates in the bulk of M parametrize an open subset $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ of the space $\tilde{\mathcal{L}}_K(M)$ of framed flat connections on M . The space $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ projects to a submanifold $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ that parametrizes the flat connections in the boundary phase space that extend to the bulk. By using sufficiently refined 3d triangulations, or taking a union over all 3d triangulations, we obtain a submanifold $\mathcal{L}_K(M, \mathbf{t})$ that only depends on \mathbf{t} .

Our main *conceptual* goal is to understand the nature of $\mathcal{L}_K(M, \mathbf{t})$ as a *Lagrangian submanifold* of $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$, with its Atiyah-Bott-Goldman symplectic structure. We will do this by expressing $\mathcal{L}_K(M, \mathbf{t}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ as a symplectic gluing of elementary symplectic pairs $\mathcal{L}_\diamond \subset \mathcal{P}_{\partial_\diamond}$ associated with octahedra. We start explaining how this should work in this section, but defer a full treatment to sections 4 and 5.

In section 3.5, we will investigate how 2–3 moves act on $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ and $\mathcal{L}_K(M, \mathbf{t}_{3d})$, and interpreted 2–3 moves in terms of the hypersimplicial K -decompositions of 4-simplices.

3.1 Boundary phase spaces

Let \mathcal{C} be an oriented surface with at least one puncture and $\chi(\mathcal{C}) < 0$ — e.g. a component of the big boundary of an admissible 3-manifold M . Our first item of business is to




Figure 9. 3-triangulation of a triangle, with triples of integers corresponding to points p_{abc} with $a + b + c = 3$, lines \mathfrak{L}_{abc} with $a + b + c = 2$ (black dots), and planes \mathfrak{P}_{abc} with $a + b + c = 1$ (white dots). The blue connectors between black and white dots stand for incidence relations between lines and planes.

review the \mathcal{X} -coordinates on the moduli space $\mathcal{X}_K(\mathcal{C})$ of framed flat $\text{PGL}(K, \mathbb{C})$ -connections on \mathcal{C} [12].¹⁰

3.1.1 Coordinates from the K -triangulation

Let us fix an ideal triangulation of \mathcal{C} . Then, given a $\text{PGL}(K)$ -vector bundle with a framed flat connection, each ideal triangle gives rise to a configuration of three flags (A, B, C) in V_K , assigned to the vertices of the triangle (as discussed in section 2.4). We assume that it is a generic configuration of flags. Then it gives rise to configurations of lines and planes in V_K , associated with the white $\Delta^{0,1}$ - and black $\Delta^{1,0}$ -triangles in the K -triangulation:

$$\Delta_{abc}^{0,1} \rightsquigarrow \mathfrak{L}_{abc} = A^a \cap B^b \cap C^c, \quad a + b + c = K - 1, \quad (3.1)$$

$$\Delta_{abc}^{1,0} \rightsquigarrow \mathfrak{P}_{abc} = A^a \cap B^b \cap C^c, \quad a + b + c = K - 2, \quad (3.2)$$

as in figure 9. The plane \mathfrak{P}_{abc} on a black triangle contains all three lines $\mathfrak{L}_{a+1,b,c}$, $\mathfrak{L}_{a,b+1,c}$, and $\mathfrak{L}_{a,b,c+1}$ on the white triangles surrounding it. If the configuration of flags (A, B, C) is generic, these are the only relations among the lines.

Next, every *internal* lattice point p_{abc} of the K -triangulation, which is labeled by three *strictly positive* integers that add to K , gives rise to a 3-space

$$\mathfrak{V}_{a-1,b-1,c-1} = A^{a-1} \cap B^{b-1} \cap C^{c-1}.$$

This space contains all three planes and all six lines on the black and white triangles surrounding it. This data allows us to define a coordinate x_{abc} associated with the internal point. It is a triple-ratio of the collection of lines and planes contained in $\mathfrak{V}_{a-1,b-1,c-1}$,

¹⁰These \mathcal{X} -coordinates were recently reviewed and generalized in [28, 29] in the context of the physical 6d (2,0) theory compactified on surfaces \mathcal{C} . The work of [28, 29] should tie in beautifully with our 3d constructions, though many details remain to be explored.




Figure 10. Two triangles glued along an edge. The points p_{0bc0} on the edge are shown in magenta, together with the incidence relations with the neighboring planes (white dots). The lines (black dots) on adjacent white triangles are identified. Internal lattice points are shown in yellow.

and there are several ways to describe it. For example, choosing any six vectors $v_{a'b'c'}$ that generate the six respective lines $\mathfrak{L}_{a'b'c'}$ that surround p_{abc} , and setting

$$\begin{aligned} l_1 &:= v_{a+1,b-1,c-1}, & l_2 &:= v_{a-1,b+1,c-1}, & l_3 &:= v_{a-1,b-1,c+1}, \\ l_4 &:= v_{a-1,b,c}, & l_5 &:= v_{a,b-1,c}, & l_6 &:= v_{a,b,c-1}, \end{aligned}$$

we can define the triple-ratio as

$$x_{abc} = \frac{\langle l_1 \wedge l_6 \wedge l_4 \rangle \langle l_2 \wedge l_4 \wedge l_5 \rangle \langle l_3 \wedge l_5 \wedge l_6 \rangle}{\langle l_1 \wedge l_5 \wedge l_4 \rangle \langle l_2 \wedge l_6 \wedge l_5 \rangle \langle l_3 \wedge l_4 \wedge l_6 \rangle}. \quad (3.3)$$

Here $\langle v \wedge v' \wedge v'' \rangle$ is defined as follows. Choose a volume form Ω_3 in a three-dimensional vector space V_3 . Given a triple of vectors v, v', v'' in V_3 , we set $\langle v \wedge v' \wedge v'' \rangle := (\Omega_3, v \wedge v' \wedge v'')$. In any unimodular basis in V_3 , it is the determinant of the matrix expressing v, v', v'' in this basis.

The triple-ratio is independent of the choice of the vectors v , since every vector occurs an equal number of times in the numerator and denominator. It is also independent of the choice of the volume form Ω_3 . The triple-ratio is the only invariant of a configuration of three flags in V_3 . More generally, the $\binom{K-1}{2}$ triple-ratios x_{abc} assigned to interior points of a K -triangulation parametrize the space of configurations of three flags (A, B, C) in V_K .

To describe framed flat connections on the surface \mathcal{C} , the triple-ratio coordinates in the interior of each triangle are supplemented by cross-ratio coordinates on edges. Every edge of the ideal triangulation belongs to a distinguished quadrilateral, formed by the two triangles sharing the edge. We parallel-transport the four flags A, B, C, D at its vertices to a point p in its interior. Then we can simultaneously decorate the K -triangulations of both triangles with a collection of lines, planes, and 3-spaces. We label the objects on the triangle ABC with quadruples of integers $(a, b, c, 0)$ and those on the triangle BCD with quadruples $(0, b, c, d)$. See figure 10.

The $K - 1$ lattice points p_{abcd} on the edge have labels $(0, b, c, 0)$ with $b, c > 0$ and $b + c = K$. They belong to the K -decomposition of both adjacent triangles. The lines on the white triangles adjacent to the edge are identified. The four lines surrounding p_{0bc0} have labels $\mathfrak{L}_{1,b-1,c-1,0}$, $\mathfrak{L}_{0,b,c-1,0}$, $\mathfrak{L}_{0,b-1,c,0}$, and $\mathfrak{L}_{0,b-1,c-1,1}$, and belong to the same copy




Figure 11. Unipotent holonomy constraints around a puncture.

of $V_2 = B^{b-1} \cap C^{c-1} \subset V_K$. To define a cross-ratio coordinate, we choose four vectors a, b, c, d inside each respective line, and set

$$x_{0bc0} = \frac{\langle a \wedge b \rangle \langle c \wedge d \rangle}{\langle a \wedge c \rangle \langle b \wedge d \rangle}. \tag{3.4}$$

Here $\langle v \wedge v' \rangle$ is defined using a volume form in V_2 .

A surface \mathcal{C} of genus g with h punctures has $-2\chi(\mathcal{C})$ triangles and $-3\chi(\mathcal{C})$ edges in any ideal triangulation. Therefore, the cross-ratios and triple-ratios form a total of

$$-((K-1)(K-2) + 3(K-1))\chi(\mathcal{C}) = -(K^2 - 1)\chi(\mathcal{C}) \tag{3.5}$$

local coordinates for the space $\mathcal{X}_K(\mathcal{C})$. The requirement that the holonomies at every puncture are unipotent imposes $(K-1)h$ independent constraints on the coordinates. Explicitly, the constraints at a given puncture say that the product of *minus* the coordinates around each of the $K-1$ loops surrounding the puncture is equal to one. This is illustrated in figure 11, in the case that \mathcal{C} is the boundary of an ideal tetrahedron, and $K=3$; the products of three coordinates ($-x_\bullet$) around the first loop and six coordinates ($-x_\bullet$) around the second loop must both equal one. (The origin of these constraints is explained in [12]; we will re-derive the constraints in section 3.3.)

Altogether, we find that cross-ratios and triple-ratios give

$$d_K(g, h) = -(K^2 - 1)\chi(\mathcal{C}) - (K-1)h \tag{3.6}$$

independent coordinates for the moduli space $\mathcal{X}_K^{\text{un}}(\mathcal{C})$. This is precisely the expected complex dimension. The coordinates are \mathbb{C}^* -valued and cover an open patch. Notice that the cross-ratios and triple-ratios never equal zero or infinity, as long as all of the invariant flags at the punctures are generic.

Dual coordinates: in [12], the coordinates for the moduli space of framed flat connections on a punctured surface were defined slightly differently than we have presented above. Rather than defining lines \mathfrak{L}_{abc} and \mathfrak{L}_{abcd} and their invariants by intersecting subspaces of flags, [12] used quotients of flags. The two approaches are very simply related by the operation of dualizing flags, as in (2.3). That is, our cross-ratio and triple-ratio




Figure 12. Poisson structure on a 4-triangulated triangle. The coordinates x and x' at the lattice points p and p' have the Poisson bracket $\{x, x'\} = +xx'$ since there is one positively oriented edge from p to p' .

coordinates defined from a collection of flags (A, B, C, \dots) (and a flat connection) on \mathcal{C} as above are equivalent to (minus) the coordinates used in [12] defined using the dual flags (A^*, B^*, C^*, \dots) . The cross-ratio defined in [12] differs from ours by an additional minus sign.

3.1.2 Symplectic structure

The Poisson structure on an open patch of $\mathcal{X}_K(\mathcal{C})$ corresponding to an ideal triangulation is described by drawing arrows circulating clockwise around every small black triangle in the K -decomposition of the triangulated surface \mathcal{C} (see figure 12). Then the Poisson bracket of two cross-ratio or triple-ratio coordinates x, x' at lattice points p, p' is

$$\{x, x'\} = (\# \text{ edges from } p \text{ to } p' - \# \text{ edges from } p' \text{ to } p) \cdot xx'. \quad (3.7)$$

After projection to the traditional moduli space $\text{Loc}_K(\mathcal{C})$, the symplectic and Poisson structures agree with the standard Goldman symplectic form on the non-singular part of the moduli space. The latter is expressed simply in terms of complex flat connections \mathcal{A} as $\Omega = \int_{\mathcal{C}} \text{Tr}(\delta\mathcal{A} \wedge \delta\mathcal{A})$.

3.2 The Lagrangian pair assigned to a 3d triangulation of M

Having described coordinates for the moduli space of framed flat connections on oriented 2d punctured surfaces \mathcal{C} , we now proceed to study the framed flat connections on admissible 3-manifolds M . In the process of building coordinates for flat 3d connections, we will find a natural way to understand the spaces $\mathcal{L}_K(M)$ of framed flat $\text{PGL}(K)$ -connections on the boundary ∂M that extend to the bulk.

Precisely, we fix an ideal triangulation \mathbf{t} of the big boundary of ∂M . Then, as we discussed in section 2.3, there is a complex torus $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$, called the coordinate phase space, realized as a Zariski open part of the symplectic moduli space $\mathcal{X}_K^{\text{un}}(\partial M)$. Pick a 3d triangulation \mathbf{t}_{3d} of M inducing the 2d triangulation \mathbf{t} . We determine explicitly a Zariski open subset $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ of framed flat connections on M and a corresponding Lagrangian subvariety

$$\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}), \quad (3.8)$$

which parametrizes the framed flat connections from $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ that extend to the bulk. Specifically, $\mathcal{L}_K(M, \mathbf{t}_{3d})$ is the (closure of) the image of the projection of $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ to $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$. We call (3.8) a *Lagrangian pair*, emphasizing both the complex symplectic torus and its Lagrangian subvariety.

Generically, $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ and $\mathcal{L}_K(M, \mathbf{t}_{3d})$ will turn out to be independent of \mathbf{t}_{3d} , depending only on the fixed boundary triangulation \mathbf{t} . For certain “degenerate” choices of \mathbf{t}_{3d} , these spaces may lose components. By taking a union over all 3d triangulations, we can obtain spaces $\tilde{\mathcal{L}}_K(M, \mathbf{t})$ and $\mathcal{L}_K(M, \mathbf{t})$ that depend only on the boundary triangulation \mathbf{t} .

The boundary of a tetrahedron, understood as an admissible 3-manifold, is a four-punctured sphere with a canonical triangulation. This triangulation, for $K = 2$, provides the most basic Lagrangian pair (3.8), which we call the *elementary Lagrangian pair* — the Lagrangian pair for the $\text{PGL}(2)$ -tetrahedron. It is our basic building block, and deserves a special notation:

$$\mathcal{L}_\diamond \subset \mathcal{P}_{\partial\diamond}. \tag{3.9}$$

We assign the elementary Lagrangian pair to each of the octahedrons of the K -decomposition of M provided by the 3d triangulation \mathbf{t}_{3d} .

3.2.1 The Lagrangian pair for the $\text{PGL}(2)$ -tetrahedron

Let us review the framed flat $\text{PGL}(2, \mathbb{C})$ -connections on a tetrahedron. We denote by \mathbf{t} the canonical triangulation of the tetrahedron. Let us choose framing flags A, B, C, D at the four punctures, as in figure 13. Each flag is just a line. Abusing notation, we denote the line defining the flag A by A , etc. Just as in section 3.1, we parametrize $\mathcal{X}_2(\partial\Delta, \mathbf{t})$ by six coordinates x_1, \dots, x_6 assigned to the edges. Each x_i is obtained by parallel-transporting the four flags to a common point p_i near the edge, and then taking a cross-ratio. For example,

$$x_1 = \left. \frac{\langle a \wedge b \rangle \langle c \wedge d \rangle}{\langle a \wedge c \rangle \langle b \wedge d \rangle} \right|_{p_1}, \tag{3.10}$$

where we emphasize the evaluation of flags at p_1 . Computing (3.10), we choose any four non-zero vectors (a, b, c, d) in the lines (A, B, C, D) . There are four relations among the x 's, reflecting the unipotent holonomy: the product of $(-x_i)$ around every vertex equals one, i.e.

$$x_1 x_2 x_3 = x_1 x_4 x_5 = x_3 x_4 x_6 = x_2 x_5 x_6 = -1. \tag{3.11}$$

The corresponding four monomials are central with respect to the Poisson bracket (3.7).

Relations (3.11) imply that coordinates on opposite edges are equal. It is then convenient to re-label the x 's as z, z', z'' (as on the right of figure 13), with $z z' z'' = -1$. We therefore find a two-dimensional complex torus sitting as a Zariski open part of the phase space $\mathcal{X}_2^{\text{un}}(\partial\Delta)$:

$$\mathcal{P}_{\partial\diamond} := \mathcal{X}_2^{\text{un}}(\partial\Delta, \mathbf{t}) \subset \mathcal{X}_2^{\text{un}}(\partial\Delta). \tag{3.12}$$

The notation emphasizes the octahedron \diamond . We call it the *octahedron phase space*. Explicitly,

$$\mathcal{P}_{\partial\diamond} = \{z, z', z'' \in \mathbb{C}^* \mid z z' z'' = -1\}, \tag{3.13}$$




Figure 13. Coordinates for framed $\text{PGL}(2, \mathbb{C})$ -connections on an ideal tetrahedron.

with Poisson brackets

$$\{z, z'\} = zz', \quad \{z', z''\} = z'z'', \quad \{z'', z\} = z''z. \quad (3.14)$$

Now let's fill in the bulk of the tetrahedron. Since the bulk is contractible, any bundle with a flat connection on the bulk is trivial. So the choice of invariant flags A, B, C, D becomes absolutely crucial — the flags are unrestricted, and they carry the only nontrivial degrees of freedom of framed flat connections in the bulk.

After filling in the tetrahedron we can parallel-transport all flags to a common point p of the tetrahedron. The previous three cross-ratios z, z', z'' can be computed at this point:

$$z = \frac{\langle a \wedge b \rangle \langle c \wedge d \rangle}{\langle a \wedge c \rangle \langle b \wedge d \rangle} \Big|_p, \quad z' = \frac{\langle b \wedge d \rangle \langle c \wedge a \rangle}{\langle b \wedge c \rangle \langle d \wedge a \rangle} \Big|_p, \quad z'' = \frac{\langle d \wedge a \rangle \langle c \wedge b \rangle}{\langle d \wedge c \rangle \langle a \wedge b \rangle} \Big|_p. \quad (3.15)$$

They are no longer independent. In addition to the relation $zz'z'' = -1$, there is a new Plücker relation

$$\mathcal{L}_\diamond := \mathcal{L}_2(\Delta, \mathbf{t}_{3d}) = \tilde{\mathcal{L}}_2(\Delta, \mathbf{t}_{3d}) : \quad z'' + z^{-1} - 1 = 0. \quad (3.16)$$

The curve defined by (3.16) parametrizes framed flat connections in the bulk (with the canonical triangulation \mathbf{t}_{3d}) for which the configuration of framing flags is generic. It is isomorphic to its image $\mathcal{L}_2(\Delta, \mathbf{t}_{3d}) \subset \mathcal{X}_2^{\text{un}}(\partial\Delta, \mathbf{t})$, the subset of connections on the boundary that extend to the bulk. So we have defined the elementary Lagrangian pair $\mathcal{L}_\diamond \subset \mathcal{P}_{\partial\Delta}$.

3.2.2 The Lagrangian pair for the $\text{PGL}(K)$ -tetrahedron

In order to generalize to $\text{PGL}(K, \mathbb{C})$ -connections, we consider a K -decomposition of the 3d tetrahedron as in figure 8, repeated for convenience on the left of figure 14.

Framed flat connections on the tetrahedron come with a choice of four invariant flags A, B, C, D near the vertices. Working on the boundary of the tetrahedron, we parallel-transport the flags to the midpoint of any of the six edges, by restricting the connection to the quadrilateral defined by the faces adjacent to the edge. The resulting configuration of four flags at this midpoint gives rise to a system of lines \mathcal{L}_{abcd} on the white triangles lying on the two tetrahedron faces adjacent to the edge. These lines then give us cross-ratio and triple-ratio coordinates for the lattice points on these faces and on the edge, just as in figure 10.




Figure 14. *Left:* 3-decomposition of a tetrahedron. *Right:* cross-ratios for an octahedron.

Repeating this for each of the six edges we construct all the cross-ratio and triple-ratio coordinates on $\mathcal{X}_K^{\text{un}}(\partial\Delta, \mathbf{t})$ assigned to the canonical triangulation of the tetrahedron boundary. There are $4 \times \binom{K-1}{2} + 6 \times (K-1)$ coordinates altogether, with $4 \times (K-1)$ relations coming from the unipotent-holonomy constraints at the four vertices. This leaves exactly $2(K-1)^2$ independent coordinates on $\mathcal{X}_K^{\text{un}}(\partial\Delta, \mathbf{t})$.

Just as in the $\text{SL}(2)$ case, if we fill in the bulk of the 3d tetrahedron, the flat $\text{PGL}(K, \mathbb{C})$ -connection must be trivial. The choice of flags is the only data remaining. Now the coordinates obey additional Plücker relations. The simplest and most beautiful way to describe the relations is to use octahedral coordinates.

A generic configuration of four flags (A, B, C, D) in V_K defines a configuration of lines in V_K assigned to the $\Delta^{0,2}$ -tetrahedra, and 2-planes assigned to the octahedra $\diamond_{abcd} := \Delta_{abcd}^{1,1}$, of the K -decomposition:

$$\Delta_{abcd}^{0,2} \rightsquigarrow \mathfrak{L}_{abcd} = A^a \cap B^b \cap C^c \cap D^d, \quad a + b + c + d = K - 1. \quad (3.17)$$

$$\diamond_{abcd} = \Delta_{abcd}^{1,1} \rightsquigarrow \mathfrak{P}_{abcd} = A^a \cap B^b \cap C^c \cap D^d, \quad a + b + c + d = K - 2. \quad (3.18)$$

The four lines sitting on the $\Delta^{0,2}$ -tetrahedra surrounding an octahedron \diamond_{abcd} are the lines

$$\mathfrak{L}_{a+1,b,c,d}, \quad \mathfrak{L}_{a,b+1,c,d}, \quad \mathfrak{L}_{a,b,c+1,d}, \quad \mathfrak{L}_{a,b,c,d+1}, \quad a + b + c + d = K - 2. \quad (3.19)$$

They belong to the 2-plane \mathfrak{P}_{abcd} . They provide the three cross-ratios *exactly* as in (3.15), replacing the (A, B, C, D) lines there with the four lines here:

$$z_{abcd}, \quad z'_{abcd}, \quad z''_{abcd}, \quad a + b + c + d = K - 2. \quad (3.20)$$

These cross-ratios are associated with the six vertices of the octahedron \diamond_{abcd} . The cross-ratios assigned to opposite vertices are equal, as on the right of figure 14 (cf. figure 13). They satisfy the standard monomial relation

$$z_{abcd} z'_{abcd} z''_{abcd} = -1, \quad (3.21)$$

as well as the Plücker relation

$$z''_{abcd} + z_{abcd}^{-1} - 1 = 0. \tag{3.22}$$

These relations encode the space of framed flat connections in the bulk of the tetrahedron and (isomorphically) the Lagrangian submanifold $\mathcal{L}_K(\Delta)$ of framed connections on the boundary that extend to the bulk. To see this, we must relate octahedron parameters to the cross-ratio and triple-ratio coordinates on the boundary.

The key to relating bulk and boundary coordinates is the realization that the 3d array of $\binom{K+2}{3}$ lines in the K -decomposition of the tetrahedron restricts to the standard 2d arrays of lines in the K -decomposition of the boundary.

Consider, for example, a lattice point p_{abcd} on an edge of the boundary. It has a label induced by the K -decomposition of the bulk, with exactly two of a, b, c, d vanishing. A unique octahedron vertex touches this point, and the cross-ratio x_{abcd} that we would assign from the boundary perspective is nothing but the bulk cross-ratio at the octahedron vertex. Explicitly, the $6 \times (K - 1)$ edge coordinates on the boundary of the tetrahedron are

$$\begin{aligned} \text{edges:} \quad & x_{0bc0} = z_{0b'c'0}, \quad x_{ab00} = z'_{a'b'00}, \quad x_{a0c0} = z''_{a'0c'0}, \\ & x_{a00d} = z_{a'00d'}, \quad x_{00cd} = z'_{00c'd'}, \quad x_{0b0d} = z''_{0b'0d'}, \end{aligned} \tag{3.23}$$

where $(a', b', c', d') = (a - 1, b - 1, c - 1, d - 1)$.

Similarly, every triple-ratio x_{abcd} on the boundary, with exactly one of a, b, c, d vanishing, is the product of the three octahedron parameters at the vertices of the three octahedra that touch p_{abcd} . For example, on face BCD , we have triple-ratio coordinates

$$\text{faces:} \quad x_{0bcd} = z_{0,b-1,c-1,d} z'_{0,b,c-1,d-1} z''_{0,b-1,c,d-1}, \tag{3.24}$$

and similarly for the other faces. There are several ways to demonstrate this relation. It is sufficient to do the case $K = 3$, using an explicit parametrization for the ten lines in the 3d K -decomposition, in a chosen basis for V_3 . (Specifying ten lines with the standard incidence relations is equivalent to specifying the configuration of four flags.) Such a parametrization is shown in figure 15, conveniently chosen to encode the four octahedron cross-ratios, called z, w, x, y . Thus, for example

$$z'' = \frac{\langle v_1 \wedge v_2 \rangle \langle v_3 \wedge v_4 \rangle}{\langle v_1 \wedge v_3 \rangle \langle v_2 \wedge v_4 \rangle}, \quad w' = \frac{\langle v_5 \wedge v_6 \rangle \langle v_2 \wedge v_8 \rangle}{\langle v_5 \wedge v_2 \rangle \langle v_6 \wedge v_8 \rangle}, \quad x = \frac{\langle v_9 \wedge v_6 \rangle \langle v_3 \wedge v_7 \rangle}{\langle v_9 \wedge v_3 \rangle \langle v_6 \wedge v_7 \rangle}. \tag{3.25}$$

We can then check that the expected triple-ratio coordinate in the middle of (say) the BCD face, namely

$$x_{0111} = \frac{\langle v_1 \wedge v_2 \wedge v_6 \rangle \langle v_5 \wedge v_6 \wedge v_3 \rangle \langle v_7 \wedge v_3 \wedge v_2 \rangle}{\langle v_1 \wedge v_3 \wedge v_6 \rangle \langle v_5 \wedge v_2 \wedge v_3 \rangle \langle v_7 \wedge v_6 \wedge v_2 \rangle}, \tag{3.26}$$

is equal to the product $xw'z''$.

Finally, let us look at the internal lattice points of the K -decomposition of a tetrahedron: points p_{abcd} with none of a, b, c, d vanishing. (This first occurs for $K = 4$.) Each




Figure 15. Ten vectors parametrizing the ten lines on the $K = 3$ tetrahedron. We have used a $\text{GL}(3, \mathbb{C})$ transformation to put the vectors in the given form; otherwise they are completely generic. We also show the corresponding octahedron parameters (z, w, x, y) , on their corresponding vertices. Here $zz'z'' = -1$ and $z'' + z^{-1} - 1 = 0$, so that $z' = (1 - z)^{-1}$, $z'' = 1 - z^{-1}$, and similarly for w, x, y .

such point coincides with six octahedron vertices, and we could try to assign to the point a coordinate c_{abcd} that is the product of the six surrounding octahedron parameters. This product, however, is automatically trivial, due to a cancellation in the cross-ratios involved. Explicitly:

$$c_{abcd} = z_{a,b-1,c-1,d} z_{a-1,b,c,d-1} z'_{a,b,c-1,d-1} z'_{a-1,b-1,c,d} z''_{a,b-1,c,d-1} z''_{a-1,b,c-1,d} \equiv 1. \quad (3.27)$$

We have shown that a generic configuration of four flags (A, B, C, D) induces octahedron parameters $z_{abcd}, z'_{abcd}, z''_{abcd}$ in the bulk of a tetrahedron that satisfy relations (3.21)–(3.22) and (3.27), and are related to the standard boundary coordinates via (3.23)–(3.24). Conversely, if we start with triples of octahedron parameters satisfying these properties, it is possible to reconstruct uniquely the configuration of four flags. One way to show this is to use the coordinates to reconstruct the array of $\binom{K+2}{3}$ lines on the upright simplices $\Delta^{0,2}$. (We will see explicitly that this can be done in section 3.3.) Identifying the lines is equivalent to specifying the configuration of flags.

Summarizing, there is a canonical bijection between the following two spaces:

1. *The space of octahedron parameters*, given by arbitrary collections of the numbers

$$z_{abcd}, z'_{abcd}, z''_{abcd} \in \mathbb{C}^* - \{1\}, \quad a + b + c + d = K - 2 \quad (3.28)$$

satisfying the relations (3.21)–(3.22) and (3.27),

2. *The subspace of generic configuration of four flags*

$$\text{Conf}_4^*(\mathcal{B}) = \tilde{\mathcal{L}}_K(\Delta, \mathbf{t}_{3d}) = \mathcal{L}_K(\Delta, \mathbf{t}_{3d}) = \mathcal{L}_K(\Delta, \mathbf{t}) \subset \mathcal{X}_K^{\text{un}}(\partial\Delta, \mathbf{t}). \quad (3.29)$$

The Lagrangian pair for the $\text{PGL}(K)$ -tetrahedron is given by (3.29). Note that in this case we again find that the projection to the boundary $\tilde{\mathcal{L}}_K(\Delta, \mathbf{t}_{3d}) \rightarrow \mathcal{L}_K(\Delta, \mathbf{t}_{3d})$ is one-to-one.

3.2.3 The coordinate phase space for a tetrahedron via symplectic reduction

The coordinate symplectic phase space related to the canonical triangulation \mathbf{t} of a tetrahedron is a complex torus $\mathcal{X}_K^{\text{un}}(\partial\Delta, \mathbf{t})$ realized as a Zariski open subset in $\mathcal{X}_K^{\text{un}}(\partial\Delta)$. We can obtain the symplectic coordinate phase space $\mathcal{X}_K^{\text{un}}(\partial\Delta, \mathbf{t})$ as a symplectic reduction of a product of the elementary *octahedron phase spaces* $\mathcal{P}_{\partial\diamond}$ defined in (3.13).

Consider the product of the octahedron phase spaces over the octahedrons of the K -decomposition of a tetrahedron, equipped with the product symplectic structure:

$$\mathcal{P}_\times := \prod_{\diamond_i \in \Delta} \mathcal{P}_{\partial\diamond_i}. \quad (3.30)$$

The edge (3.23) and face (3.24) coordinates on the boundary of the tetrahedron, written in the octahedron parameters, can be viewed as a monomials on the complex torus \mathcal{P}_\times . We call them the edge and face coordinate monomials. The Poisson brackets between them coincide with the Poisson brackets of the corresponding edge and face coordinates for the Atiyah-Bott-Goldman Poisson structure (3.7) on $\mathcal{X}_K^{\text{un}}(\partial\Delta)$. The octahedron relations $zz'z'' = -1$ imply that they satisfy monomial relations given by the unipotent holonomy constraints around each puncture, as in figure 11.

The products (3.27) of six octahedron parameters around the $\binom{K-1}{3}$ internal points of the K -decomposition are no longer equal to one. They are monomial functions c_k on \mathcal{P}_\times . The c_k Poisson-commute with each other, as well as with the boundary coordinates. A nice graphical way to see that the boundary monomials and the c_k have the expected commutation relations is to draw arrows on each (oriented) octahedron, as on the right of figure 13. After gluing together the octahedra, the Poisson bracket between product coordinates on any two points of the K -decomposition is proportional to the net number of arrows going from one to the other. All arrows in the bulk of the tetrahedron cancel out in pairs, while the arrows on the boundary reproduce the configuration in figure 12.

It is also easy to see that the $\binom{K-1}{3}$ monomial functions c_k are all independent. Indeed, since $\text{Conf}_4^*(\mathcal{B})$ has complex dimension $(K-1)^2$ and (from section 3.2.2) $\text{Conf}_4^*(\mathcal{B}) = \tilde{\mathcal{L}}_K(\Delta, \mathbf{t}_{3d}) = (\prod_{\diamond_i} \mathcal{L}_{\diamond_i}) \cap \{c_k\}_{\text{vertices } k}$, there must be at least $(\# \text{ octahedra}) - \dim \text{Conf}_4^*(\mathcal{B}) = \binom{K+1}{3} - (K-1)^2 = \binom{K-1}{3}$ independent c_k 's.

It follows that there is a canonical isomorphism of symplectic spaces

$$\mathcal{X}_K^{\text{un}}(\partial\Delta, \mathbf{t}) = \mathcal{P}_\times // (c_k - 1). \quad (3.31)$$

The notation on the right denotes the symplectic reduction corresponding to the family of commuting moment maps c_k at the level set $c_k = 1$. Each moment map c_k provides an independent Hamiltonian flow that amounts to an action of a copy of the group \mathbb{C}^* . We take the quotient by the resulting action of the complex torus, and then set $c_k = 1$. Note that $\dim \mathcal{P}_\times // (c_k - 1) = 2(K-1)^2$, as expected.

We identify $\mathcal{L}_K(\Delta, \mathbf{t}_{3d})$, the space of generic configurations of four flags from (3.29), as the image of the product Lagrangian $\prod_{\diamond_i \in \Delta} \mathcal{L}_{\diamond_i} \subset \mathcal{P}_\times$ under the symplectic reduction (3.31). It follows that

$$\mathcal{L}_K(\Delta, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial\Delta, \mathbf{t}) \quad (3.32)$$

is Lagrangian as well.




Figure 16. Lattice point on a pair of glued faces.




Figure 17. Hexagonal bipyramid obtained by gluing six tetrahedra along a common edge.

3.2.4 Gluing constraints for 3-manifolds

Now let M be an admissible 3-manifold with an ideal triangulation \mathbf{t}_{3d} and a framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connection. Then we can define the parameters z_i, z'_i, z''_i for every octahedron in the K -decomposition of M .

In section 3.2.2 we found that at every internal point in the K -decomposition of a tetrahedron the product of octahedron parameters is one. We also saw that the product of parameters on a boundary point is equal to the cross-ratio or triple-ratio coordinate in the boundary phase space. Both of these facts generalize. Now the product at every internal point in the K -decomposition of a 3-manifold — including points on faces and edges of glued tetrahedra — is one; while products at external points are equal to big-boundary coordinates. We explain briefly why this is so. (The situation for small torus boundaries is slightly more subtle, and will be deferred to section 3.4.)

First consider a point p on a big internal face in M , shared by two tetrahedra, as in figure 16. As a boundary point for the top tetrahedron, p is assigned a triple-ratio x_{top} formed from flags A, C, E . As a boundary of the bottom tetrahedron, p is also assigned a triple-ratio x_{bottom} formed from A, C, E , but in the opposite orientation. Therefore, $x_{\text{top}}x_{\text{bottom}} = 1$. Moreover, the quantity $x_{\text{top}}x_{\text{bottom}}$ is itself the product of the six octahedron parameters at vertices that touch p , so the product of these octahedron parameters must be one.




Figure 18. *Left:* configuration of octahedra around a point on a trivalent internal edge. *Right:* point on an external edge.

Next, consider points on an internal edge E in M shared by N tetrahedra. We picture this arrangement as an N -gonal bipyramid, as in figure 17, and use the flags at the $N + 2$ vertices to define octahedron parameters throughout the bipyramid.¹¹ Each of the $K - 1$ lattice points on E is touched by N octahedron vertices. The product of the N octahedron cross-ratios at each lattice point involves numerators and denominators that cancel in pairs, and so automatically equals one. For example, a situation with $N = 3$ is shown on the left of figure 18. The three octahedron parameters touching a point on the internal edge are

$$z = \frac{\langle v_4 \wedge v_1 \rangle \langle v_3 \wedge v_2 \rangle}{\langle v_4 \wedge v_3 \rangle \langle v_1 \wedge v_2 \rangle}, \quad w = \frac{\langle v_4 \wedge v_3 \rangle \langle v_5 \wedge v_2 \rangle}{\langle v_4 \wedge v_5 \rangle \langle v_3 \wedge v_2 \rangle}, \quad y = \frac{\langle v_4 \wedge v_5 \rangle \langle v_1 \wedge v_2 \rangle}{\langle v_4 \wedge v_1 \rangle \langle v_5 \wedge v_2 \rangle}, \quad (3.33)$$

so that $zwy = 1$.

Finally, let us consider external lattice points on a big boundary. For a point on a boundary face, we already know from section 3.2.2 that its triple-ratio coordinate is a product of three octahedron parameters. For a boundary edge E , let us again suppose that E is shared by N tetrahedra in the bulk. The boundary cross-ratio parameter x for a point p on E can be computed as a product of octahedron parameters that touch p . Just like in the case of internal edges above, there will be a cancellation between numerators and denominators, but the cancellation is incomplete, and simply leaves behind the single expected cross-ratio x . For example, in the situation depicted on the right of figure 18 (for $N = 2$), the octahedron product is

$$y \cdot w = \frac{\langle v_4 \wedge v_5 \rangle \langle v_1 \wedge v_2 \rangle}{\langle v_4 \wedge v_1 \rangle \langle v_5 \wedge v_2 \rangle} \cdot \frac{\langle v_4 \wedge v_3 \rangle \langle v_5 \wedge v_2 \rangle}{\langle v_4 \wedge v_5 \rangle \langle v_3 \wedge v_2 \rangle} = \frac{\langle v_4 \wedge v_3 \rangle \langle v_1 \wedge v_2 \rangle}{\langle v_4 \wedge v_1 \rangle \langle v_3 \wedge v_2 \rangle}, \quad (3.34)$$

which is independent of v_5 and agrees with the boundary cross-ratio.

Notice that the octahedron parameters constructed from a framed flat connection on M are well-defined and take values in $\mathbb{C}^* - \{1\}$ as long as the configurations of four flags parallel-transported to the interior of every tetrahedron are generic. Conversely, we claim that starting with a collection of octahedron parameters $z_i, z'_i, z''_i \in \mathbb{C}^* - \{1\}$ in the K -decomposition of M that satisfy the octahedron equations (3.21)–(3.22) and the condition that the product at any internal point is one, we can reconstruct a unique framed flat connection on M . So:

¹¹Note that the N tetrahedra forming this bipyramid need not all be distinct. Multiple edges of a single tetrahedron might be glued to E . This does [mc: not] alter the conclusion of the current argument; one just needs to be careful about identifying octahedron cross-ratios with the proper orientations.

Theorem 3.1. *The octahedron parameters satisfying the equations (3.21)–(3.22) and the monomial gluing relations parametrize an open subset $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ of the space of framed flat connections $\tilde{\mathcal{L}}_K(M)$.*

We will prove this in section 3.3.

We would also like to show that the phase space $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ is a symplectic reduction of octahedron phase spaces (3.12). Observe that if we define gluing monomials c_k on the product phase space

$$\mathcal{P}_\times = \prod_{\diamond_i \in M} \mathcal{P}_{\partial \diamond_i} \tag{3.35}$$

as the product of octahedron parameters at every internal lattice point of the K -decomposition of M , then the c_k all Poisson-commute with each other. This follows easily from a cancellation-of-arrows argument as in section 3.2.3. Moreover, for the lattice points on the big boundary of M , we define functions x_j to be the products of octahedron parameters there. Then they satisfy the standard boundary commutation relations and commute with all the c_k , because while octahedron arrows cancel out in the bulk they leave behind the standard arrangements of figure 12 on the boundary. We then expect that there is an equivalence

$$\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}) = \mathcal{P}_\times // (c_k - 1). \tag{3.36}$$

Moreover, the Lagrangian submanifold $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ should be obtained as the image of a product of octahedron Lagrangians $\mathcal{L}_\times = \prod_{\diamond_i \in M} \mathcal{L}_{\diamond_i}$ under the reduction (3.36). We will formalize this in section 4.1.

3.3 Holonomy representation

We would like to show that, given octahedron parameters in the K -decomposition of an admissible manifold M that satisfy the octahedron equations (3.21)–(3.22) as well as the gluing equations at every internal lattice point of the K -decomposition, we can reconstruct a unique framed flat $\text{PGL}(K, \mathbb{C})$ -connection on M — and thus prove theorem 3.1. To do so, we will extend the 2d notion of local projective bases labeled by “snakes” [12] to three dimensions.

We will begin with a framed flat connection on M , and derive the rules for constructing projective bases and the transformations from one projective basis to another (which all involve octahedron parameters). Then we will run the argument in reverse, starting from abstract octahedron parameters and using 3d snakes to reconstruct projective bases and a framed flat connection.

3.3.1 3d snakes in a tetrahedron

We first consider a single ideal tetrahedron and its K -decomposition.

We start with a generic configuration of four flags in V_K . It gives rise to $\binom{K+2}{3}$ lines \mathcal{L}_{abcd} labeled by upright tetrahedra $\Delta_{abcd}^{0,2}$. We are going to construct projective bases in V_K labeled by geometric objects that we call *3d snakes*.

To define a 3d snake, we begin with an oriented path on the edges of the $(K - 1)$ -decomposition of a tetrahedron, which starts at any of the tetrahedron vertices and ends




Figure 19. 3d snakes in the $(K - 1)$ -decomposition of a tetrahedron, from the top vertex to the opposite face (here for $K = 4$).




Figure 20. Correspondence between the K and $(K - 1)$ -decompositions, for $K = 3$. (It is much more natural to draw snakes on the latter.)

at a lattice point on the opposite face, consisting of $K - 1$ segments (the small edges it traverses). We orient the segments toward the opposite face. We decorate each segment of the path by a *fin*. To define a fin observe that a path segment is an edge of a unique $\Delta^{0,2}$ tetrahedron; the fin is a face of that $\Delta^{0,2}$ tetrahedron containing the edge. We visualize it by saying that a fin points from the midpoint of the segment to one of the two unoccupied vertices of the $\Delta^{0,2}$ tetrahedron (figure 19). We define a 3d snake as a path as above whose segments are decorated by fins.

It may help to observe that every $\Delta_{abcd}^{0,2}$ tetrahedron in the K -decomposition corresponds to a lattice point of the $(K - 1)$ -decomposition; thus the vertices of snakes are labeled by lines \mathcal{L}_{abcd} . Moreover, every octahedron in the K -decomposition corresponds to a black $\Delta^{0,2}$ tetrahedron in the $(K - 1)$ -decomposition, and octahedron parameters z, z', z'' label the *edges* of these black $\Delta^{0,2}$ tetrahedra.

The K lines \mathcal{L}_{abcd} that label the vertices of a 3d snake are generic since the original configuration of four flags is generic. Let us associate a basis (v_1, \dots, v_K) to the snake so that the vector v_i lies in the line at the i -th vertex of the snake (numbered from tail to head), and the vector $v_i \pm v_{i+1}$ belongs to the line at the vertex to which the i -th fin points to. Here we take a sum if the fin points to the right and a difference if the fin points to the




Figure 21. Using the i -th segment of a snake to select part of a projective basis. Here there are three lines $\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3$ associated with the vertices of a small black tetrahedron. The indicated vectors must belong to these lines.




Figure 22. The four fundamental snake moves. In each case we go from the orange snake to the green snake. Move IV is shown for $K = 4$.

left, as in figure 21. These conditions determine the basis uniquely, up to overall scaling; thus a snake defines a projective basis. To construct the basis we can choose any v_1 at the tail of the snake, and use the fin relations to determine every subsequent v_i .

There are four elementary moves that can be used to move a 3d snake to any position within a tetrahedron. They are shown in figure 22. Each move acts as a simple $\text{PGL}(K, \mathbb{C})$ transformation on the projective basis associated with a 3d snake. The moves are expressed in terms of two kinds of $\text{GL}(K, \mathbb{C})$ matrices. The first are the elementary unipotent matrices $F_i := \text{Id} + E_{i+1,i}$ (for $i = 1, \dots, K - 1$), where $E_{n,k}$ is the matrix with entry 1 on the n -th row and k -th column and zero elsewhere. The second are the diagonal matrices $H_i(x) = \text{diag}(1, \dots, 1, x, \dots, x)$, with the $(i + 1)$ -st through K -th entries equal to x . Their arguments are always the negatives of octahedron parameters. Observe that F_i and $H_j(x)$ commute unless $i = j$.

Move I can only occur at the head of a snake, adjacent to the face where it ends. It moves the head across a small black tetrahedron, keeping the fins on the same tetrahedron

face, and acts as F_{K-1} , sending

$$\begin{pmatrix} v_1 \\ \vdots \\ v_{K-1} \\ v_K \end{pmatrix} \mapsto F_{K-1} \begin{pmatrix} v_1 \\ \vdots \\ v_{K-1} \\ v_K \end{pmatrix} = \begin{pmatrix} v_1 \\ \vdots \\ v_{K-1} \\ v_{K-1} + v_K \end{pmatrix}. \tag{3.37}$$

(Here we act on the left on column vectors containing the v_i .) Move II is a fin-flip that can occur at any segment along the snake; it simply multiplies v_i (at the head of the segment) and all subsequent vectors by (minus) the octahedron parameter associated with the segment's edge. Move III can act on any two consecutive segments of a snake in the position shown; note that the fins on the second segments are pointing inwards. Move III is a sort of combination of Moves I and II, and involves the octahedron parameter on the edge that the snake slides across.¹² Finally, Move IV is a reversal of orientation that makes sense only when a snake lies entirely along an edge of a big tetrahedron; the fins are unmodified. Move IV reverses the order of the basis vectors, and negates every even-numbered vector, thus multiplying by an anti-diagonal matrix

$$S = \begin{pmatrix} \cdots & 0 & 0 & 1 \\ \cdots & 0 & -1 & 0 \\ \cdots & 1 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \end{pmatrix}. \tag{3.38}$$

The first three moves preserve a flag; i.e. they are lower-triangular matrices. Indeed, any move keeping the tail of a snake attached to a specific vertex of the big tetrahedron preserves the flag associated with that vertex. An example of the sequence of moves that can be used to move a snake from one edge of a big tetrahedron to another (with fins oriented the same way) appears in figure 23. Multiplying the matrices in this example shows that the transformations of 3d snakes along a 2d surface are equivalent to the snake transformations of [12] — given the usual identification of octahedron parameters and boundary cross- and triple-ratios.

The single most important thing to observe about snakes on a tetrahedron is that the projective bases associated with them are well-defined — i.e. the $\mathrm{PGL}(K, \mathbb{C})$ holonomy inside the tetrahedron can be trivialized — *if and only if the octahedron relations $zz'z'' = -1$, $z'' + z^{-1} - 1 = 0$, and the sextic internal gluing constraints are satisfied.*

On one hand, if we derive both octahedron parameters and projective bases by using a configuration of four flags, then the bases are automatically well defined and the relations are automatically satisfied. Conversely, if we simply place a triple of unconstrained coordinates $z_{abcd}, z'_{abcd}, z''_{abcd} \in \mathbb{C}^* - \{1\}$ on every octahedron and *define* the $\mathrm{PGL}(K, \mathbb{C})$ transformations from one snake to another using the rules of figure 22, we will find that the holonomy around a closed path in the space of snakes is trivial if and only if the standard relations are obeyed.

¹²This transformation is perhaps the most nontrivial to derive. It suffices to apply it to snakes on a $K = 3$ tetrahedron, in which case the parametrizations of figure 15 easily lead to the answer.




Figure 23. Moving a snake from one edge to another, for $K = 3$.




Figure 24. Moving a snake around a black tetrahedron to obtain coordinate constraints.

For example, to move a snake all the way around a black tetrahedron with parameters z, z', z'' , as shown in figure 24, we can use a combination of Moves II (fin-flips) and either Moves I or III. Either way, the total holonomy is diagonal aside from a 2×2 block conjugate to

$$\begin{pmatrix} 1 & 0 \\ -zz'(z'' + z^{-1} - 1) & -zz'z'' \end{pmatrix}. \quad (3.39)$$

Having trivial eigenvalues requires $zz'z'' = -1$, and having a trivial holonomy requires $z'' + z^{-1} - 1 = 0$. If $K \geq 4$, then the $(K - 1)$ -decomposition of a tetrahedron contains $\binom{K-1}{3}$ upside-down $\Delta^{2,0}$ tetrahedra in addition to the upright black ones. We may move snakes around these upside-down tetrahedra in a similar way, and find that the holonomy is trivial if and only if the sextic gluing constraints are satisfied — i.e. the product of six octahedron vertices at corresponding internal points in the K -triangulation is one.

Once we know that the $\text{PGL}(K, \mathbb{C})$ holonomy inside a big tetrahedron is trivial, we can use octahedron parameters and snake moves to reconstruct the configuration of framing flags. We first choose *any* projective basis (v_1, \dots, v_K) for \mathbb{C}^K associated with a snake along one of the edges of the tetrahedron. The basis is unique up to $\text{PGL}(K, \mathbb{C})$ isomorphism.




Figure 25. The fifth snake move: jumping between glued tetrahedra, fins along the common face.

Then snake moves determine vectors (hence lines \mathcal{L}_{abcd}) on every other lattice point of the $(K - 1)$ -decomposition, which uniquely reconstruct the configuration of flags.

3.3.2 3d snakes in a 3-manifold

Now we extend 3d snakes and projective bases to an entire triangulated 3-manifold M , and prove that octahedron parameters are coordinates for an open subset of framed flat connections.

We simply need one more elementary move to transport a snake from one tetrahedron to another, across a common face (figure 25). It can be derived by looking at a bipyramid, as in section 3.2.4 (see figure 16). We assume that all fins lie along the common glued face, and don't change position during the jump between tetrahedra. The lines determining the snakes' projective bases are all built from the three flags at the vertices of the common glued face (in figure 16, these are flags A, C, E), but they come with opposite orientation on the two sides. Therefore, the snake move is implemented by a diagonal matrix J of alternating ± 1 's.

We can now run the argument backwards. Suppose that we assign arbitrary parameters

$$z, z', z'' \in \mathbb{C}^* - \{1\}, \quad \text{with} \quad zz'z'' = -1, \quad z'' + z^{-1} - 1 = 0 \quad (3.40)$$

to each octahedron of the K -decomposition of a triangulated 3-manifold M . Consider the set of all snakes in the manifold. We assign to each snake s a standard K -dimensional vector space \mathbb{C}_s^K with a basis. We treat the latter as a projective basis, i.e. consider bases that differ by common \mathbb{C}^* rescaling as equivalent. If two snakes s_1 and s_2 are related by an elementary move of type I-V from figures 22 and 25, we define a linear map $\mathbb{C}_{s_1}^K \rightarrow \mathbb{C}_{s_2}^K$ using the snake move matrices and the octahedral coordinates (3.40). The collection of the vector spaces \mathbb{C}_s^k and the maps between them determines a $\text{PGL}(K, \mathbb{C})$ -connection. The holonomy of this connection inside of a tetrahedra is trivial if and only if the octahedron parameters satisfy the two relations in (3.40), and the sextic gluing constraints for internal lattice points of tetrahedra. Now by moving snakes among different tetrahedra we can show that the $\text{PGL}(K, \mathbb{C})$ holonomy around any contractible cycle for the snakes in M is trivial if and only if the face and edge gluing constraints of section 3.2.4 are also obeyed.

For example, to find the face constraints we can compute the holonomy around the loop shown in figure 26. The loop encircles a lattice point p of the K -decomposition




Figure 26. Moving a snake around a lattice point of the K -triangulation that lies on a pair of glued faces. The four images on the left (orange snakes) are from one tetrahedron, and the images on the right (green snakes) are from a second.

that lies on the common glued face (corresponding to a white triangle in the $(K - 1)$ -triangulation of the face). It is easy to see that corresponding holonomy matrix is $\text{diag}(1, \dots, 1, z''w'xy'v''t, \dots, z''w'xy'v''t)$, with $K - i$ nontrivial entries. For a trivial holonomy the product $z''w'xy'v''t$ of the six octahedron parameters that touch p must equal one — i.e. the face constraint must be satisfied.

In a similar way, we may analyze the internal edge constraints by moving snakes in a circle around that edge. The corresponding holonomy matrix is a product of diagonal fin-flips and J -matrices. Requiring trivial holonomy forces the products of octahedra coordinates at every point on the edge to equal one. To make this explicit, suppose that N tetrahedra are glued around an edge and in each tetrahedron there are $K - 1$ octahedra touching the edge, with parameters $z_i^{(I)}$ (where $1 \leq i \leq K - 1$, $1 \leq I \leq N$). Define $D(x_1, \dots, x_{K-1}) := H_1(x_1)H_2(x_2) \cdots H_{K-1}(x_{K-1}) = \text{diag}(1, x_1, x_1x_2, \dots, x_1 \cdots x_{K-1})$. Then the total snake holonomy around the edge is

$$\prod_{I=1}^N [D(-z_1^{(I)}, \dots, -z_{K-1}^{(I)}) \cdot J] = \prod_{I=1}^N D(z_1^{(I)}, \dots, z_{K-1}^{(I)}) . \quad (3.41)$$

The matrix is trivial if and only if $\prod_{I=1}^N z_i^{(I)} = 1$ for all i . The minus signs that accompany octahedron parameters are cancelled out by the signs in J -matrices!

In combination with loops inside tetrahedra, the face and edge loops just discussed are enough to generate all the homotopically trivial transformations of snakes. Thus, we




Figure 27. The associahedron, and canonical snakes corresponding to three of its 24 vertices.




Figure 28. Connecting glued associahedra.

find that imposing all the standard octahedron gluing constraints allows to define a flat $\mathrm{PGL}(K)$ -connection on the space of snakes, which induces a flat $\mathrm{PGL}(K)$ -connection on the 3-manifold M .

To be more explicit about the last step, we may “doubly truncate” every big tetrahedron Δ in the triangulation of M (figure 27). This polyhedron has 24 vertices. It has hexagons in places of the original vertices and faces, and rectangles in place of the original edges. It is isomorphic to the convex hull of the orbit of a generic point in the space \mathbb{R}^3 , realized as a hyperplane $x_1 + x_2 + x_3 + x_4 = 1$ in \mathbb{R}^4 , under the action of the group S_4 permuting the coordinates (x_1, \dots, x_4) . It is nothing else but the 3d associahedron.¹³

Now, we mark the 24 vertices of each doubly-truncated tetrahedron in the triangulation of M . Each marked point labels one of 24 canonical snakes on the edges of a given tetrahedron, with all fins going along the same tetrahedron face. The location of the point indicates: 1) the edge choice; 2) the snake orientation; and 3) the fins.

At every marked point s in M we fix a standard K -dimensional vector space \mathbb{C}_s^K with a projective basis. The snake holonomy matrices then define a $\mathrm{PGL}(K, \mathbb{C})$ transformation between these vector spaces, determined by the octahedron parameters. Within a single

¹³A similar “cutting” procedure in \mathbb{R}^d leads to a d -dimensional associahedron with $(d + 1)!$ vertices. It is isomorphic to the convex hull of the S_{d+1} -orbit of a generic point in the hyperplane $\mathbb{R}^d \subset \mathbb{R}^{d+1}$. For example, for $d = 2$ we get a hexagon.

tetrahedron we use Moves I–IV to obtain holonomies along the edges of the associahedron, and we know that any closed cycle in the graph has trivial holonomy. We can represent this by filling in the closed cycles with 2-cells. We then extend the graph from one tetrahedron to another by using Move V, and fill in all the closed cycles corresponding to face and edge gluing constraints with 2-cells as well, since these cycles also have trivial holonomy. In the end, we obtain a 2-complex that has the same fundamental group as M , so we have produced a flat $\mathrm{PGL}(K, \mathbb{C})$ -connection on M . Finally, as at the end of section 3.3.1, we use the projective bases assigned to snakes within each tetrahedron to reconstruct the framing flags; and, by construction, the flags at all small boundaries of M will be preserved by the flat $\mathrm{PGL}(K, \mathbb{C})$ -connection. So we arrive at a framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connection on M .

Notice again that the space $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ of octahedron parameters parametrizes framed flat connections for which the 4-tuples of framing flags for each tetrahedron in \mathbf{t}_{3d} are generic. For certain choices of \mathbf{t}_{3d} , the space $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ may lose components. For example, if \mathbf{t}_{3d} includes a tetrahedron glued to itself along two faces, forming a univalent internal edge, the space $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ will be empty. (It is impossible to satisfy the gluing constraints along the univalent edge with octahedron parameters in $\mathbb{C}^* - \{1\}$.) By taking a union over all 3d triangulation \mathbf{t}_{3d} we can define a space $\tilde{\mathcal{L}}_K(M, \mathbf{t})$ that depends only on the boundary triangulation. For any sufficiently refined \mathbf{t}_{3d} , the algebraic closures of $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ and $\tilde{\mathcal{L}}_K(M, \mathbf{t})$ must coincide.

3.4 Small torus boundaries

Let us take a moment to spell out how the boundary phase space $\mathcal{X}_K^{\mathrm{un}}(\partial M, \mathbf{t})$ gets enriched in the presence of small torus boundaries.

Suppose that ∂M consists of a single small torus boundary, and choose A and B cycles for it. The 3d snakes of section 3.3 show us how to express the holonomy eigenvalues around the A and B cycles in terms of octahedron parameters. We compute the full holonomy around a cycle by fixing the tail of a snake on the torus boundary, and swinging the head around the desired cycle (figure 29). Only snake moves I, II, III, and V are used, which leads to a lower-triangular holonomy matrix; this must be so, since the holonomy must preserve the flag on the torus. Therefore, the eigenvalues can simply be read off from the diagonal of the matrix, and take the form of Laurent monomials in the octahedron parameters. A more intuitive and constructive method for finding these monomials will be described in section 5.

For now, we make the following claims. As coordinates on the product phase space $\mathcal{P}_\times = \prod_{\diamond_i \in M} \mathcal{P}_{\partial \diamond_i}$, the holonomy eigenvalues Poisson-commute with all of the internal gluing functions c_k . Moreover, they satisfy the expected Atiyah-Bott-Goldman Poisson bracket for the torus boundary. Explicitly, let us parametrize the $\mathrm{PGL}(K, \mathbb{C})$ eigenvalues as

$$\mathrm{Hol}(\mathrm{A}) = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ & m_1 & 0 & \cdots & 0 \\ & & m_1 m_2 & \cdots & 0 \\ & * & & \ddots & \\ & & & & m_1 m_2 \cdots m_{K-1} \end{pmatrix}, \quad \mathrm{Hol}(\mathrm{B}) = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ & \ell_1 & 0 & \cdots & 0 \\ & & \ell_1 \ell_2 & \cdots & 0 \\ & * & & \ddots & \\ & & & & \ell_1 \ell_2 \cdots \ell_{K-1} \end{pmatrix}, \tag{3.42}$$




Figure 29. Computing the holonomy around a cycle of the torus boundary.

The functions l_a and m_b are the characters of the Cartan group of $\text{PGL}(K)$ provided by the simple positive roots. The Poisson brackets between them are described by the Cartan matrix κ_{ab} :

$$\{l_a, m_b\} = \kappa_{ab} l_a m_b, \quad \{l_a, l_b\} = \{m_a, m_b\} = 0. \quad (3.43)$$

$$\kappa_{ab} = \begin{cases} 2 & a = b \\ -1 & |a - b| = 1 \\ 0 & |a - b| > 1 \end{cases} \quad (3.44)$$

Here we assumed that the intersection number of the A and B cycles on the torus $\langle A, B \rangle = 1$, with the orientation of the torus induced from that of M . Otherwise the formula reads $\{l_a, m_b\} = \langle A, B \rangle \kappa_{ab} l_a m_b$.

We again expect that the boundary phase space associated to a small torus¹⁴ can be expressed as a symplectic quotient of a product of octahedron spaces:

$$\mathcal{X}_K^{\text{un}}(\partial M) \simeq \mathcal{P}_\times // (c_k - 1). \quad (3.45)$$

Notice that here the result of the (linear) symplectic reduction is a complex torus. Thus, we hope to obtain the smooth part $\mathcal{X}_K^{\text{un}}(\partial M) \subset \mathcal{X}_K(T^2)$, isomorphic to $(\mathbb{C}^*)^{2K-2}$, discussed in section 2.3.

We may motivate the symplectic reduction by counting degrees of freedom. An ideal triangulation of M with N tetrahedra has exactly N internal edges (this can be verified by computing the Euler character) and $2N$ internal faces. Therefore, the number of internal gluing functions is

$$\# c_k = N \binom{K-1}{3} + (2N) \binom{K-1}{2} + N \binom{K-1}{1} = N \binom{K+1}{3}. \quad (3.46)$$

This is the same as the number of octahedra in the K -decomposition of M . Thus, if all the c_k 's were independent moment maps, the quotient $\mathcal{P}_\times // (c_k - 1)$ would be zero-dimensional. However, in the presence of a small torus boundary, it turns out that there

¹⁴Although we continue to use the notation “ $\mathcal{X}_K^{\text{un}}(\partial M)$ ” to discuss boundary phase spaces associated with small tori, we emphasize again that we *never* require unipotent holonomy around the cycles of a torus. It is only a condition at the holes in the big boundary of M .

are $K - 1$ automatic relations among the c_k . We will demonstrate this in section 5. Then — assuming that the remaining c_k are independent — $\mathcal{P}_\times // (c_k - 1)$ does have dimension $2(K - 1)$, as required.

If we take the image of the product octahedron Lagrangian \mathcal{L}_\times under the reduction $\mathcal{P}_\times // (c_k - 1)$, then we will obtain a Lagrangian submanifold $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K(T^2)$. It is typically just a component of the variety of framed flat connections on the boundary that extend to the bulk. In particular, the requirement that the framing flags for each big tetrahedron in M be generic (so that octahedron parameters are well defined) implies that $\mathcal{L}_K(M, \mathbf{t}_{3d})$ only parametrizes flat connections whose boundary holonomy has finite stabilizer. Whenever the boundary holonomies have coincident eigenvalues, these eigenvalues must fit in a single Jordan block.

Finally, recall that if ∂M consists of multiple big components \mathcal{C}_i and small torus components T_i^2 , the boundary phase space is just a straightforward product $\mathcal{X}^{\text{un}}(\partial M, \mathbf{t}) = \prod_{i=1}^b \mathcal{X}_K^{\text{un}}(\mathcal{C}_i, \mathbf{t}) \times \prod_{i=1}^t \mathcal{X}_K(T_i^2)$. The expected equivalence $\mathcal{X}_K(\partial M, \mathbf{t}) \simeq \mathcal{P}_\times // (c_k - 1)$ should generalize in the obvious way.

3.5 2–3 moves via 4d cobordism

As a final application of the ideas in this section, we describe how 2–3 moves act on K -decompositions and octahedron parameters.

Recall that a 2–3 Pachner move (shown in figure 1 on page 2) affects a triangular bipyramid in the triangulation of a 3-manifold M . It changes the bulk triangulation of the bipyramid, leaving its boundary intact. Any two bulk triangulations of M can be connected a sequence of 2–3 moves [2, 3]. So if M is an admissible 3-manifold, a 2–3 Pachner move leaves the big-boundary triangulation fixed.

In general, alterations of d -dimensional triangulations (Pachner moves) can be thought of as elementary cobordisms that amount to gluing $(d + 1)$ -dimensional simplices to d -triangulations. For example, in two dimensions, changing a triangulation of a quadrilateral by performing a *flip* is equivalent to gluing a 3d tetrahedron onto the quadrilateral. We should think about the 2-3 move as a cobordism provided by a 4-simplex.

Studying framed flat $\text{PGL}(K)$ -connections on a d -manifold, it is essential to realize that d -simplices are no longer the most elementary pieces of the manifold: we rather have to consider K -decompositions of simplices in a d -dimensional triangulation. Therefore Pachner moves, understood as cobordisms provided by $(d + 1)$ -simplices, should no longer be the most elementary cobordisms. Instead, the elementary cobordisms should reflect the hypersimplices of the K -decomposition of a $(d + 1)$ -simplex. The Pachner move needs to be decomposed into a sequence of such elementary cobordisms.

In two dimensions, this led to the decomposition of a flip into $\binom{K+1}{3}$ elementary mutations. Each mutation corresponded to one of the octahedra $\Delta^{1,1}$ in a 3d tetrahedron [12]. Now we would like to proceed similarly, and decompose a 2–3 move as a sequence of elementary moves that reflect the hypersimplicial K -decomposition of a 4-simplex.

The proof of the invariance of the motivic volume map under the 2–3 moves suggests that the effect of the Pachner 2–3 move on the framed flat $\text{PGL}(K)$ -connections should be presented as a composition of the elementary moves of two types, parametrized by the set of all $\Delta^{1,2}$ - and $\Delta^{2,1}$ -hypersimplices in the K -decomposition of a 4-simplex. This includes




Figure 30. The two elementary 2–3 moves for octahedra, corresponding to cobordisms through $\Delta^{1,2}$ hypersimplices (left) and $\Delta^{2,1}$ hypersimplices (right).

- the $\binom{K+2}{4}$ elementary moves via cobordisms through the $\Delta^{1,2}$ -hypersimplices; and
- the $\binom{K+1}{4}$ elementary moves from cobordisms through $\Delta^{2,1}$ -hypersimplices.

Altogether, these are $\frac{1}{2}K^2(K^2 - 1)$ elementary “2–3 moves” on octahedra $\Delta^{1,1}$ in the K -decomposition.

This is exactly what we find. Topologically, the two types of moves (figure 30) look identical. Indeed, the simplices $\Delta^{1,2}$ and $\Delta^{2,1}$ are isomorphic, both having five 3d octahedra on their boundary. Geometrically, we will see (momentarily) that the first elementary move changes coordinates on the configuration space of five lines in a 2-dimensional space V_2 , while the second move changes coordinates on the space of five planes in a 3-dimensional space V_3 ,

$$\begin{aligned} \Delta^{1,2} &\leftrightarrow \text{Configurations of 5 lines in } V_2 \\ \Delta^{2,1} &\leftrightarrow \text{Configurations of 5 planes in } V_3. \end{aligned} \tag{3.47}$$

These two moduli spaces are canonically isomorphic.

The order in which the 2–3 elementary moves are performed is important, but not unique. Every choice of 3d slicing of a 4d simplex gives a different ordering. One systematic way to organize the elementary moves is shown in figure 31 for $K = 3$ and in figure 32 (below) for $K = 5$. It can be described as follows.

We start with a bipyramid made from two tetrahedra. We start on the “2” side. In the first step, we take the $\binom{K}{2}$ pairs of octahedra that are glued along black faces (at the common face of the two big tetrahedra) and apply the 2–3 move on the left of figure 30 to each pair. This creates $3 \times \binom{K}{2}$ octahedra in a single “equatorial” plane. Among them, there are $\binom{K-1}{2}$ triples of octahedra glued along three white faces and sharing an internal vertex. For the second step, we apply to these triples the move on the right of figure 30, in the $3 \rightarrow 2$ direction. Then we repeat. At the i -th step, for i odd, there should be $\binom{i+1}{2} \times \binom{K-(i-1)/2}{2}$ pairs of octahedra glued along common black faces. They appear in $\frac{i+1}{2}$ horizontal planes. We apply the $2 \rightarrow 3$ move on the left of figure 30 to them. At the i -th step for i even, there should be $\frac{i}{2} \times \binom{K-i/2}{2}$ triples of octahedra glued along three white faces, with an internal trivalent vertex. They appear in $i/2$ horizontal planes. We apply the $3 \rightarrow 2$ move on the right of figure 30 to them. After $2K - 3$ steps we are done: we arrive at a big bipyramid triangulated into three ideal tetrahedra.

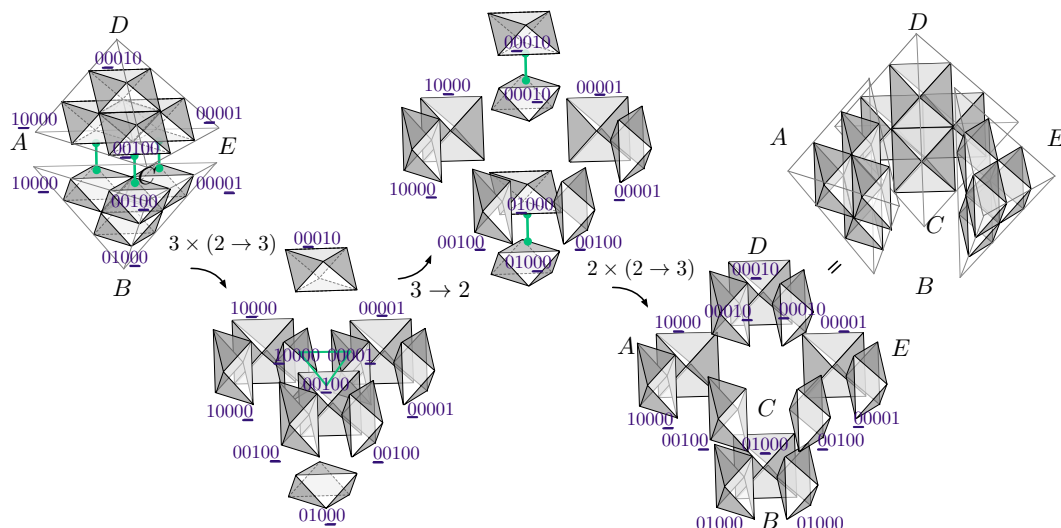


Figure 31. The sequence of elementary 2–3 moves for $K = 3$ (labels are for future reference).

3.5.1 Coordinates

Geometrically, the decomposed 2–3 move gives us many different coordinate systems for a generic configuration of five flags in V_K . They are the five flags (A, B, C, D, E) at the vertices of the bipyramid, or (better yet) the five flags at the vertices of a 4d simplex. Then every full slice of the 4d simplex — i.e. every collection of octahedra that fit together into a bipyramid, at some stage in the sequence of 2–3 moves — provides a complete coordinate system. The octahedra are assigned standard cross-ratio coordinates.

To demonstrate this, let us label 4d hypersimplices by 5-tuples in the usual way,

$$\Delta_{abcde}^{1,2} \rightsquigarrow \mathfrak{P}_{abcde} = A^a \cap B^b \cap C^c \cap D^d \cap E^e, \quad a + b + c + d + e = K - 2, \quad (3.48)$$

$$\Delta_{abcde}^{2,1} \rightsquigarrow \mathfrak{V}_{abcde} = A^a \cap B^b \cap C^c \cap D^d \cap E^e, \quad a + b + c + d + e = K - 3. \quad (3.49)$$

These hypersimplices are associated with planes $\mathfrak{P} \subset V_K$ and 3-spaces $\mathfrak{V} \subset V_K$, built from the flags. Similarly, we may label the $5 \times \binom{K+2}{4}$ octahedra that are involved in the sequence of 2–3 moves by a 5-tuple $abcde$ ($a + b + c + d + e = K - 2$) with one of the five indices “underlined” or “decorated.” The decoration serves (in part) to encode how octahedra occur as boundaries of hypersimplices: the boundaries of $\Delta^{1,2}$ and $\Delta^{2,1}$ contain

$$\partial \Delta_{abcde}^{1,2} : \quad \diamond_{\underline{a}bcde}, \quad \diamond_{a\underline{b}cde}, \quad \diamond_{abc\underline{c}de}, \quad \diamond_{abcd\underline{d}e}, \quad \diamond_{abcde\underline{e}}; \quad (3.50)$$

$$\partial \Delta_{abcde}^{2,1} : \quad \diamond_{(\underline{a+1})bcde}, \quad \diamond_{a(\underline{b+1})cde}, \quad \diamond_{ab(\underline{c+1})de}, \quad \diamond_{abc(\underline{d+1})e}, \quad \diamond_{abcd(\underline{e+1})}. \quad (3.51)$$

Notice that the same octahedron may be on the boundary of *both* $\Delta^{1,2}$ and $\Delta^{2,1}$ hypersimplices. On the other hand, the $5 \times \binom{K+1}{3}$ octahedra that appear at the beginning and end of the full 2–3 move (composing either 2 or 3 complete tetrahedra) are distinguished by the fact that the decorated index vanishes; thus they can only be boundaries of $\Delta^{1,2}$'s.

Every octahedron is associated with a 2-plane, regardless of the decoration, e.g.

$$\diamond_{abcde} \rightsquigarrow \mathfrak{P}_{abcde} = A^a \cap B^b \cap C^c \cap D^d \cap E^e, \quad a + b + c + d + e = K - 2. \quad (3.52)$$




Figure 32. The sequence of elementary 2-3 moves for $K = 5$.

It is *also* associated with four lines, corresponding to the *undecorated* indices,

$$\diamond_{abcde} \rightsquigarrow \mathfrak{L}_{(a+1)bcde}, \mathfrak{L}_{ab(c+1)de}, \mathfrak{L}_{abc(d+1)e}, \mathfrak{L}_{abcd(e+1)}. \quad (3.53)$$

Then the octahedron inherits standard vertex coordinates z, z', z'' formed from cross-ratios of these lines in \mathfrak{P}_{abcde} . It is easy to check that on any slice of the 4-simplex (at any step in the sequence of 2-3 moves) the product of octahedron parameters at any internal lattice point equals one. That is, *the standard gluing equations are obeyed*. By the usual methods, one checks that the octahedron parameters on a slice are coordinates for the configuration space of five flags.

Moreover, if we isolate a cluster of five octahedra involved in one of the elementary moves of figure 30, the products of parameters z, z', z'' on the nine external vertices of this cluster are unchanged during the move. The proof follows from analyzing the cross-ratios involved. This is a local version of the universal statement that boundary coordinates are products of bulk coordinates, no matter how the bulk is decomposed.

It should be clear that the first type of elementary move relates coordinates for five lines in V_2 : the five octahedra on the left of figure 30 carry subsets of four of the five lines

$$\mathfrak{L}_{(a+1)b c d e}, \mathfrak{L}_{a(b+1) c d e}, \mathfrak{L}_{a b(c+1) d e}, \mathfrak{L}_{a b c(d+1) e}, \mathfrak{L}_{a b c d(e+1)} \subset \mathfrak{P}_{a b c d e}, \quad (3.54)$$

and cross-ratios thereof. On the other hand, the second elementary move relates invariants of the configuration of five planes

$$\mathfrak{P}_{(a+1) b c d e}, \mathfrak{P}_{a(b+1) c d e}, \mathfrak{P}_{a b(c+1) d e}, \mathfrak{P}_{a b c(d+1) e}, \mathfrak{P}_{a b c d(e+1)} \subset \mathfrak{V}_{a b c d e}. \quad (3.55)$$

Each of the octahedra in this second move carries the four lines (cf. (3.53)) that are intersections of one of these five planes with the other four! For example, $\diamond_{a b c d e}$ carries the intersection of $\mathfrak{P}_{(a+1) b c d e}$ with the other planes, and the cross-ratios formed from the resulting lines.

The moduli spaces of generic configurations of five lines in V_2 and five planes in V_3 are canonically isomorphic, or better to say canonically dual to each other (see the end of section 2.6). This is why the elementary moves on octahedra look so similar. The “ K -decomposition” of a full 2–3 move into elementary moves on octahedra is a much refined version of the collection of the “pentagons” assigned to a configuration of five flags in section 2.6.

Finally, we note that the symplectic constructions of phase spaces and Lagrangian submanifolds that we have discussed throughout this section commute beautifully with 2–3 moves — simply because they commute with the elementary moves on octahedra. For example, once we know that $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ is a symplectic reduction of a product phase space \mathcal{P}_\times for any given triangulation \mathbf{t}_{3d} , it follows that it can also be obtained by symplectic reduction for any other triangulation \mathbf{t}'_{3d} , or indeed for any “partial” triangulation obtained by doing some series of elementary 2–3 moves on individual octahedra. (Note that the 2–3 moves never change the boundary triangulation \mathbf{t} !) We will say more about this in section 5.5. Moreover, the elementary 2–3 moves induce birational maps on the spaces $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ and $\mathcal{L}_K(M, \mathbf{t}_{3d})$. The maps are one-to-one so long as the configurations of five flags involved in the move are generic.

4 K_2 -Lagrangians and the Bloch complex

In this section and the next, we study the symplectic nature of gluing, and the Lagrangian properties of the moduli space $\mathcal{L}_K(M)$ of framed flat connections on ∂M that can be extended to M .

4.1 Perspective

The basic result we are after can be phrased in several equivalent ways. Let M be an admissible 3-manifold with a triangulation \mathbf{t}_{3d} , inducing a triangulation \mathbf{t} of the big boundary. Fix K , and let N denote the number of octahedra in the K -decomposition of \mathbf{t}_{3d} . In section 3, we introduced the three moduli spaces

$$\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d}) \rightarrow \mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}), \tag{4.1}$$

and argued that they could be constructed by “gluing” the elementary Lagrangian pairs

$$\begin{array}{ccc} \mathcal{L}_{\diamond_i} & \subset & \mathcal{P}_{\partial \diamond_i} \\ \parallel & & \parallel \\ \{z''_i + z_i^{-1} = 1\} & \subset & \{z_i z'_i z''_i = -1\} \simeq (\mathbb{C}^*)^2 \end{array} \tag{4.2}$$

associated to individual octahedra. Let $2d$ be the dimension of $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$, so that $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}) \simeq (\mathbb{C}^*)^{2d}$. What we have shown so far is that

- There exists a non-canonical projection $\tilde{\pi}$ from $\prod_{i=1}^N \mathcal{P}_{\partial \diamond_i} \simeq (\mathbb{C}^*)^{2N}$ to $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$. In particular, the coordinates x_a on $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ are monomials of the octahedron parameters z_i, z'_i, z''_i .
- $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ is the subvariety in the product Lagrangian $\prod_{i=1}^N \mathcal{L}_{\diamond_i} \subset (\mathbb{C}^*)^{2N}$ cut out by gluing equations $\{c_j = 1\}_{j=1}^r$, where c_j are monomials of the octahedron parameters.
- There is a canonical projection $\pi : \prod_{i=1}^N \mathcal{P}_{\partial \diamond_i} \cap \{c_j = 1\}_{j=1}^r \rightarrow \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$.
- $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ is the image of $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ under either the projection $\tilde{\pi}$ or π .

Taking a union of $\tilde{\mathcal{L}}_K(M, \mathbf{t}_{3d})$ (or $\mathcal{L}_K(M, \mathbf{t}_{3d})$) over finitely many 3d triangulations, or by sufficiently refining the triangulation, we lose the dependence on \mathbf{t}_{3d} , keeping only the dependence on big-boundary triangulation \mathbf{t} .

This description of gluing, however, ignores symplectic structures. The spaces $\mathcal{P}_{\partial \diamond_i}$ and $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ have natural holomorphic symplectic structures, given by $\Omega_{\partial \diamond_i} = d \log z_i \wedge d \log z'_i$ and the canonical 2-form Ω on the space of framed flat unipotent connections, respectively. As anticipated in section 3, we propose that

Conjecture 4.1 (Symplectic Gluing). *The symplectic space $(\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}), \Omega)$ is isomorphic to a holomorphic symplectic quotient of $(\prod_{i=1}^N \mathcal{P}_{\partial \diamond_i}, \sum_i \Omega_{\partial \diamond_i})$ for the Hamiltonian action of the group $(\mathbb{C}^*)^{N-d}$, whose Hamiltonians are the gluing monomials $\{c_j\}$. Thus*

$$\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}) \simeq \prod_{i=1}^N \mathcal{P}_{\partial \diamond_i} // (\mathbb{C}^*)^{N-d} = [\prod_{i=1}^N \mathcal{P}_{\partial \diamond_i} \cap (c_j = 1)] // (\mathbb{C}^*)^{N-d}. \tag{4.3}$$

(Sometimes we denote this quotient as $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t}) \simeq \prod_{i=1}^N \mathcal{P}_{\partial \diamond_i} // (c_j = 1)$.)

This is equivalent to saying that there exist monomials γ_j in the octahedron parameters (possibly with fractional exponents) such that the holomorphic symplectic forms on $\mathcal{X}_K(\partial M, \mathbf{t})$ and the octahedron space $\prod_{i=1}^N \mathcal{P}_{\partial\phi_i}$ are related as

$$\sum_{i=1}^N d \log z_i \wedge d \log z_i'' = \Omega + \sum_j d \log c_j \wedge d \log \gamma_j. \tag{4.4}$$

In sections 4.2.3 and 4.3 we consider the Milnor group K_2 of the fields of functions on $\prod_{i=1}^N \mathcal{P}_{\partial\phi_i}$ and $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$, and K_2 -avatars \mathcal{W} of the symplectic forms Ω . Relation (4.4) between the symplectic forms is upgraded to a relation between their K_2 -avatars:

$$\sum_{i=1}^N z_i \wedge z_i'' = \mathcal{W} + \sum_j c_j \wedge \gamma_j. \tag{4.5}$$

Since the $\mathcal{L}(M, \mathbf{t}_{3d})$ is the projection of $\tilde{\mathcal{L}}(M, \mathbf{t}_{3d}) = \prod_{i=1}^N \mathcal{L}_{\phi_i} \cap (c_j = 1)$ to $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$, we re-interpret $\mathcal{L}(M, \mathbf{t}_{3d})$ as the image of the product Lagrangian $\prod_{i=1}^N \mathcal{L}_{\phi_i} \subset \prod_{i=1}^N \mathcal{P}_{\partial\phi_i}$ under the symplectic quotient (4.3). It follows that

Corollary 4.1. *Conjecture 4.1 implies that $\mathcal{L}(M, \mathbf{t}_{3d})$ is a K_2 -Lagrangian, and hence a holomorphic Lagrangian, subvariety of $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$. In other words $\mathcal{L}(M, \mathbf{t}_{3d})$ is d -dimensional and the restrictions of \mathcal{W} , and hence Ω , to $\mathcal{L}(M, \mathbf{t}_{3d})$ vanish.*

Since the K_2 -class and hence the symplectic structures are preserved by 2–3 moves, the refined $\mathcal{L}(M, \mathbf{t})$ that depends only on big-boundary triangulation is K_2 -Lagrangian as well.

In order to quantize $\mathcal{L}(M, \mathbf{t})$ and to construct the gauge theories $T_K[M]$, we will a slightly generalized version of Conjecture 4.1, involving logarithmic lifts of the various \mathbb{C}^* coordinates. We discuss this in section 5.

We will not prove Conjecture 4.1 in full generality in this paper. We will, however, prove some special cases, and provide insight into the general case.

In the rest of this section, we assume that the boundary ∂M consists entirely of big boundary, with holes filled in by small discs (no small annuli or tori). In this case, we use the algebraic approach of [66], involving a canonical homomorphism from the complex of configurations of decorated flags to the Bloch complex, to prove that $\mathcal{L}(M, \mathbf{t}_{3d})$ is a K_2 -isotropic variety of $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ — meaning that the restriction of \mathcal{W} to $\mathcal{L}(M, \mathbf{t}_{3d})$ vanishes, but that $\mathcal{L}(M, \mathbf{t}_{3d})$ is not necessarily mid-dimensional. In section 5 and appendix B we will use a combinatorial/topological approach to prove a weaker version of Conjecture 4.1 for any admissible M , showing that $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ is a symplectic reduction of a multiplicative subspace of $\prod_{i=1}^N \mathcal{P}_{\partial\phi_i}$ — not necessarily of the entire space; see Proposition 5.1 on page 82.¹⁵

¹⁵Since the first version of this paper appeared, Garoufalidis and Zickert [71] have adapted the combinatorial approach to prove Conjecture 4.1 in the case that ∂M consists entirely of small tori. Similar arguments appeared in [80].

The canonical homomorphism from the complex of configurations of decorated flags to the Bloch complex, which we are about to review in detail, has the following form:

$$\begin{array}{ccccccc}
 \dots & \xrightarrow{d} & A_4^{(K)} & \xrightarrow{d} & A_3^{(K)} & \xrightarrow{d} & A_2^{(K)} \\
 & & \downarrow 0 & & \downarrow \alpha_4 & & \downarrow \alpha_3 \\
 \dots & \longrightarrow & 0 & \longrightarrow & B_2(F) & \xrightarrow{\delta} & \wedge^2(F^*) .
 \end{array} \tag{4.6}$$

Here d is the standard simplicial differential. For instance, it maps a configuration of four flags in $A_3^{(K)}$ (thought of as located at the four vertices of a tetrahedron) to a signed sum over configurations of three flags in $A_2^{(K)}$ (the triangles on the boundary of the tetrahedron). The map δ sends an element $\{x\}_2$ of the Bloch group (satisfying the five-term relation) to $(1-x) \wedge x$. The maps α_\bullet are easy to describe for $K = 2$. In this case, a flag is just a line and the map α_4 is given by the cross-ratio of a configuration of four vectors

$$\alpha_4(l_1, l_2, l_3, l_4) := \frac{\langle l_1 \wedge l_2 \rangle \langle l_3 \wedge l_4 \rangle}{\langle l_2 \wedge l_3 \rangle \langle l_4 \wedge l_1 \rangle} . \tag{4.7}$$

On the other hand, α_3 is given by the following element of $\Lambda^2 \mathbb{C}^*$:

$$\alpha_3(l_1, l_2, l_3) := \langle l_1 \wedge l_2 \rangle \wedge \langle l_2 \wedge l_3 \rangle + \langle l_2 \wedge l_3 \rangle \wedge \langle l_3 \wedge l_1 \rangle + \langle l_3 \wedge l_1 \rangle \wedge \langle l_1 \wedge l_2 \rangle . \tag{4.8}$$

(To relate this to (3.4) or (3.10), set $(l_1, l_2, l_3, l_4) = (a, b, d, c)$.) The homomorphism of complex α_\bullet for $K > 2$ is defined by using the hypersimplicial K -decomposition and the related biGrassmannian complex.

If ∂M has just the big boundary, we use (4.6) to prove that $\mathcal{L}(M, \mathbf{t}_{3d})$ is K_2 -isotropic via the following argument. First, the configurations of decorated flags parameterizing generic points of $\tilde{\mathcal{L}}(M, \mathbf{t}_{3d})$ give rise to elements of $A_3^{(K)}$ in (4.6); while the configurations parametrizing generic points of $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ relate to elements of $A_2^{(K)}$. The projection of $\tilde{\mathcal{L}}(M, \mathbf{t}_{3d})$ to $\mathcal{L}(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ is reflected by the map $d : A_3^{(K)} \rightarrow A_2^{(K)}$. Moving to the bottom row of (4.6), let F be the field of rational functions on $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$. The group $K_2(F)$ is the cokernel of δ , i.e. $K_2(F) = \wedge^2(F^*)/\delta(B_2(F))$. The K_2 -class \mathcal{W} of the symplectic form on $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ was constructed in [12, section 15]. By the very definition, its pre-image in $\wedge^2(F^*)$ is obtained as follows: we apply the map α_3 to the sum, over all triangles in the triangulation \mathbf{t} of ∂M , of the elements of $A_2^{(K)}$ corresponding to the configurations of decorated flags assigned to the triangles. The restriction of \mathcal{W} to $\mathcal{L}(M, \mathbf{t}_{3d})$ is precisely the image under the composition $\alpha_3 \circ d : A_3^{(K)} \rightarrow \wedge^2(F^*)$ of the element in $A_3^{(K)}$ describing the generic point of the space $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$. By the commutativity of the right square in (4.6), this can be rewritten as the image under the map $\alpha_3 \circ d = \delta \circ \alpha_4$. Since $K_2(F)$ is the cokernel of δ , this implies that the K_2 -class \mathcal{W} must restrict to zero on $\mathcal{L}_K(M, \mathbf{t}_{3d})$. Similar considerations also lead to a formula for the differential of the volume in a family of framed flat connections.

Let us now present the details.

4.2 From decorated flags to the Bloch complex

4.2.1 Decorated flags

Recall that a *decorated flag* \tilde{F}^\bullet (often called an affine flag) in V_K is a flag F^\bullet equipped with additional data: a choice of a non-zero vector $f^i \in F^{i-1}/F^i$ for each $i = 1, \dots, K$. Denote

by \mathcal{A}_K the space of all decorated flags in V_K . It is the principal affine space for $\mathrm{GL}(K)$. Denote by $\mathrm{Conf}_{n+1}^*(\mathcal{A}_K)$ the space of generic configurations of $n + 1$ decorated flags.

Denote by $\mathrm{Conf}_m^*(k)$ the space of generic configurations of m vectors in a vector space of dimension k . Given a partition \mathbf{a} as in (2.20), or, equivalently, a hypersimplex $\Delta_{\mathbf{a}}^{p,q}$ of the hypersimplicial K -decomposition of an n -simplex, we define a canonical projection

$$\mathcal{C}_{\mathbf{a}} : \mathrm{Conf}_{n+1}^*(\mathcal{A}_K) \longrightarrow \mathrm{Conf}_{n+1}^*(p+1). \quad (4.9)$$

Namely, take a generic configuration of decorated flags $(\tilde{F}_0^\bullet, \dots, \tilde{F}_n^\bullet)$ in V_K . Intersecting the $(p+1)$ -dimensional subspace

$$\mathfrak{F}_{\mathbf{a}} := F_0^{a_0} \cap \dots \cap F_n^{a_n} \quad (4.10)$$

with the “previous” flag subspace in each of the flags, we get a configuration of $n + 1$ codimension-one subspaces (hyperplanes) in this vector space:

$$\mathfrak{H}_k := F_k^{a_k+1} \cap \mathfrak{F}_{\mathbf{a}} \subset \mathfrak{F}_{\mathbf{a}}, \quad k = 0, \dots, n. \quad (4.11)$$

So far we just recovered the configuration of $n + 1$ hyperplanes in (2.22). Now we use the decorations to produce a configuration of $n + 1$ vectors in $(\mathfrak{F}_{\mathbf{a}})^*$. The decoration $f_{a_k} \in F_k^{a_k}/F_k^{a_k+1}$ determines a unique linear functional l_k on $F_k^{a_k}/F_k^{a_k+1}$ taking the value 1 on the vector f_{a_k} . It provides a linear functional on the $(p+1)$ -space $\mathfrak{F}_{\mathbf{a}}$ annihilating the hyperplane \mathfrak{H}_k . Altogether, the linear functionals $l_0, \dots, l_n \in (\mathfrak{F}_{\mathbf{a}})^*$ provide a configuration $(l_0, \dots, l_n) \in \mathrm{Conf}_{n+1}^*(p+1)$. So we set

$$\mathcal{C}_{\mathbf{a}}(\tilde{F}_0^\bullet, \dots, \tilde{F}_n^\bullet) := (l_0, \dots, l_n). \quad (4.12)$$

Definition 4.1. *Given a generic configuration of $n + 1$ decorated flags $(\tilde{F}_0^\bullet, \dots, \tilde{F}_n^\bullet)$ in V_K , we assign to a hypersimplex $\Delta_{\mathbf{a}}^{p,q}$ the configuration (4.12) of $n + 1$ vectors in a $(p + 1)$ -dimensional vector space:*

$$\Delta_{\mathbf{a}}^{p,q} \longmapsto \mathcal{C}_{\mathbf{a}}(\tilde{F}_0^\bullet, \dots, \tilde{F}_n^\bullet) \in \mathrm{Conf}_{n+1}^*(p+1), \quad \mathbf{a} = (a_0, \dots, a_n). \quad (4.13)$$

Its projectivization is the configuration of $n + 1$ points in a p -dimensional projective space from section 2.6.

Remark: We use the convention dual to the one in [66]. Notice that a decorated flag in a vector space V_K determines the dual decorated flag in the dual space V_K^* . To match the two conventions, replace the decorated flags by the dual decorated flags. So intersections of flag subspaces here correspond to quotients by direct sums of flag subspaces in [66].

4.2.2 BiGrassmannian complex

Our next goal is to relate, when n varies, different configurations of $n + 1$ vectors assigned to

- generic configurations of $(n + 1)$ decorated flags in V_K , and
- the collections of hypersimplices of the K -decompositions of an n -simplex.

Given a configuration (l_1, \dots, l_m) in V_n , we can use two operations to get a new configuration of $m - 1$ vectors:

1. Forgetting the i -th vector l_i , we get a map

$$f_i : \text{Conf}_m^*(n) \longrightarrow \text{Conf}_{m-1}^*(n), \quad (l_1, \dots, l_m) \longmapsto (l_1, \dots, \widehat{l_i}, \dots, l_m). \quad (4.14)$$

2. Projecting the vectors (l_1, \dots, l_m) to the quotient $V_n/(l_j)$ along the one-dimensional subspace spanned by l_j , we get a map

$$p_j : \text{Conf}_m^*(n) \longrightarrow \text{Conf}_{m-1}^*(n-1), \quad (l_1, \dots, l_m) \longmapsto (\bar{l}_1, \dots, \widehat{l_j}, \dots, \bar{l}_m). \quad (4.15)$$

Here \bar{l}_s stands for the projection of the vector l_s to the quotient.

Following [66], we organize the spaces $\text{Conf}_m^*(n)$ into a single object, the *biGrassmannian*:

$$\begin{array}{ccccccc}
 & & & & \cdots & \xrightarrow{\quad} & \mathbf{G}_3^0 \\
 & & & & & \downarrow \vdots \downarrow & \downarrow \vdots \downarrow \\
 & & & & \cdots & \xrightarrow{\quad} & \mathbf{G}_2^1 \xrightarrow{\quad} \mathbf{G}_2^0 \\
 & & & & & \downarrow \vdots \downarrow & \downarrow \vdots \downarrow \\
 & & & & \cdots & \xrightarrow{\quad} & \mathbf{G}_1^2 \xrightarrow{\quad} \mathbf{G}_1^1 \xrightarrow{\quad} \mathbf{G}_1^0 \\
 & & & & & \downarrow \vdots \downarrow & \downarrow \vdots \downarrow \\
 & & & & \cdots & \xrightarrow{\quad} & \mathbf{G}_0^3 \xrightarrow{\quad} \mathbf{G}_0^2 \xrightarrow{\quad} \mathbf{G}_0^1 \xrightarrow{\quad} \mathbf{G}_0^0
 \end{array} \quad (4.16)$$

The name is explained by the canonical isomorphism

$$\text{Conf}_{p+q+2}^*(p+1) = \mathbf{G}_p^q \quad (4.17)$$

with the Grassmannian \mathbf{G}_p^q parametrising the $(p + 1)$ -dimensional subspaces in a vector space of dimension $p + q + 2$ with a given basis (e_1, \dots, e_{p+q+2}) , which are in generic position to the coordinate hyperplanes. Namely, a generic $(p + 1)$ -plane π gives rise to a configuration of vectors in π^* given by the restrictions of the coordinate linear functional x_i dual to the basis. In the diagram (4.16), given a Grassmannian \mathbf{G}_p^q , there are $p + q + 2$ horizontal arrows for the maps f_i , and $p + q + 2$ vertical arrows for the maps p_j . Notice that the hypersimplex $\Delta^{p,q}$ match the Grassmannian \mathbf{G}_p^q , as the notation suggest. The arrows f_i and p_j match the two kinds of boundaries of the hypersimplex $\Delta^{p,q}$ discussed in section 2.5.

There is a canonical isomorphism

$$\mathbf{G}_p^q = \mathbf{G}_q^p. \quad (4.18)$$

It assigns to a $(p + 1)$ -dimensional subspace π in a coordinate vector space of dimension $p + q + 2$ its annihilator π^\perp in the dual vector space, equipped with the dual basis.

Given a set X , we assign to it a free abelian group $\mathbb{Z}[X]$ with the generators $\{x\}$ parametrized by the elements of X . Applying this construction to the sets of complex points of the configuration spaces (identified with Grassmannians), we arrive at abelian groups

$$C_m(n) := \mathbb{Z}[\text{Conf}_m^*(n)(\mathbb{C})]. \tag{4.19}$$

They are organized into the *Grassmannian bicomplex*

$$\begin{array}{ccccccc}
 & & & & \cdots & \xrightarrow{f} & C_5(4) \\
 & & & & & \downarrow p & \downarrow p \\
 & & & \cdots & \xrightarrow{f} & C_5(3) & \xrightarrow{f} & C_4(3) \\
 & & & & \downarrow p & \downarrow p & \downarrow p & \\
 \cdots & \xrightarrow{f} & C_5(2) & \xrightarrow{f} & C_4(2) & \xrightarrow{f} & C_3(2) & \\
 & \downarrow p & \downarrow p & \downarrow p & \downarrow p & & & \\
 \cdots & \xrightarrow{f} & C_5(1) & \xrightarrow{f} & C_4(1) & \xrightarrow{f} & C_3(1) & \xrightarrow{f} & C_2(1)
 \end{array} \tag{4.20}$$

Here the maps f and p are the alternating sums of the maps f_i , and p_j :

$$f = \sum_{i=0}^m (-1)^i f_i, \quad p = \sum_{j=0}^m (-1)^j p_j. \tag{4.21}$$

Denote by BC_* the sum of the groups on the diagonals, shifting the index m by one:

$$BC_{m-1} := \bigoplus_{n=1}^{m-1} C_m(n). \tag{4.22}$$

As usual, we change the sign of the differentials f in all odd columns, and then define the total differential $D : BC_m \rightarrow BC_{m-1}$ by $D := f + p$. Then $D^2 = 0$ (thanks to the sign change). We get the *biGrassmannian complex*

$$\cdots \xrightarrow{D} BC_4 \xrightarrow{D} BC_3 \xrightarrow{D} BC_2 \xrightarrow{D} BC_1. \tag{4.23}$$

Finally, consider the complex of generic configurations of decorated flags in V_K :

$$\cdots \rightarrow A_4^{(K)} \rightarrow A_3^{(K)} \rightarrow A_2^{(K)} \rightarrow A_1^{(K)}, \quad A_m^{(K)} := \mathbb{Z}[\text{Conf}_{m+1}^*(\mathcal{A}_K)(\mathbb{C})] \tag{4.24}$$

with the standard simplicial differential

$$(\tilde{F}_0, \dots, \tilde{F}_n) \rightarrow \sum_{s=0}^n (-1)^s (\tilde{F}_0, \dots, \widehat{\tilde{F}_s}, \dots, \tilde{F}_n). \tag{4.25}$$

Let us now recall the main construction of [66, section 2].

Definition 4.2. *Given a configuration $(\tilde{F}_0, \dots, \tilde{F}_n)$ of $n + 1$ generic decorated flags in V_K , assigned to the vertices of the n -dimensional simplex, take the hypersimplicial K -decomposition Δ_K^n of the simplex, assign to every hypersimplex $\Delta_{\mathbf{a}}^{p,q}$ in Δ_K^n a configuration*

of vectors $\mathcal{C}_{\mathbf{a}}(\tilde{F}_0, \dots, \tilde{F}_n)$, thought of as a point of the Grassmannian \mathbf{G}_p^q , and take the sum over all hypersimplices:

$$c_n : (\tilde{F}_0, \dots, \tilde{F}_n) \longrightarrow \sum_{\mathbf{a}} \mathcal{C}_{\mathbf{a}}(\tilde{F}_0, \dots, \tilde{F}_n) \in BC_n. \quad (4.26)$$

Finally, extend the map to a homomorphism of abelian groups

$$c_n : A_n^{(K)} \longrightarrow BC_n. \quad (4.27)$$

The following crucial result was proved in Key Lemma 2.1 from [66].

Theorem 4.1. *The collection of maps c_n gives rise to a homomorphism of complexes*

$$\begin{array}{ccccccc} \dots & \longrightarrow & A_4^{(K)} & \longrightarrow & A_3^{(K)} & \longrightarrow & A_2^{(K)} & \longrightarrow & A_1^{(K)} \\ & & \downarrow c_4 & & \downarrow c_3 & & \downarrow c_2 & & \downarrow c_1 \\ \dots & \longrightarrow & BC_4 & \longrightarrow & BC_3 & \longrightarrow & BC_2 & \longrightarrow & BC_1 \end{array} \quad (4.28)$$

Proof. We claim that, given a configuration of decorated flags $(\tilde{F}_0, \dots, \tilde{F}_n)$, the map

$$\mathcal{C} : \Delta_{\mathbf{a}}^{p,q} \longrightarrow \mathcal{C}_{\mathbf{a}}(\tilde{F}_0, \dots, \tilde{F}_n) \quad (4.29)$$

that assigns to a hypersimplex $\Delta_{\mathbf{a}}^{p,q}$ the corresponding configuration of vectors intertwines the operation of taking the boundary of the hypersimplex with the differential in the Grassmannian bicomplex. This would immediately imply the theorem since we apply the map \mathcal{C} to the sum of all hypersimplices of the decomposition of a simplex Δ_K^n , and so the total boundary is given by the total boundary $d\Delta_K^n$ of the simplex. The latter is given by the alternating sum of the $n+1$ boundary $(n-1)$ -dimensional simplices, which matches the differential in the complex of generic configurations of decorated flags.

Let us now formulate the claim precisely. Recall that a hypersimplex $\Delta^{p,q}$ has two kinds of boundaries, call them ∂ - and ∂' -boundaries, one consisting of $(n+1)$ hypersimplices of type $\Delta^{p-1,q}$, and another consisting of $(n+1)$ hypersimplices of type $\Delta^{p,q-1}$. So, the boundary $d\Delta_{\mathbf{a}}^{p,q}$ of the hypersimplex $\Delta_{\mathbf{a}}^{p,q}$ can be written, taking into account the orientations, as

$$d[\Delta_{\mathbf{a}}^{p,q}] = \sum_{i=0}^n (-1)^i [\Delta_{\partial_i \mathbf{a}}^{p-1,q}] + \sum_{j=0}^n (-1)^j [\Delta_{\partial'_j \mathbf{a}}^{p,q-1}]. \quad (4.30)$$

Here $\partial_i \mathbf{a}$ and $\partial'_j \mathbf{a}$ are the partitions which label the boundary hypersimplices. We assert that taking the ∂_i -boundary (resp. ∂'_j -boundary) of a hypersimplex $\Delta_{\mathbf{a}}^{p,q}$, and then assigning to the obtained hypersimplices the configurations of n vectors we get the same result as first assigning to the hypersimplex $\Delta_{\mathbf{a}}^{p,q}$ the configuration of $(n+1)$ vectors, and then applying the map f_i (resp. p_j). Precisely, there are two commutative diagrams:

$$\begin{array}{ccc} \Delta_{\mathbf{a}}^{p,q} & \xrightarrow{\partial_i} & \Delta_{\partial_i \mathbf{a}}^{p-1,q} & & \Delta_{\mathbf{a}}^{p,q} & \xrightarrow{\partial'_j} & \Delta_{\partial'_j \mathbf{a}}^{p,q-1} \\ \downarrow \mathcal{C} & & \downarrow \mathcal{C} & & \downarrow \mathcal{C} & & \downarrow \mathcal{C} \\ \mathcal{C}_{\mathbf{a}}(\tilde{F}_0, \dots, \tilde{F}_n) & \xrightarrow{p_i} & \mathcal{C}_{\partial_i \mathbf{a}}(\tilde{F}_0, \dots, \tilde{F}_n) & & \mathcal{C}_{\mathbf{a}}(\tilde{F}_0, \dots, \tilde{F}_n) & \xrightarrow{f_j} & \mathcal{C}_{\partial'_j \mathbf{a}}(\tilde{F}_0, \dots, \tilde{F}_n) \end{array} \quad (4.31)$$

The commutativity of these two diagrams is straightforward to verify, and the theorem follows immediately from this.

4.2.3 Bloch complex

Let F be a field. Recall first the definition of the Bloch group $B_2(F)$ [81]. Denote by $\mathbb{Z}[F^* - \{1\}]$ the free abelian group generated by elements $\{z\}$ assigned to $z \in F^* - \{1\}$. Let us define a subgroup $R_2(F) \subset \mathbb{Z}[F^* - \{1\}]$ as follows. Take any generic configuration of 5 points (z_0, \dots, z_4) on $\mathbb{P}^1(F)$. Then the subgroup $R_2(F)$ is generated by the five-term relations, where r is the cross-ratio:

$$\sum_{i=0}^4 (-1)^i \{r(z_0, \dots, \widehat{z}_i, \dots, z_4)\} \subset \mathbb{Z}[F^* - \{1\}]. \tag{4.32}$$

Definition 4.3. *The Bloch group $B_2(F)$ is the quotient of the abelian group $\mathbb{Z}[F^* - \{1\}]$ by the subgroup of the five-term relations:*

$$B_2(F) := \frac{\mathbb{Z}[F^* - \{1\}]}{R_2(F)}. \tag{4.33}$$

We denote by $\{z\}_2$ the projection of the generator $\{z\}$ to the Bloch group.

Next, there is a well-known homomorphism of abelian groups

$$\delta : B_2(F) \longrightarrow F^* \wedge F^*, \quad \{x\}_2 \longmapsto (1-x) \wedge x. \tag{4.34}$$

The group $K_2(F)$ is the cokernel of this map

$$K_2(F) := F^* \wedge F^* / \delta(B_2(F)). \tag{4.35}$$

Let us construct a canonical map of complexes [67, sections 2.3, 2.4]:

$$\begin{array}{ccccccc} \cdots & \longrightarrow & BC_4 & \longrightarrow & BC_3 & \longrightarrow & BC_2 \\ & & \downarrow b_5 & & \downarrow b_4 & & \downarrow b_3 \\ \cdots & \longrightarrow & R_2(\mathbb{C}) & \longrightarrow & \mathbb{Z}[F^* - \{1\}] & \xrightarrow{\delta} & \Lambda^2 F^* \end{array} \tag{4.36}$$

We define it by looking at the Grassmannian bicomplex (4.20), and defining the map row by row.

- First, we send the bottom row $\cdots \rightarrow C_3(1) \rightarrow C_2(1)$ to zero.
- The map on the next row is the crucial part of the construction. It amounts to a construction of the map of complexes

$$\begin{array}{ccccccc} \cdots & \longrightarrow & C_5(2) & \longrightarrow & C_4(2) & \longrightarrow & C_3(2) \\ & & \downarrow \gamma_5 & & \downarrow \gamma_4 & & \downarrow \gamma_3 \\ \cdots & \longrightarrow & R_2(F) & \longrightarrow & \mathbb{Z}[F^* - \{1\}] & \xrightarrow{\delta} & \Lambda^2 F^* \end{array} \tag{4.37}$$

- Given a configuration of vectors (l_1, l_2, l_3) in a two-dimensional space V_2 , set

$$\gamma_3(l_1, l_2, l_3) := \langle l_1 \wedge l_2 \rangle \wedge \langle l_2 \wedge l_3 \rangle + \langle l_2 \wedge l_3 \rangle \wedge \langle l_3 \wedge l_1 \rangle + \langle l_3 \wedge l_1 \rangle \wedge \langle l_1 \wedge l_2 \rangle \in \mathbb{C}^* \wedge \mathbb{C}^* .$$
(4.38)

Here to define $\langle l_1 \wedge l_2 \rangle$ we choose a volume form $\Omega \in \det V_2^*$ and set $\langle l_1 \wedge l_2 \rangle := (\Omega, l_1 \wedge l_2)$, as in (3.3) or (3.4). The invariant $\gamma_3(l_1, l_2, l_3)$ does not depend on the choice of volume form Ω .

- The map γ_4 is given by the cross-ratio:

$$\gamma_4(l_1, l_2, l_3, l_4) = r(l_1, l_2, l_3, l_4) := \frac{\langle l_1 \wedge l_2 \rangle \langle l_3 \wedge l_4 \rangle}{\langle l_2 \wedge l_3 \rangle \langle l_4 \wedge l_1 \rangle} .$$
(4.39)

- The map γ_5 assigns to a configuration $(l_1, l_2, l_3, l_4, l_5)$ the corresponding configuration of five points on the projective line, which by the very definition is an element of R_2 .

- Finally, we define a map

$$C_5(3) \longrightarrow R_2$$

by assigning to a configuration of five vectors in V_3 the corresponding projective configuration of five points, and then identifying them with five points on a line as in the end of 2.6.

- All the other maps are set to zero.

To prove that we get a map of complexes one shows ([67], section 2.3) that the right square in (4.37) commutes. The rest is obvious. We arrive at a homomorphism of complexes

$$\begin{array}{ccccccc} \dots & \longrightarrow & BC_4 & \longrightarrow & BC_3 & \longrightarrow & BC_2 \\ & & \downarrow & & \downarrow b_4 & & \downarrow b_3 \\ \dots & \longrightarrow & 0 & \longrightarrow & B_2(F) & \xrightarrow{\delta} & \Lambda^2 F^* \end{array}$$
(4.40)

4.2.4 Conclusion: from decorated flags complex to the Bloch complex

Combining the homomorphism (4.28) from the complex of decorated flags to the biGrassmannian complex with the homomorphism (4.36) from the biGrassmannian complex to the Bloch complex, we arrive at our final output — a homomorphism α_\bullet from the complex of decorated flags in V_K to the Bloch complex:

$$\begin{array}{ccccccc} \dots & \longrightarrow & A_4^{(K)} & \longrightarrow & A_3^{(K)} & \longrightarrow & A_2^{(K)} \\ & & \downarrow \alpha_5 & & \downarrow \alpha_4 & & \downarrow \alpha_3 \\ \dots & \longrightarrow & R_2(F) & \longrightarrow & \mathbb{Z}[F^* - \{1\}] & \xrightarrow{\delta} & \Lambda^2 F^* \end{array}$$
(4.41)

Of course, the complex $R_2(F) \longrightarrow \mathbb{Z}[F^* - \{1\}] \xrightarrow{\delta} \Lambda^2 F^*$ is quasi-isomorphic to the Bloch complex, so we can cast the map (4.41) in a more traditional form, where we abuse notation slightly by denoting the middle map by α_4 again:

$$\begin{array}{ccccccc} \dots & \longrightarrow & A_4^{(K)} & \longrightarrow & A_3^{(K)} & \longrightarrow & A_2^{(K)} \\ & & \downarrow 0 & & \downarrow \alpha_4 & & \downarrow \alpha_3 \\ \dots & \longrightarrow & 0 & \longrightarrow & B_2(F) & \xrightarrow{\delta} & \Lambda^2 F^* \end{array}$$
(4.42)

This resulting map of complexes is the main ingredient of the construction of the second motivic Chern class in [66], which is a class

$$C_2^{\mathcal{M}} \in H^4(BGL_K(\mathbb{C}), \mathbb{Z}_{\mathcal{M}}(2)). \tag{4.43}$$

In section 4.3–4.4 we show that the K_2 -Lagrangian property and the variation formula for the volume of a framed flat $PGL(K, \mathbb{C})$ -connection on an admissible 3-manifold in the absence of the small toric and annular boundaries follow immediately from these homological considerations.

4.3 $\mathcal{L}_K(M)$ is a K_2 -isotropic subvariety

The symplectic structure on the moduli space $\mathcal{X}_K^{\text{un}}(\mathcal{C})$ of unipotent framed flat $PGL(K, \mathbb{C})$ -connections on a punctured surface \mathcal{C} is given by a 2-form Ω , the Atiyah-Bott-Goldman form. The form Ω has been identified in [12, section 15] as the image under the $d \log \wedge d \log$ map of a class \mathcal{W} in K_2 of the space $\mathcal{X}_K^{\text{un}}(\mathcal{C})$, called the *motivic avatar of the symplectic form*:

$$\Omega = (d \log \wedge d \log)(\mathcal{W}). \tag{4.44}$$

Recall that $K_2(F)$ is a quotient of $\Lambda^2 F^*$. To define the class \mathcal{W} in K_2 , a $\Lambda^2 F^*$ -representative was assigned to every ideal triangulation \mathbf{t} of \mathcal{C} in [12],

$$W_K(\mathcal{C}; \mathbf{t}) \in \mathbb{C} \left(\mathcal{X}_K^{\text{un}}(\mathcal{C}) \right)^* \wedge \mathbb{C} \left(\mathcal{X}_K^{\text{un}}(\mathcal{C}) \right)^*. \tag{4.45}$$

Its image in K_2 does not depend on the choice of triangulation, and defines the class \mathcal{W} . Specialising it to a generic unipotent framed flat $PGL(K, \mathbb{C})$ -connection $\mathcal{V}_{\mathcal{C}}$ on \mathcal{C} we get an element

$$W_K(\mathcal{V}_{\mathcal{C}}; \mathbf{t}) \in \mathbb{C}^* \wedge \mathbb{C}^*. \tag{4.46}$$

Let us recall the definition of $W_K(\mathcal{V}_{\mathcal{C}}; \mathbf{t})$. Pick an invariant decorated flag \tilde{F}_v near each puncture v on \mathcal{C} . So each triangle t of the ideal triangulation \mathbf{t} with the vertices $v_1(t), v_2(t), v_3(t)$ gives rise to a configuration of three decorated flags $(\tilde{F}_{v_1(t)}, \tilde{F}_{v_2(t)}, \tilde{F}_{v_3(t)})$. We get the element $W_K(\mathcal{V}_{\mathcal{C}}; \mathbf{t})$ by applying the map α_3 from (4.41) to it, and taking the sum over all triangles:

$$W_K(\mathcal{V}_{\mathcal{C}}; \mathbf{t}) := \sum_{t \in \mathbf{t}} \alpha_3(\tilde{F}_{v_1(t)}, \tilde{F}_{v_2(t)}, \tilde{F}_{v_3(t)}). \tag{4.47}$$

(The definition given in [12, section 15] is equivalent to this one, and in fact was originally obtained by calculating (4.47).)

Assume that M is an admissible 3-manifold without small torus or annulus components. Taking the sum over all big boundary components \mathcal{C}_i of ∂M , we set

$$W_K(\partial M; \mathbf{t}) := \sum_i W_K(\mathcal{C}_i; \mathbf{t}).$$

Now let us run a similar procedure on an admissible 3-manifold M with an ideal triangulation \mathbf{t}_{3d} . The moduli spaces $\mathcal{L}_K(M)$ and $\tilde{\mathcal{L}}_K(M)$ may have several irreducible components.

Definition 4.4. A framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connection \mathcal{V} on M is generic for a triangulation \mathfrak{t}_{3d} if its restriction to any tetrahedron Δ of the triangulation \mathfrak{t}_{3d} provides a generic configuration of four flags $(F_1(\Delta), \dots, F_4(\Delta))$.

The \mathcal{V} is generic if there exists a triangulation \mathfrak{t}_{3d} of M for which it is generic.

A component of $\tilde{\mathcal{L}}_K(M)$ is generic if it contains a generic framed flat connection.

Being generic is an open condition: framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connections close to a generic one are also generic. This is why it is enough to have a single generic one in a component.

Examples. 1. Let M be a polyhedron. So it is a ball whose boundary is a sphere, punctured in n points. Then the moduli space $\tilde{\mathcal{L}}_K(M)$ has a single irreducible component, and this component is generic. Indeed, by the very definition it parametrizes arbitrary configurations of n flags in a K -dimensional vector space. And evidently $\tilde{\mathcal{L}}_K(M) = \mathcal{L}_K(M)$.

2. Let M be a hyperbolic threefold with big boundary. Then it has an Epstein-Penner triangulation into ideal tetrahedra. Notice that a configuration of four points on the projective line is generic if all the points are distinct. So an ideal tetrahedron with distinct vertices, even a flat one, provides a generic configuration of points on $\mathbb{C}P^1$. Therefore the framed flat $\mathrm{PGL}(2, \mathbb{C})$ -bundle describing the hyperbolic structure is generic. So the component of $\tilde{\mathcal{L}}_K(M)$ containing the hyperbolic structures on M is generic. We call it the hyperbolic component.

3. Take the standard embedding $\mathrm{PGL}(2) \rightarrow \mathrm{PGL}(K)$ induced by the linear representation of $\mathrm{SL}(2)$ in polynomials in two variables of degree $K - 1$. Then any hyperbolic threefold M gives rise to a framed flat $\mathrm{PGL}(K, \mathbb{C})$ -bundle on M : it is obtained by taking the framed flat $\mathrm{PGL}(2, \mathbb{C})$ -connection on M assigned to the hyperbolic structure, and using the standard embedding above. The component of $\tilde{\mathcal{L}}_K(M)$ containing this representation is generic. Here is a proof. A flag in a two-dimensional vector space V_2 , that is a linear subspace $L \subset V_2$, determines a flag F_L in $S^{K-1}V_2$, given by the nested subspaces $\{L^m S^{K-m-1}(V_2)\}$, $m = K - 1, \dots, 0$. It is the unique flag fixed by the Borel subgroup of PGL_2 stabilizing the L . Then four distinct subspaces (L_1, \dots, L_4) give rise to a generic configuration of four flags $(F_{L_1}, \dots, F_{L_4})$. To see this, one can argue geometrically. For example, if $K = 3$, consider a conic in $P^2 = P(S^2(V_2))$ given by the projectivisation of the set of decomposable vectors v^2 . Take a point l on the conic. Then the flag given by the point l and the tangent line to the conic at the point l is the flag F_L . Evidently, given four distinct points (l_1, \dots, l_4) on the conic, the associated configuration of flags $(F_{L_1}, \dots, F_{L_4})$ is generic.

Take a *generic* framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connection \mathcal{V} on M . Given a tetrahedron Δ from a triangulation \mathfrak{t}_{3d} of M , and applying the map α_4 to the configuration of four flags assigned to Δ , we get

$$\alpha_4(F_1(\Delta), \dots, F_4(\Delta)) = \sum_{\mathbf{a}} \{z_{\mathbf{a}}\}_2 \in B_2(\mathbb{C}). \tag{4.48}$$

Here $\{z_a\}$ are the octahedron parameters. Taking the sum over all tetrahedra Δ_i of the triangulation, we get a $B_2(\mathbb{C})$ -valued invariant:

$$\alpha_4(\mathcal{V}; \mathbf{t}_{3d}) := \sum_i \alpha_4(F_1(\Delta_i), \dots, F_4(\Delta_i)) \in B_2(\mathbb{C}). \tag{4.49}$$

Theorem 4.2. *Let M be an admissible 3-fold without small torus or annulus components. Let $\mathcal{V} \in \tilde{\mathcal{L}}_K(M)$ be a generic framed flat $PGL_K(\mathbb{C})$ -connection on M , and $\text{Res}_{\partial M}(\mathcal{V})$ its restriction to the boundary.*

(i) *If a triangulation \mathbf{t} of ∂M is the boundary of a triangulation \mathbf{t}_{3d} of M , then*

$$W_K(\text{Res}_{\partial M}(\mathcal{V}); \mathbf{t}) = \delta \circ \alpha_4(\mathcal{V}; \mathbf{t}_{3d}). \tag{4.50}$$

So the restriction of the K_2 -class $W_K(\partial M)$ to the component of $\mathcal{L}_K(M)$ containing \mathcal{V} is zero. Therefore the component of $\mathcal{L}_K(M)$ containing \mathcal{V} is K_2 -isotropic.

(ii) *If M is a convex polyhedron, then $\mathcal{L}_K(M)$ is a K_2 -Lagrangian subvariety of $\mathcal{X}_K^{\text{un}}(\partial M)$.*

(iii) *The element $\alpha_4(\mathcal{V}; \mathbf{t}_{3d}) \in B_2(\mathbb{C})$ does not depend on the choice of a triangulation \mathbf{t}_{3d} , although it does depend on the boundary triangulation \mathbf{t} . Therefore we get an invariant of $(\mathcal{V}, \mathbf{t})$, which we call the motivic volume invariant:*

$$\alpha_4(\mathcal{V}; \mathbf{t}) \in B_2(\mathbb{C}). \tag{4.51}$$

Proof. (i) The second claim follows from the first, since K_2 is the cokernel of δ . The first follows from the commutativity of the right square in (4.41). Indeed, calculating $\delta \circ \alpha_4(\mathcal{V})$ we pick a tetrahedron of the ideal triangulation of M , and calculate the map α_4 followed by δ for the corresponding configuration of four decorated flags. This means taking these four flags and going down and to the right in (4.41). Thanks to the commutativity of the right square in (4.41), we get the same result by applying the differential to the configuration of four decorated flags (going to the right in (4.41)), and then apply the map α_3 (going down). After taking the sum over all tetrahedra of the ideal triangulation of M , the contributions of the internal big faces of the tetrahedra cancel pairwise, and the contribution of the other faces match the definition of $W_K(\text{Res}_{\partial M}(\mathcal{V}); \mathbf{t})$ given by (4.47).

(ii) This was proved right after formulation of the theorem 4.2 in the Introduction.

(iii) It is clear from (4.41) that the element $\alpha_4(\mathcal{V}; \mathbf{t}_{3d}) \in B_2(\mathbb{C})$ does not change under the 2–3 moves. Indeed, the change is described by the composition $A_4^{(K)} \rightarrow A_3^{(K)} \rightarrow B_2(\mathbb{C})$, which is zero by (4.42). So it does not depend on the choice of the 3d triangulation \mathbf{t}_{3d} of M .

The theorem is proved. □

Small boundaries: here is how to extend formula (4.50) when M has small toric boundary components. On each of the toric components T_i^2 of the boundary, the symplectic form is naturally lifted to a $\Lambda^2\mathbb{C}^*$ -invariant $W_K(T_i^2)$. Namely, choose a pair of generators A, B of $H_1(T^2, \mathbb{Z})$ with intersection number $\langle A, B \rangle = 1$. Present the diagonal parts of the

holonomies of a flat $\mathrm{PGL}(K, \mathbb{C})$ -connection \mathcal{V} on the torus around the A- and B-cycles as in (2.13), using the values ℓ_a, m_b , where $a, b = 1, \dots, K-1$, of the positive simple roots on the Cartan group. Then

$$W_K(T^2)(\mathcal{V}) := \sum_{a,b=1}^{K-1} (\kappa^{-1})_{ab} \ell_a \wedge m_b \in \mathbb{C}^* \wedge \mathbb{C}^*. \quad (4.52)$$

Here κ_{ab} is the Cartan matrix. See sections 3.4 and 5.3, where we elaborate on the Poisson structure for the symplectic form defined by this class.

The Poisson brackets for eigenvalues derived combinatorially in appendix B are almost sufficient to imply (4.50) for toric boundary. See section 5.3 for further remarks on this. Formula (4.50) can also be generalized to admissible 3-manifolds with small annular boundaries (or, equivalently, laminations), but this will be discussed elsewhere (cf. [15] for $K = 2$).

Remark: another way to get theorem 4.2 is to use the result of [12, section 10] that a flip of a 2d triangulation on a surface \mathcal{C} amounts to a cluster transformation on the moduli space $\mathcal{A}_{\mathrm{SL}(K), \mathcal{C}}$. Indeed, it was shown in [17, section 6], that any cluster transformation preserves the class $\mathcal{W}_{\mathcal{A}}$ in K_2 of a cluster \mathcal{A} -variety, and that the difference of the $\Lambda^2 \mathbb{C}^*$ -invariants $\mathcal{W}_{\mathcal{A}}(\mathbf{s}_1) - \mathcal{W}_{\mathcal{A}}(\mathbf{s}_2)$ assigned to two seeds \mathbf{s}_1 and \mathbf{s}_2 related by a cluster transformation is

$$\mathcal{W}_{\mathcal{A}}(\mathbf{s}_1) - \mathcal{W}_{\mathcal{A}}(\mathbf{s}_2) = \sum_{i=1}^n (1 - z_i) \wedge z_i.$$

Here the cluster transformation $\mathbf{s}_1 \rightarrow \mathbf{s}_2$ is written as a composition of n mutations, and the parameters z_1, \dots, z_n are the \mathcal{X} -coordinates related to these mutations. In the case of a flip of triangulation, these are the octahedron parameters $\{z_{\mathbf{a}}\}$ in (4.48). These two approaches have the same origin: the homomorphism of complexes (4.36), which was in fact used in [12] to obtain the cluster picture just outlined. However, the map α_4 has S_4 skew-symmetry built in its definition. To see it by decomposing a flip into a composition of mutations is a non-trivial problem.

4.4 Differential of the volume of a framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connection

Let G be a complex semi-simple Lie group. Let M be a closed manifold and \mathcal{V} a G -bundle with flat connection on M . The flat G -bundle is described by a map to the classifying space BG of the group G made discrete:

$$\varphi_{\mathcal{V}} : M \longrightarrow BG^{\delta}. \quad (4.53)$$

It is well defined modulo homotopy. It maps the fundamental class $[M]$ of M to a class $\varphi_{\mathcal{V}}[M] \in H_3(BG^{\delta}, \mathbb{Z})$. Combined with the regulator map

$$r_2 : H_3(BG^{\delta}, \mathbb{Z}) \longrightarrow \mathbb{C}/(2\pi i)^2 \mathbb{Z} \quad (4.54)$$

we get an invariant of a G -bundle with flat connection \mathcal{V} on M :

$$c_2(\mathcal{V}) := r_2(\varphi_{\mathcal{V}}[M]) \in \mathbb{C}/(2\pi i)^2 \mathbb{Z}. \quad (4.55)$$

In particular, taking the imaginary part of $c_2(\mathcal{V})$ we recover the volume invariant of \mathcal{V} . The name is suggested by the representations to $\mathrm{PGL}(2, \mathbb{C})$ provided by hyperbolic 3-manifolds M , when it is the hyperbolic volume. All this is well-known.

Now let M have a boundary. Then, although the fundamental class of M does not exist, for a unipotent G -vector bundle with flat connection \mathcal{V} on M one can still define an invariant

$$c_2(\mathcal{V}) \in \mathbb{C}/(2\pi i)^2\mathbb{Z} \tag{4.56}$$

using the “relative fundamental class” construction from [82, section 3], or [26].

For generic *non-unipotent* local systems on an admissible 3-manifold M , the $B_2(\mathbb{C})$ -valued invariant (4.51) is a motivic refinement of the volume invariant. In fact it amounts to the pullback of the explicit construction of the second motivic Chern class (4.43).

To recover the \mathbb{R} -valued volume one applies to (4.51) the Bloch-Wigner dilogarithm

$$\mathcal{L}i_2(z) := \mathrm{Im}(\mathrm{Li}_2(z) + \log(1 - z) \log |z|). \tag{4.57}$$

It satisfies the five-term relation, which just means that it provides a homomorphism

$$\mathcal{L}i_2 : B_2(\mathbb{C}) \mapsto \mathbb{R}, \quad \{z\}_2 \mapsto \mathcal{L}i_2(z). \tag{4.58}$$

Notice that even in the simplest possible case when M is the tetrahedron and thus the flat connection itself is trivial, the volume is a nontrivial invariant. In this particular case it is nothing else but the measurable 3-cocycle of $\mathrm{PGL}(K, \mathbb{C})$ defined in [83]: given a flag F ,

$$(g_1, \dots, g_4) \in \mathrm{PGL}(K, \mathbb{C})^4 \mapsto \mathcal{L}i_2(\alpha_4(g_1F, \dots, g_4F)). \tag{4.59}$$

Let us show now how formula (4.50) leads to a variation of the volume formula.

Given a family $\mathcal{V}(t)$ of framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connections on an admissible 3-manifold M , depending on a parameter t , we would like to calculate the differential of the $\mathrm{PGL}(K, \mathbb{C})$ -volume Vol_K in this family.

Let X be a complex algebraic variety. There is a homomorphism to the space $\mathcal{A}^1(X)$ of real valued 1-forms, which are smooth at the generic point of X :

$$\log d \arg : \mathbb{C}(X)^* \wedge \mathbb{C}(X)^* \longrightarrow \mathcal{A}^1(X), \quad f \wedge g \longmapsto \log |f| d \arg g - \log |g| d \arg f. \tag{4.60}$$

The differential of the Bloch-Wigner dilogarithm can be written as a composition of the map δ with the map $\log d \arg$:

$$d\mathcal{L}i_2(z) := \log |1 - z| d \arg z - \log |z| d \arg(1 - z) = \log d \arg \circ \delta \{z\}_2. \tag{4.61}$$

This shows that to calculate $\sum_i d\mathcal{L}i_2(z_i)$ one needs to calculate $\delta \sum_i \{z_i\}_2 = \sum_i (1 - z_i) \wedge z_i$.

Let us return to the calculation of the differential of the volume. Assume that M is as in theorem 4.2. Choose an ideal triangulation \mathbf{t} of the big boundary. As explained above,

$$d\mathrm{Vol}_K(\mathcal{V}(t); \mathbf{t}) = d\mathcal{L}i_2(\alpha_4(\mathcal{V}(t); \mathbf{t})) = \log d \arg \circ \delta \circ \alpha_4(\mathcal{V}(t); \mathbf{t}). \tag{4.62}$$

Applying the $\log d \arg$ map to the right hand side of (4.50), we get the differential of the volume:

$$d\mathrm{Vol}_K(\mathcal{V}(t); \mathbf{t}) = \log d \arg \left(W_K(\mathrm{Res}_{\partial M}(\mathcal{V}(t)); \mathbf{t}) \right). \tag{4.63}$$

So to calculate the differential $d\text{Vol}_K(\mathcal{V}(t))$, we restrict the family $\mathcal{V}(t)$ to the boundary ∂M , and then evaluate the universal 1-form $\log d \arg \left(W_K(\text{Res}_{\partial M}(\mathcal{V}(t)); \mathbf{t}) \right)$ on the resulting family of the framed flat connections on the boundary.

We conclude that although the motivic volume of a framed flat $\text{PGL}(K, \mathbb{C})$ -connection on M is determined by the connection on M , the differential of the volume is calculated using the class \mathcal{W} only.

We emphasize that the motivic volume invariant does depend on the choice of the boundary triangulation \mathbf{t} of M . In some situations there is a preferred triangulation, e.g. in the case when M is a hyperbolic manifold with big geodesic boundary. So the motivic volume of a generic framed flat connection on a hyperbolic manifold is an invariant.

5 Combinatorics of K -decompositions

In this section we introduce a complementary graphical approach for describing the combinatorics of coordinates for framed flat $\text{PGL}(K)$ -connections that were developed in section 3. We will use an algebra of paths on slices in the K -decomposition of a 3-manifold M to encode the relationship between octahedron parameters on one hand, and gluing functions and boundary coordinates on the other. This generalizes the work of [70] in the context of hyperbolic geometry. The approach is especially useful for describing holonomy eigenvalues around small torus boundaries and analyzing their Poisson brackets. In turn, the Poisson brackets lead to a proof of a weak version of the Symplectic Gluing Conjecture 4.1 (Proposition 5.1, page 82). It seems promising that with additional effort the combinatorial analysis here could be upgraded to produce a full proof of the Conjecture.

The algebra of paths also has a second important advantage: it allows us define a consistent logarithmic lift of all the phase space coordinates, promoting them to additive \mathbb{C} -valued (rather than \mathbb{C}^* -valued) functions. Using \mathbb{C}^* -coordinates as in section 3, we know that all gluing functions c_k and boundary coordinates are monomials in the octahedron parameters z_i, z'_i, z''_i . Assuming Conjecture 4.1, we could encode the exponents of these monomials in a symplectic matrix g , which holds all the necessary information for expressing a patch of $\mathcal{X}_K^{\text{un}}(\partial M)$ as a symplectic reduction of the product \mathcal{P}_\times of octahedron phase spaces, as well as reproducing the Lagrangian submanifold $\mathcal{L}_K(M) \subset \mathcal{X}_K^{\text{un}}(\partial M)$. However, this information is not sufficient for quantization or for defining theories $T_K[M]$. A consistent quantization (e.g. in the sense of [9]) requires keeping track of an extra vector of integers σ that translate to q -corrections in the quantum Lagrangian operators $\hat{\mathcal{L}}_K(M)$, as well as in any quantum wavefunctions. In terms of 3d $\mathcal{N} = 2$ theories $T_K[M]$, the extra vector of integers encodes a consistent UV R-charge assignment, which is necessary for compactification of $T_K[M]$ on any curved backgrounds. The formal algebra of path coordinates in this section is used to specify, unambiguously, the vector σ .

The consistent definition of path coordinates fundamentally requires all small boundary components of M to admit a Euclidean 2d structure. Indeed, the path coordinates will look like Euclidean angle structures on slices of M parallel to the small boundary. This consideration motivated our original definition of admissible manifolds, back in section 2.1.




Figure 33. K -decompositions of an ideal tetrahedron, for $K = 2, 3, 4$.




Figure 34. Labeling of vertices.

Ultimately, we will also want to choose polarizations Π and Π_\times for the logarithmic lifts of phase spaces $\mathcal{X}_K^{\text{un}}(\partial M)$ and the product octahedron space \mathcal{P}_\times . Then the gluing data g and σ should be combined into an affine symplectic transformation $(g, \sigma) \in \text{ISp}(2N, \mathbb{Z})$, or sometimes $\text{ISp}(2N, \mathbb{Q})$ (where N is the number of octahedra), which relates the two polarizations and is the fundamental piece of information needed to define $T_K[M]$ and any quantum wavefunctions associated with M .

Much of this section can be read without knowledge of the framed flat connections of sections 2–3. In this sense it is a combinatorial guide to producing the data for quantization and for theories $T_K[M]$.

In this section, we will suppress the dependence of various moduli spaces (coordinate phase spaces and Lagrangians) on 2d triangulations \mathbf{t} and 3d triangulations \mathbf{t}_{3d} . We assume throughout that \mathbf{t} and \mathbf{t}_{3d} are fixed. We use the notation $\mathcal{P}_K(\partial M)$ to denote the logarithmic version of coordinate phase spaces $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$, and use uppercase variables (e.g. Z, Z', X, \dots) to denote the logarithmic \mathbb{C} -valued lifts of the corresponding \mathbb{C}^* -valued parameters (e.g. $z = \exp Z, z' = \exp Z', x = \exp X, \dots$). (We abuse notation, continuing to use $\mathcal{P}_{\partial\Diamond}$ for the canonical logarithmic lifts of octahedron phase spaces; we also will not distinguish between Lagrangian submanifolds $\mathcal{L}_K(M, \mathbf{t}_{3d}) \subset \mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ and their preimages in $\mathcal{P}_K(\partial M)$.)

5.1 Octahedra and slices

Recall that an admissible 3-manifold M can be decomposed into ideal tetrahedra, each of which has a further K -decomposition. In particular, the K -decomposition cuts an ideal tetrahedron into $\binom{K+1}{3} = \frac{1}{6}K(K^2 - 1)$ octahedra \Diamond_i , which are our central players.

Each octahedron is oriented and its faces can be canonically assigned two distinct colors, according to whether they align with big faces of a tetrahedron (black) or with the truncations of vertices (white). The colors alternate on neighboring faces. We label the six vertices of every octahedron with formal complex parameters Z_i, Z'_i, Z''_i , equal on opposite vertices and occurring in a clockwise order around every black face. (Thus




Figure 35. Four parallel sets of slices in the $K = 4$ decomposition.




Figure 36. Continuing slices from one tetrahedron to another; the sets of slices on the right here are associated with vertices A , B , and C , respectively.

there are three distinct ways to label any octahedron, related by a cyclic permutation $Z_i \rightarrow Z'_i \rightarrow Z''_i \rightarrow Z_i$.)

Notice that the decomposition into octahedra naturally identifies four families of $K - 1$ parallel *slices* of every tetrahedron Δ , as shown in figure 35. Each family of slices is centered around one of the four vertices (a small boundary component). Moreover, every slice is densely tiled by black and white faces of octahedra: white faces from octahedra below the slice and black faces from octahedra above the slice. The three dihedral angles of these small black and white triangles can all be labeled by the vertex parameters Z_i, Z'_i, Z''_i of the octahedra that they came from. We have attempted to indicate this in figure 35, for the four triangles associated with the top octahedron.

When ideal tetrahedra are glued together to form M , the slices can be continued from one tetrahedron to the next, as indicated in figure 36. Indeed, every small boundary component of ∂M acquires a family of $K - 1$ global slices throughout M centered around that small boundary and having the same topology. For example, a small torus boundary comes with $K - 1$ toroidal slices in the K -decomposition.

5.2 Coordinates on the product phase space

We now want to revisit the symplectic gluing of elementary Lagrangian pairs for octahedra to form the Lagrangian pair for an admissible 3-manifold M . We will assume that the boundary ∂M does not contain small annuli (i.e. it only has small tori and the small

discs that fill in holes on the big boundary). The extension to small annuli follows in a straightforward way by combining results of [15] with the present analysis.

Suppose that M is triangulated into N ideal tetrahedra, so that the K -decomposition has $\frac{1}{6}NK(K^2 - 1)$ total octahedra. To each octahedron \diamond_i we associate a formal logarithmic phase space

$$\mathcal{P}_{\partial \diamond_i} = \{(Z_i, Z'_i, Z''_i) \mid Z_i + Z'_i + Z''_i = i\pi\}, \tag{5.1}$$

with holomorphic symplectic form $\Omega_{\partial \diamond_i} = dZ_i \wedge dZ'_i$, or equivalently with Poisson bracket

$$\{Z_i, Z'_i\} = \{Z'_i, Z''_i\} = \{Z''_i, Z_i\} = 1. \tag{5.2}$$

These coordinates Z_i, Z'_i, Z''_i are the octahedron vertex parameters, logarithmic lifts of the cross-ratios of section 3.2. We also associate with \diamond_i the canonical Lagrangian submanifold

$$\mathcal{L}_{\diamond_i} = \{z''_i + z_i^{-1} - 1 = 0\} \subset \mathcal{P}_{\partial \diamond_i}, \tag{5.3}$$

with $z_i = \exp Z_i$, etc.; and we recall that both the phase space and Lagrangian are invariant under cyclic permutations $Z_i \rightarrow Z'_i \rightarrow Z''_i \rightarrow Z_i$.

For the entire manifold M , we build a product phase space

$$\mathcal{P}_\times := \prod_{i=1}^{\frac{1}{6}NK(K^2-1)} \mathcal{P}_{\partial \diamond_i}, \tag{5.4}$$

with independent Poisson brackets for (Z_i, Z'_i, Z''_i) and (Z_j, Z'_j, Z''_j) when $i \neq j$. This space contains the canonical product Lagrangian

$$\mathcal{L}_\times := \prod_{i=1}^{\frac{1}{6}NK(K^2-1)} \mathcal{L}_{\diamond_i} \subset \mathcal{P}_\times. \tag{5.5}$$

We conjectured in section 4.1 that the phase space $\mathcal{P}_K(\partial M)$ and Lagrangian $\mathcal{L}_K(M)$ associated with the glued manifold M itself are symplectic reductions of $\mathcal{P}_\times, \mathcal{L}_\times$. Let us describe how this works logarithmically.

5.2.1 Paths on slices

It is extremely useful to encode various affine linear functions on \mathcal{P}_\times in terms of paths on the slices of the octahedral decomposition of a manifold M .

Recall that there are $K - 1$ global slices associated with every small boundary component of M , which extend from one tetrahedron to the next. We always view the slices from *above*, that is, from the point of view of the small boundary. On any one of these slices, we consider oriented paths made of segments that begin at the midpoint of a side of a small black or white triangle, curve to the left or right inside the triangle, and end at the midpoint of another side. Two examples are shown in figure 37. The entire path may either be closed or open; if open we require it to begin and end on faces of big tetrahedra.

To any such oriented path γ , we can assign an affine linear function on the product phase space \mathcal{P}_\times by




Figure 37. Open (γ_1) and closed (γ_2) paths on a slice in M .

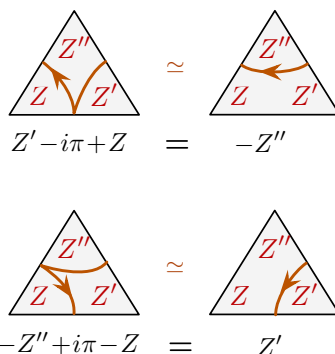


Figure 38. Path homotopy.

1. adding octahedron vertex parameters (Z_i, Z'_i, Z''_i) for the dihedral angles subtended counterclockwise by segments of γ ;
2. subtracting vertex parameters for dihedral angles subtended clockwise;
3. adding (resp., subtracting) $i\pi$ for a 180-degree “bounce” in a counterclockwise (resp., clockwise) direction.

We will call these functions path-coordinates. For example, the open path γ_1 in figure 37 has path-coordinate $-Z_1 + Z_2 + Z_3 + Z_4 + Z_5'' - i\pi + Z_5'$, while the closed path γ_2 has path-coordinate $Z_9 + Z_8 + Z_7 + Z_6 + Z_5 + Z_3' + Z_2''$.

The constraint that $Z_i + Z'_i + Z''_i = i\pi$ (for all octahedra \diamond_i) in the product phase space implies that path-coordinates are invariant under homotopies inside the small black and white triangles, as shown in figure 38.

5.2.2 Gluing functions

Now we proceed to define the new coordinates on \mathcal{P}_\times .

For every internal lattice point v_k in the octahedral decomposition of M , we define an affine linear function C_k on the product phase space to be the sum of octahedron parameters (Z_i, Z'_i , or Z''_i) assigned to octahedron vertices identified with v_k , minus an overall factor of $2\pi i$. These *gluing functions* C_k include:

- $\binom{K-1}{3} = \frac{1}{6}(K-1)(K-2)(K-3)$ points inside every ideal tetrahedron (for example, there is one internal point in the $K = 4$ tetrahedron of figure 33);




Figure 39. Gluing functions from loops on a slice of M .

- $\binom{K-1}{2} = \frac{1}{2}(K-1)(K-2)$ points on every pair of big faces of ideal tetrahedra that are glued together (i.e. on every internal tetrahedron face);
- $\binom{K-1}{1} = K-1$ points on every internal tetrahedron edge.

For example, for the two $K=3$ tetrahedra glued together in figure 36, there are no points inside tetrahedra, and there are no internal tetrahedron edges; but there is one lattice point in the center of the internal face.

In general, suppose that the big boundary of M (tilled by unglued faces of tetrahedra) has genus g and h holes. Let us define the quantity $d_K(g, h) = 2(K^2 - 1)(g - 1) + K(K - 1)h$ as in (2.7), which is the expected complex dimension of the phase space associated with that boundary. Then a straightforward Euler-character argument shows that the total number of internal points in the octahedral decomposition of M is exactly

$$\# C_k\text{'s} = \frac{1}{6}NK(K^2 - 1) - \frac{1}{2}d_K(g, h). \tag{5.6}$$

(If there are multiple disjoint big boundaries, each contributes $-\frac{1}{2}d_K(g_i, h_i)$ to (5.6).) In addition, for every small torus boundary of M there are $K - 1$ linear relations among the C_k 's, described momentarily. Thus, if there are t torus boundaries one *expects*

$$\# \text{ independent } C_k\text{'s} = \frac{1}{6}NK(K^2 - 1) - \frac{1}{2}d_K(g, h) - (K - 1)t. \tag{5.7}$$

Every function C_k is the path-coordinate (minus $2\pi i$) for a small closed loop surrounding a lattice point on a slice of M . The three types of vertices described above correspond to the three types of closed loops shown in figure 39: γ_1 surrounds a point inside a tetrahedron, γ_2 surrounds a point on the face between tetrahedra, and γ_3 surrounds a point on a shared tetrahedron edge. Conversely, every path-coordinate (minus $2\pi i$) for such a small closed loop on a slice defines a C_k , though the same C_k may arise from multiple loops on different slices.

When homotoping a path through a lattice point, as in figure 40, its path-coordinate simply picks up a contribution $\pm C$ for the gluing function assigned to that point. A nice application of this fact is to illustrate the $K - 1$ linear relations among the C_k coming from every small torus boundary of M . Consider any of the $K - 1$ global slices of M associated




Figure 40. Homotopy through a vertex p .




Figure 41. Using path homotopy to obtain relations among C_k 's.

with a small torus boundary — each having the topology of a torus. Choose two small adjacent triangles on this slice, as in figure 41, and draw a closed path γ_1 as shown, so that everywhere outside the two triangles segments of γ_1 always come in pairs with opposite orientations. Due to cancellations from these segments, the path-coordinate for γ_1 is just $2\pi i$. We can then homotope γ_1 to a second path γ_2 that just surrounds a single vertex p , as shown on the left of figure 41. The difference in the path-coordinates for γ_1 and γ_2 is the sum of C_k 's corresponding to every lattice point on this slice except for p . The path-coordinate for γ_2 is $C_p + 2\pi i$ (where C_p corresponds to p). Therefore, we find that *the sum of C_k 's for all the vertices lying on this slice (including C_p) must be zero*. Repeating the argument for the $K - 1$ slices parallel to any small torus boundary produces $K - 1$ linear constraints, which modify the counting in (5.7).

5.2.3 Big boundary coordinates

Recall that the ideal triangulation of M induces a 2d ideal triangulation of the big boundary. Moreover, every big triangle (the face of a tetrahedron) on the big boundary acquires a 2d K -triangulation.

For every lattice point p_j in the K -triangulation of the big boundary we define the function \mathcal{X}_j to be the sum of octahedron parameters (Z_i, Z'_i, Z''_i) at octahedron vertices incident to p_j . This includes $\frac{1}{2}(K - 1)(K - 2)$ functions for every big tetrahedron face on the boundary, and $K - 1$ functions for every edge on the boundary. These are logarithmic versions of the boundary triple-ratios and cross-ratios of section 3.1.

The functions \mathcal{X}_j can be interpreted as path-coordinates for open paths on the slices associated with small disc boundaries of M . This is illustrated in figure 42, in the case of




Figure 42. The big boundary functions \mathcal{X}_j as path-coordinates on slices.




Figure 43. Longitude and meridian paths on the second slice of a torus boundary.

three vertices on a face and one on the edge of a single tetrahedron. Just as in the case of gluing functions C_k , a single \mathcal{X}_j may be represented by multiple paths on multiple slices.

5.2.4 Small torus coordinates

Finally, we define $2(K - 1)$ functions on \mathcal{P}_\times for every small torus boundary of M . Given a torus boundary, we begin by choosing a basis of A and B homology cycles for it, with intersection number $\langle A, B \rangle = 1$. If M is a knot complement in S^3 , there is a distinguished basis given by the *meridian* and *longitude* of the torus boundary, and we will sometimes use this terminology. If M is not a knot complement, the basis of cycles can be chosen at will.

We consider the $K - 1$ torus slices in the octahedral decomposition of M that are associated with the small torus boundary. Let us number the slices in order of their proximity to the torus boundary — so that inside a given tetrahedron Δ the a -th slice is tiled by a^2 small triangles. On each slice, we draw two non-self-intersecting paths μ_a and λ_a such that their projections to the actual torus boundary are in the homology classes of the A and B cycles, respectively. An example of such paths is shown in figure 43. The path-coordinates U_a and V_a of the paths μ_a and λ_a provide the desired functions on \mathcal{P}_\times .

Notice that the definition of U_a and V_a depends on precisely which paths μ_a and λ_a are chosen to represent a given homology class. However, changing the choice of paths can only modify U_a and V_a by multiples of the gluing functions C_k , since any two paths in a fixed homology class are homotopic (using the moves in figures 38 and 40).

After setting $C_k = 0$, the path coordinates U_a and V_a are completely unambiguous. They are logarithms of the holonomy eigenvalues called m_a and ℓ_a in section 3.4! It is quite easy to show this by using 3d snakes to compute the holonomies, as depicted in figure 29.




Figure 44. Arrows representing the Poisson bracket on big boundaries.

A snake rotating around a torus boundary simultaneously sweeps out $K - 1$ paths (say) μ_a on the $K - 1$ slices parallel to that boundary, and the octahedron parameters that the snake moves pick up are exactly the exponentiated versions of the corresponding path coordinates.

5.3 Poisson brackets and the quotient

For a triangulated manifold M , whose tetrahedra are decomposed into octahedra, we have defined three sets of affine linear functions on \mathcal{P}_\times : the gluing functions C_k associated with internal vertices in the octahedral decomposition; the functions \mathcal{X}_j associated with external vertices (on big boundaries); and the $2(K - 1)$ functions $U_a^{(\nu)}$ and $V_a^{(\nu)}$ associated with cycles on every small torus boundary. Here we've indexed the torus boundaries by $\nu = 1, \dots, t$.

These functions on the product phase space \mathcal{P}_\times satisfy very special Poisson brackets. We can summarize the result as follows. First, the gluing functions all commute with each other and with the boundary functions:

$$\{C_k, C_{k'}\} = \{C_k, \mathcal{X}_j\} = \{C_k, U_a^{(\nu)}\} = \{C_k, V_a^{(\nu)}\} = 0. \quad (5.8a)$$

Second, the Poisson bracket of two functions \mathcal{X}_j and $\mathcal{X}_{j'}$ associated with vertices p_j and $p_{j'}$ on the big boundary is equal to the number of octahedron edges going from p_j to $p_{j'}$, counted with orientation:

$$\{\mathcal{X}_j, \mathcal{X}_{j'}\} = \# \text{ edges from } p_j \text{ to } p_{j'} - \# \text{ edges from } p_{j'} \text{ to } p_j. \quad (5.8b)$$

The orientation is clockwise around every black octahedron face. Thus, on every external tetrahedron face, the r.h.s. of (5.8b) can be represented by arrows drawn as in figure 44 (cf. section 3.1). Note that the \mathcal{X}_j coordinates on two disjoint big boundaries necessarily commute. Third, the big boundary coordinates commute with torus coordinates,

$$\{\mathcal{X}_j, U_a^{(\nu)}\} = \{\mathcal{X}_j, V_a^{(\nu)}\} = 0, \quad (5.8c)$$

and the torus coordinates satisfy

$$\{V_a^{(\nu)}, U_{a'}^{(\nu')}\} = \delta^{\nu\nu'} \kappa_{aa'}, \quad (5.8d)$$

$$\{V_a^{(\nu)}, V_{a'}^{(\nu')}\} = \{U_a^{(\nu)}, U_{a'}^{(\nu')}\} = 0, \quad (5.8e)$$

where

$$\kappa_{aa'} = \begin{cases} 2 & a = a' \\ -1 & |a - a'| = 1 \\ 0 & \text{otherwise,} \end{cases} \quad (5.9)$$

is the Cartan matrix of $SL(K)$. Namely, the logarithmic A and B cycle coordinates on different torus boundaries commute; while on the slices corresponding to a single boundary the Poisson bracket is proportional to the intersection number of the cycles, and vanishes if parallel slices are more than one layer apart. These Poisson brackets generalize a fundamental theorem of Neumann and Zagier [8] for $K = 2$.

All of the brackets above can be understood intuitively as statements about path-coordinates. In particular, the coordinate of a closed contractible path γ commutes with the coordinate of any other open or closed path γ' , so long as γ' (if open) begins and ends at the boundary of M . The coordinates of two general paths that intersect have a bracket (roughly) proportional to their intersection number.

We have already proved that the C_k commute with each other and with all other coordinates in section 3.2. We also proved there that big-boundary coordinates \mathcal{X}_j have the expected Poisson bracket (5.8b). The trickiest commutation relations to show are the ones for the A and B cycle holonomies around small tori; we present a full proof in appendix B.

We may recognize $\mathcal{X}_j, U_a^{(\nu)}, V_a^{(\nu)}$ as coordinates on the universal cover $\mathcal{P}_K(\partial M) = \mathcal{P}_K(\partial M, \mathbf{t})$ of the space $\mathcal{X}_K^{\text{un}}(\partial M, \mathbf{t})$ of framed flat $PGL(K)$ -connections on the boundary. Thus, on one hand, we have a projection

$$\mathcal{P}_\times \rightarrow \mathcal{P}_K(\partial M). \quad (5.10)$$

On the other hand, the Poisson brackets (5.8b)–(5.8d) precisely match the Atiyah-Bott-Goldman bracket on $\mathcal{X}_K^{\text{un}}(\partial M)$, cf. (2.14), (3.7). This means that we may (non-canonically) *embed*

$$\mathcal{P}_K(\partial M) \subset \mathcal{P}_\times. \quad (5.11)$$

as a complex symplectic subspace. Moreover, the brackets $\{C_k, C_{k'}\} = 0$ among gluing functions allow us to define the holomorphic symplectic quotient $\mathcal{P}_\times // (C_k)$, using the C_k 's as a collection of commuting moment maps. Since our space \mathcal{P}_\times is affine-linear, the quotient is rather trivial to describe: each C_k generates a \mathbb{C} translation action; we quotient out by these translations and then set $C_k = 0$. Since the C_k commute with all boundary coordinates $\mathcal{X}_j, U_a^{(\nu)}, V_a^{(\nu)}$, the boundary coordinates descend to well-defined functions on the symplectic quotient,¹⁶ and we obtain

Proposition 5.1. *The boundary phase space $\mathcal{P}_K(\partial M)$ can be identified (non-canonically) as a subspace of the symplectic quotient*

$$\mathcal{P}_K(\partial M) \subset \mathcal{P}_\times // (C_k). \quad (5.12)$$

¹⁶Note in particular that on the quotient $\mathcal{P}_\times // (C_k)$, the functions $U_a^{(\nu)}$ and $V_a^{(\nu)}$ no longer have any dependence on the choice of paths drawn on slices, since the C_k are all set to zero. These functions depend only on the homology classes of the paths.

Equivalently, we may factor $\mathcal{P}_\times \simeq \mathcal{P}_K(\partial M) \times \mathcal{P}_C \times \mathcal{P}_C^\perp$, with symplectic form $\Omega = \sum_i dZ \wedge dZ' = \Omega_{\text{WP}} + \sum_k dC_k \wedge d(\dots) + \Omega^\perp$ on the three factors. The quotient above equals $\mathcal{P}_\times // (C_k) \simeq \mathcal{P}_K(\partial M) \times \mathcal{P}_C^\perp$.

The Symplectic Gluing Conjecture 4.1 amounts to the statement that $\mathcal{P}_K(\partial M) \simeq \mathcal{P}_\times // (C_k)$, or in other words that \mathcal{P}_C^\perp is trivial. If we assume that the number of independent gluing functions C_k is exactly the expected number (5.7), then the triviality of \mathcal{P}_C^\perp follows from an elementary counting of dimensions. Namely, the boundary space $\mathcal{P}_K(\partial M)$ must have complex dimension

$$\dim_{\mathbb{C}} \mathcal{P}_K(\partial M) = 2(K - 1)t + \sum_{i=1}^b d_K(g_i, h_i), \tag{5.13}$$

where t is the number of small torus boundaries, g_i, h_i are the genus and number of holes for each of the b big boundaries, and $d_K(g, h) = 2(K^2 - 1)(g - 1) + K(K - 1)h$ as in (2.7). Then from (5.7) we find

$$\dim_{\mathbb{C}} \mathcal{P}_K(\partial M) = \dim_{\mathbb{C}} \mathcal{P}_\times - 2(\# \text{ independent } C_k \text{'s}), \tag{5.14}$$

leaving no room for \mathcal{P}_C^\perp .

Corollary 4.1 would then tell us that the image of \mathcal{L}_\times under the quotient (5.12) defines the Lagrangian submanifold $\mathcal{L}_K(M) \subset \mathcal{P}_K(\partial M)$. This involves projecting with respect to the flows of the moment maps C_k , then intersecting with $C_k = 0$. (Technically speaking, $\mathcal{L}_K(M)$ is the logarithmic lift of the space we have previously denoted $\mathcal{L}_K(M, \mathfrak{t}_{3d})$.)

Unfortunately, the combinatorics required to prove that the exact number of independent gluing functions agrees with (5.7) are rather complicated for general admissible M . Neumann succeeded in proving this for $K = 2$ in [70]. After the first version of our paper, Garoufalidis-Zickert proposed a proof for general $K \geq 2$ [71].

In the case of hyperbolic 3-manifolds and $\text{PGL}(2, \mathbb{C})$ -connections, the independence of the gluing functions was originally proved by [8] by (in our language) analyzing the images of Lagrangian submanifolds under symplectic reduction, and using Mostow rigidity. The argument could easily be generalized here, if one knew about rigidity for framed flat $\text{PGL}(K, \mathbb{C})$ -connections. However, the statement about independence of the C_k 's is ultimately topological in nature, and should hold independent of any rigidity properties of connections in the bulk.

5.3.1 Unipotent conditions

The functions $U_a^{(\nu)}$ and $V_a^{(\nu)}$ are independent on $\mathcal{P}_K(\partial M)$, but the \mathcal{X}_j are not — we know from section 3.1 (cf. figure 11) that requiring unipotent holonomy around holes on the big boundary leads to constraints among the $x_j = \exp(\mathcal{X}_j)$. It is easy to understand these constraints in terms of paths and slices. We briefly explain how this works.

For every hole in the big boundary of M (filled in by a small disc) we expect $K - 1$ central relations among the \mathcal{X}_j . The hole in the big boundary is filled in by a small disc, and in the octahedral decomposition of M there are $K - 1$ slices parallel to this disc. The




Figure 45. Loops surrounding a hole on the big boundary (left), and a contractible path on the 2nd slice parallel to the hole that demonstrates the constraint on eight X_j functions (right).

path-coordinate for a path on the periphery of such a slice, as on the right of figure 45, is a sum of functions $(\mathcal{X}_j - i\pi)$; but since the path is a contractible clockwise loop its total coordinate (after setting $C_k = 0$) must equal $-2\pi i$. Therefore, we find that the sum of functions $(\mathcal{X}_j - i\pi)$ around each of the $K - 1$ loops surrounding a hole (as on the left of figure 45) must equal $-2\pi i$,

$$\sum_{p_j \text{ around a hole}} (\mathcal{X}_j - i\pi) = -2\pi i. \quad (5.15)$$

Exponentiating these relations, we recover the constraints introduced in section 3.1.

5.3.2 Symplectic data

Finally, let us summarize the combinatorial data needed to define a class- \mathcal{R} theory $T_K[M]$ associated with M (as well as any quantum wavefunctions). This requires making a few more choices. Most importantly, we must choose a *polarization* for $\mathcal{P}(\partial M)$.

For the part of $\mathcal{P}(\partial M)$ coming from the big boundary, a choice of polarization amounts to a splitting of the \mathcal{X}_j into linearly independent “positions” X_j and “momenta” P_j that obey canonical Poisson brackets

$$\{P_j, X_{j'}\} = \delta_{jj'}. \quad (5.16)$$

For the part of $\mathcal{P}(\partial M)$ coming from small tori, the choice of meridian and longitude (or A and B) cycles already amounts to a polarization. In addition, it is useful to choose (rational) linear combinations \tilde{V}_a of the V_a ’s that have canonical Poisson brackets with the meridian coordinates,

$$\{\tilde{V}_a, U_{a'}\} = \delta_{aa'}. \quad (5.17)$$

Given (5.8d), it suffices to take $\tilde{V}_a = \sum_b (\kappa^{-1})_{ab} V_b$.

Altogether, we call these choices a polarization Π for the boundary ∂M . The theory $T_K[M, \Pi]$ depends on Π .

We also make two inconsequential choices: we choose a maximal independent set of gluing functions C_k , and a set of affine linear functions Γ_k on \mathcal{P}_\times that are canonically conjugate to these independent C_k , i.e. $\{\Gamma_k, C_{k'}\} = \delta_{kk'}$.

If we assume the Symplectic Gluing Conjecture 4.1, then the independent C_k , along with the Γ_k and the functions $X_j, P_j, U_a^{(\nu)}, \tilde{V}_a^{(\nu)}$, define a new set of symplectic coordinates on \mathcal{P}_\times . Their relation to the original coordinates (Z_i, Z'_i, Z''_i) can be encoded in the affine symplectic transformation

$$\begin{pmatrix} X_j \\ U_a^{(\nu)} \\ C_k \\ P_j \\ \tilde{V}_a^{(\nu)} \\ \Gamma_k \end{pmatrix} = g \cdot \begin{pmatrix} Z_i \\ Z''_i \end{pmatrix} - i\pi \sigma, \tag{5.18}$$

where g is a symplectic matrix in $Sp(\frac{1}{3}NK(K^2 - 1), \mathbb{Q})$, and σ is a $\frac{1}{3}NK(K^2 - 1)$ -dimensional vector of rational numbers. (Recall that N is the number of tetrahedra in the triangulation of M .)

The rows of g and σ corresponding to the gluing functions C_k are quite special; their properties are discussed further in appendix A. Here it is at least clear that they only contain integer entries. Often (but not always) it is possible to choose polarization wisely so that the entire g and σ contain only integer entries, i.e. $(g, \sigma) \in ISp(\frac{1}{3}NK(K^2 - 1), \mathbb{Z})$. For example, this seems possible for any knot complement in S^3 , for any K .

Given the symplectic data (g, σ) corresponding to the triangulation of M and a choice of polarization, one can attempt to define a gauge theory $T_K[M, \Pi]$ as well as various quantum wavefunctions associated with M , following [1, 9, 54, 55, 59, 60, 65], etc. Each of these objects requires that the triangulation satisfy a few extra properties in order to be well-defined.

5.4 Remarks on quantization

We have claimed that constructing a logarithmic lift of functions on \mathcal{P}_\times and $\mathcal{P}_K(\partial M)$ using path coordinates is necessary for quantization. We may briefly illustrate this idea, in terms of quantizing the Lagrangian $\mathcal{L}_K(M) = \mathcal{L}_K(M, \mathfrak{t}_{3d})$ — i.e. promoting $\mathcal{L}_K(M, \mathfrak{t}_{3d})$ to a left ideal for the quantized algebra of functions on $\mathcal{X}_K(\partial M, \mathfrak{t})$. We follow the formalism of [9].

(We note that, while constructing theories $T_K[M, \Pi]$ and wavefunctions requires a choice of polarization, the quantization of the Lagrangian $\mathcal{L}_K(M)$ as a left ideal is independent of Π . Thus, the quantization of $\mathcal{L}_K(M)$ provides the simplest possible example of quantization and the necessity of logarithmic lifts.)

For every octahedron, quantization promotes the coordinates Z_i, Z'_i, Z''_i to operators with commutation relations

$$[\hat{Z}_i, \hat{Z}'_i] = [\hat{Z}'_i, \hat{Z}''_i] = [\hat{Z}''_i, \hat{Z}_i] = \hbar, \tag{5.19}$$

and a central constraint

$$\hat{Z}_i + \hat{Z}'_i + \hat{Z}''_i = i\pi + \hbar/2. \tag{5.20}$$

We may also exponentiate, setting $\hat{z}_i = \exp \hat{Z}_i$, etc., as well as $q = \exp \hbar$. Then the commutation relations are $\hat{z}_i \hat{z}'_i = q \hat{z}'_i \hat{z}_i$, etc., and the central constraint is $\hat{z}_i \hat{z}'_i \hat{z}''_i = -q$. The

octahedron Lagrangian becomes an operator

$$\hat{\mathcal{L}}_{\diamond_i} : (\hat{z}_i'' + \hat{z}_i^{-1} - 1), \tag{5.21}$$

which generates a left ideal in the q -Weyl algebra of exponentiated operators. Notice that the left ideal is invariant under cyclic permutations $\hat{z}_i \rightarrow \hat{z}_i' \rightarrow \hat{z}_i'' \rightarrow \hat{z}_i$, due to the \hbar correction in the central constraint (5.20).

Now we can promote the rules for path coordinates of section 5.2 to define new linear operators $\hat{C}_k, \hat{\mathcal{X}}_j, \hat{U}_a, \hat{V}_a$ in the algebra generated by the $\hat{Z}_i, \hat{Z}_i', \hat{Z}_i''$ for all octahedra \diamond_i . Along a given path, we simply add octahedron parameters with the appropriate signs, and modify every contribution of $\pm i\pi$ to $\pm(i\pi + \hbar/2)$. Moreover, we set \hat{C}_k to be the sum of octahedron parameters around a closed loop, minus $(2\pi i + \hbar)$. The consistent modification of factors of $i\pi$ to $i\pi + \hbar/2$ leads to a consistent prescription for quantum path coordinates that are invariant under path homotopy.

In order to glue the octahedra together, we perform the quantum-mechanical version of symplectic reduction. We form the product left ideal $\hat{\mathcal{L}}_{\times} = (\hat{z}_i'' + \hat{z}_i^{-1} - 1)_{i=1}^{\#\text{octs}}$, and take the centralizer of this ideal with respect to the quantum gluing functions \hat{C}_k . This leads to an ideal that involves only the exponentiated operators $\hat{c}_k = \exp \hat{C}_k, \hat{x}_j = \exp \hat{\mathcal{X}}_j, \hat{m}_a = \exp \hat{U}_a$, and $\hat{\ell}_a = \exp \hat{V}_a$. In the centralizer, we can unambiguously set $\hat{c}_k = 1$, and obtain a quantization of the Lagrangian $\hat{\mathcal{L}}_K(M)$. By construction, upon setting $q \rightarrow 1$, this Lagrangian should reduce to the classical ideal $\mathcal{L}_K(M)$.

This procedure illustrates that the purpose of logarithmic path coordinates was to properly keep track of powers of q . The proposed construction of $\hat{\mathcal{L}}_K(M)$ here suffers from a some technical difficulties (for example, the operation of taking the centralizer of $\hat{\mathcal{L}}_{\times}$ with respect to the \hat{C}_k is not guaranteed to succeed in general). However, modulo these difficulties, it was shown in [9] that the prescription for assigning q -corrections here is absolutely unique. It is the only prescription that allows cyclic invariance of the (5.21) and allows $\hat{\mathcal{L}}_K(M)$ to be invariant under 2–3 moves.

5.5 2–3 moves and path coordinates

We now return briefly to the 2–3 moves of section 3.5. We demonstrate how our logarithmic path coordinates change under elementary 2–3 moves, and why they can be consistently defined as the triangulation of an admissible manifold M changes. In the process, we explicitly provide the birational map between varieties $\mathcal{L}_K(M, \mathbf{t}_{3d}), \mathcal{L}_K(M, \mathbf{t}'_{3d})$ associated to different triangulations.

(The discussion here is parallels the treatment of 2–3 moves and their quantization in [9, section 5.4]; we simply replace $K = 2$ tetrahedra replaced by octahedra.)

5.5.1 Elementary moves

Recall the two elementary moves on octahedra from section 3.5. We show them again in figure 46. From a combinatorial perspective, they are absolutely identical.

Let us assign logarithmic parameters Z_i, Z_i', Z_i'' to the octahedra on the “2” side of the move, and parameters W_i, W_i', W_i'' to the octahedra on the “3” side. We found previously




Figure 46. The two 2–3 moves on octahedra. The second is equivalent to the first, viewed from behind, with white and black faces exchanged.

that the coordinates associated with the nine points on the “boundary” of the clump of octahedra involved in the move did not change under the move — they were products of octahedron cross-ratios before and after. Logarithmically, this leads us to the nine relations

$$\begin{aligned}
 Z_1 &= W_2'' + W_3', & Z_1' &= W_3'' + W_1', & Z_1'' &= W_1'' + W_2', \\
 Z_2 &= W_2' + W_3'', & Z_2' &= W_1' + W_2'', & Z_2'' &= W_3' + W_1'', \\
 W_1 &= Z_1 + Z_2, & W_2 &= Z_1' + Z_2'', & W_3 &= Z_1'' + Z_2'.
 \end{aligned}
 \tag{5.22}$$

On the “3” side, there is also an internal trivalent vertex, giving a gluing function $C = W_1 + W_2 + W_3 - 2\pi i$. Therefore, the dimension of the boundary phase space associated with this clump of octahedra is 4 on both sides. For example, if we choose a polarization $(X_1, X_2; P_1, P_2) = (Z_1, Z_2; Z_1'', Z_2'')$ for the product phase space $\mathcal{P}_{\partial\Diamond}^{(Z_1)} \times \mathcal{P}_{\partial\Diamond}^{(Z_2)}$ on the “2” side and a consistent polarization $(X_1, X_2, C; P_1, P_2, \Gamma) = (W_2'' + W_3', W_2' + W_3'', W_1 + W_2 + W_3 - 2\pi i; W_1'' + W_2', W_3' + W_1'', W_1'')$ for the product phase space $\mathcal{P}_{\partial\Diamond}^{(W_1)} \times \mathcal{P}_{\partial\Diamond}^{(W_2)} \times \mathcal{P}_{\partial\Diamond}^{(W_3)}$ on the “3” side, then it is clear that

$$\mathcal{P}_{\partial(\text{clump})} \simeq \mathcal{P}_{\partial\Diamond}^{(Z_1)} \times \mathcal{P}_{\partial\Diamond}^{(Z_2)} = \{X_1, X_2, P_1, P_2\} = (\mathcal{P}_{\partial\Diamond}^{(W_1)} \times \mathcal{P}_{\partial\Diamond}^{(W_2)} \times \mathcal{P}_{\partial\Diamond}^{(W_3)}) // (C = 0). \tag{5.23}$$

We can also compute the Lagrangians corresponding to the clump, and check invariance. Let us set $x_i = e^{X_i}$, $p_i = e^{P_i}$ as usual. On the “2” side we trivially find

$$\mathcal{L}_{\text{clump},2} = \{p_1 + x_1^{-1} - 1 = 0, p_2 + x_2^{-1} - 1 = 0\}, \tag{5.24}$$

while the Lagrangian on the “3” side is given, after eliminating e^Γ and setting $e^C = 1$, by

$$\mathcal{L}_{\text{clump},3} = \{(1 - x_1 x_2)(p_1 + x_1^{-1} - 1) = 0, (1 - x_1 x_2)(p_2 + x_2^{-1} - 1) = 0\}. \tag{5.25}$$

The ideals (5.24) and (5.25) are equivalent, provided that $x_1 x_2 = e^{W_1} \neq 1$. This is precisely the condition that octahedra are non-degenerate, i.e. their cross-ratio parameters take values in $\mathbb{C}^* \setminus \{1\}$, which is ensured if classical configurations of framing flats are generic. As we already know, in general, the map of Lagrangians is birational.




Figure 47. The sequence of moves on octahedra that realizes a full 2–3 move for $K = 3$.




Figure 48. Transforming paths during the 2–3 move.

5.5.2 Composing elementary moves

In order to perform a full 2–3 move on ideal tetrahedra that have K -decompositions, we iterate the elementary 2–3 moves. We recall in figure 47 how this worked for $K = 3$.

Notice that path-coordinates as discussed in section 5.2.1 can be translated from one side of the full 2–3 move to the other. Recall that the path-coordinates are affine linear combinations of octahedron parameters. Then, at each step in the sequence of elementary 2–3 moves, the relations (5.22) provide a dictionary for transforming the path-coordinates. These transformations are well-defined, modulo the addition of gluing constraints C_k . Graphically, homotopy classes of paths on slices on one side of the 2–3 move are mapped to homotopy classes of paths on the other side. In figure 48 we give an example of how this works on two sets of slices of the $K = 3$ bipyramid, outlined in red and in blue in figure 47. Note, however, that slices of the bipyramid are *not* uniquely defined during the intermediate steps of the full 2–3 move; they are only canonically defined at the beginning and at the end.

6 Tetrahedron and polyhedron theories

In this section we present the first and simplest examples of $\mathcal{N} = 2$ gauge theories $T_K[M]$ associated with 3-manifolds with $K > 2$. We will also compute some of their observables and the corresponding geometric quantities for the 3-manifolds, such as the supersymmetric parameter space $\mathcal{L}_{\text{SUSY}} \simeq \mathcal{L}_K(M)$.

We start with a survey of tetrahedron theories $T_K[\Delta, \Pi]$ for $K = 2, 3, 4, 5$. These beautifully illustrate the types of complexities that arise at higher K and how they can be handled. With a judicious choice of boundary polarization Π , the theories can be described as

K	$T_K[\Delta]$	flavor sym.
2	1 free chiral	$U(1)$
3	4 free chirals	$U(1)^4$
4	10 chirals + degree-six superpotential	$U(1)^9$
5	20 chirals, $U(1)^2$ gauge group, $W = \sum(\text{four monopole ops})$	$U(1)^{16}$

(6.1)

In general, the class- \mathcal{R} methods of this paper construct $T_K[\Delta]$ as a theory with manifest flavor symmetry¹⁷ $U(1)^{(K-1)^2}$, where $2(K-1)^2$ is the dimension of the phase space $\mathcal{P}_K[\partial\Delta]$. The theory has $\frac{1}{6}K(K^2 - 1)$ chirals (the number of octahedra in the K -decomposition) and $\frac{1}{6}(K-1)(K-2)(K-3)$ operators added to its superpotential (the number of internal lattice points). The dynamical gauge group depends on the choice of boundary polarization, though for $K \geq 5$ there will always be a nontrivial gauge group; likewise, for $K \geq 5$ the superpotential will always contain some monopole operators.

We also consider the bipyramid theory for $K = 3$. The triangular bipyramid can be triangulated in two different ways, leading to two different theories $T_3[\text{bip}, \Pi]$ that are mirror symmetric. With a suitable choice of polarization, one theory (coming from two tetrahedra) can be described as a collection of eight chirals with a single degree-six superpotential; while the other (coming from three tetrahedra) is a $U(1)^6$ gauge theory with twelve chirals and five operators in its superpotential, including three monopole operators. The mirror symmetry is quite non-trivial! Nevertheless, we know mirror symmetry must hold, and it can systematically be broken down into a sequence of fundamental XYZ \leftrightarrow SQED mirror symmetries just by looking at the geometric decomposition of a 2-3 move into octahedra, as in figure 47. This is a nice example of how a complicated physical duality is encoded by relatively simple geometry.

The polyhedron theories $T_K[M, \Pi]$ in this section all depend on a choice of polarization, and in general there is no canonical choice.¹⁸ This makes them somewhat ill-suited for a study of large- K behavior, despite the nice hint of K^3 scaling for degrees of freedom that comes from counting octahedra. On the other hand, we should keep in mind that in

¹⁷With a proper choice of background Chern-Simons terms (part of the polarization data), these flavor symmetries may occasionally be enhanced to non-abelian groups of the same rank. We will only write down the maximal abelian part of the flavor symmetry group.

¹⁸Put differently, in the spirit of [30], there is a great deal of freedom in choosing asymptotic boundary conditions for K M5-branes wrapped on a polyhedron.

many circumstances the choice of polarization is not important. For example, the choice dissolves once the theories are re-interpreted as boundary conditions for $U(1)$ gauge fields in four dimensions (i.e. by weakly gauging $U(1)$ flavor symmetries in a 4d bulk). This was explained in [1]. In three dimensions, it is natural to think of theories $T_K[M, \Pi]$ as coming in families or equivalence classes whose elements are related by $Sp(2N, \mathbb{Z})$ transformations. While some observables of a theory (such as partition functions or indices) depend on a choice of representative in the class, and transform under $Sp(2N, \mathbb{Z})$, others (such as $\mathcal{L}_{\text{SUSY}} \simeq \mathcal{L}_K(M)$) do not.

In this section, we will usually work with the logarithmic lifts $\mathcal{P}_K(\partial M)$ of phase spaces $\mathcal{X}_K(\partial M)$, as in section 5. We alternatively use uppercase logarithmic coordinates (e.g. Z_i) and the lowercase exponentiated ones (e.g. $z_i = \exp Z_i$), as appropriate.

A summary of the systematic rules used to obtain a 3d gauge theory in class \mathcal{R} from the symplectic gluing data (g, σ) of an octahedral decomposition (assuming the Symplectic Gluing Conjecture 4.1) is given in appendix A. We use these rules implicitly to derive gauge theories below, but focus the discussion on more intuitive aspects of the geometry-physics correspondence.

6.1 Review: the octahedron theory

To begin, let us review the theory T_\diamond associated with an octahedron — or, equivalently a tetrahedron for $K = 2$. It was constructed in [1].

Recall that the vertices of an octahedron are labeled by a triplet of parameters (Z, Z', Z'') , occurring in the same clockwise cyclic order around every black face. Classically they satisfy $Z + Z' + Z'' = i\pi$ and the Poisson brackets $\{Z, Z'\} = \{Z', Z''\} = \{Z'', Z\} = 1$. Once we fix a labeling, the octahedron has a canonical initial polarization with Z a position and Z'' a momentum. Typically we will also choose the final polarization $\Pi_Z = (X; P) = (Z; Z'')$ to agree with the initial one. Then $T_\diamond := T_2[\diamond, \Pi_Z]$ is the theory of a free $\mathcal{N} = 2$ chiral multiplet.

Specifically, we define

$$T_\diamond : \left\{ \begin{array}{l} \text{Free chiral } \phi, \\ U(1)_F \text{ flavor symmetry with complex mass parameter } Z; \\ \text{charges: } \begin{array}{|c|c|} \hline & \phi \\ \hline F & 1 \\ \hline R & 0 \\ \hline \end{array} \quad \text{CS levels: } \begin{array}{|c|c|c|} \hline & F & R \\ \hline F & -\frac{1}{2} & \frac{1}{2} \\ \hline R & \frac{1}{2} & * \\ \hline \end{array} \end{array} \right. \quad \begin{array}{c} \text{Diagram of an octahedron with vertices labeled } Z, Z', Z'' \text{ and } \phi. \text{ A blue dot is labeled } X \text{ and a blue arrow is labeled } P. \end{array} \tag{6.2}$$

The theory has a $U(1)_F$ flavor symmetry rotating the phase of the chiral, and we add $-1/2$ units of Chern-Simons coupling¹⁹ for a background vector multiplet associated with this flavor symmetry. This is a Chern-Simons contact term, as in [86, 87]. The theory also has a $U(1)_R$ R-symmetry. We specify an R-charge assignment, as well as mixed background

¹⁹As usual, chiral matter in three dimensions produces one-loop half-integer contributions to effective Chern-Simons levels for gauge or flavor symmetries, cf. [4, 84, 85]. Thus bare Chern-Simons must sometimes take half-integer values, as in (6.2), in order for the full theory to be gauge invariant.

Chern-Simons couplings between R-symmetry and flavor-symmetry fields. These are irrelevant in flat space, but become important when one compactifies the theory on S^1 , S^3 , $S^2 \times S^1$, etc. Similarly, in flat 3d space the mass parameter associated with the flavor symmetry (a scalar in the vector multiplet) is real, but gets complexified in any of these compactifications. This mass parameter is identified with Z . On the 3-manifold side, it is the logarithmic lift of phase-space coordinates that lets us consistently keep track of R-charge assignments and complexified masses Z in the compactifications. (These comments apply not just to the octahedron theory, but to all 3-manifold theories obtained from triangulations.)

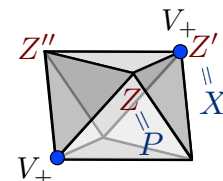
Notice that T_\diamond has a chiral operator ϕ . Geometrically, it can be associated with both of the vertices with final position coordinate $X = Z$ in (6.2). We are also at liberty to choose *different* final polarizations for the octahedron theory, related to Π_Z by an $\text{Sp}(2, \mathbb{Z})$ transformation. In any given polarization $\Pi = (X; P)$, the octahedron vertices commuting with the position coordinate X (if any exist) will be assigned chiral operators.

For example, if we choose polarization $\Pi_{Z'} = (Z'; Z)$ with Z' as a position, the theory becomes $T'_\diamond := T_2[\diamond, \Pi_{Z'}]$

$$T'_\diamond : \left\{ \begin{array}{l} \text{U}(1)_G \text{ gauge theory w/ a chiral } \phi', \\ \text{U}(1)_J \text{ topological flavor symmetry w/ complex mass } Z'; \end{array} \right.$$

	ϕ'	V_+
charges:	G	J
	1	0
	0	1
	0	0

	G	J	R
CS levels:	$\frac{1}{2}$	1	$-\frac{1}{2}$
	1	0	0
	$-\frac{1}{2}$	0	*



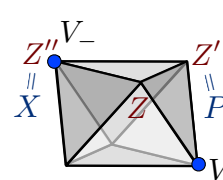
(6.3)

It has a gauge-invariant monopole operator V_+ “sitting” on the Z' vertices. Similarly, in polarization $\Pi_{Z''} = (Z''; Z')$ we have $T''_\diamond := T_2[\diamond, \Pi_{Z''}]$

$$T''_\diamond : \left\{ \begin{array}{l} \text{U}(1)_G \text{ gauge theory w/ a chiral } \phi'', \\ \text{U}(1)_J \text{ topological flavor symmetry w/ complex mass } Z''; \end{array} \right.$$

	ϕ''	V_-
charges:	G	J
	1	0
	0	1
	0	0

	G	J	R
CS levels:	$-\frac{1}{2}$	-1	$\frac{1}{2}$
	-1	-1	1
	$\frac{1}{2}$	1	*



(6.4)

with an anti-monopole operator V_- on the Z'' vertices.

The three theories above are actually all mirror-symmetric:

$$T_\diamond \simeq T'_\diamond \simeq T''_\diamond. \tag{6.5}$$

The flavor symmetry $\text{U}(1)_F$ of T_\diamond matches the topological $\text{U}(1)_J$ symmetries of T'_\diamond and T''_\diamond . Moreover, the respective monopole (V_+) and anti-monopole (V_-) operators of T'_\diamond and




Figure 49. Longitudinal polarization on mirror-symmetric clumps of octahedra.

T''_{\diamond} match the chiral operator ϕ of T_{\diamond} . This special mirror symmetry shows that in a given 3-manifold M , where final polarizations have been fixed, it makes no difference how the octahedra are labeled by (Z, Z', Z'') as long as the cyclic order around black faces is preserved — there are three equivalent choices for every octahedron.

In *any* polarization, the octahedron theory will have a supersymmetric parameter space²⁰ isomorphic to $\mathcal{L}_{\text{SUSY}} \simeq \mathcal{L}_{\diamond} = \{p + x^{-1} - 1 = 0\}$. Changing the polarization just acts as an $\text{Sp}(2, \mathbb{Z})$ transformation on $(X; P)$, with $x = \exp X$ and $p = \exp P$ as usual.

6.2 Warmup: elementary 2–3 move theories

As a warmup exercise, we derive the theories associated with the clusters of either 2 or 3 octahedra that enter elementary 2–3 moves. Alternatively, these are theories $T_2[\text{bip}]$ for a $K = 2$ bipyramid, as investigated in [1]. Versions of these theories appear as building blocks in almost all more interesting 3-manifold theories. They also illustrate the fundamental mirror symmetry associated with the 2–3 move on octahedra, which (putatively) is used to build up all the mirror symmetries in class \mathcal{R} .

In all cases, the theories have $U(1)^2$ flavor symmetry. Put differently, the dimension of the phase spaces $\mathcal{P}_{\partial(\text{cluster})}$ corresponding to these clumps of octahedra is always four. We label octahedra by parameters (z_1, z_2) on the 2-side and (w_1, w_2, w_3) on the 3-side, as shown in figures 49–50 (just as in section 5.5). Then we have

$$\mathcal{P}_{\partial(\text{cluster})} \simeq \mathcal{P}_{\partial_{\diamond} z_1} \times \mathcal{P}_{\partial_{\diamond} z_2} \simeq (\mathcal{P}_{\partial_{\diamond} w_1} \times \mathcal{P}_{\partial_{\diamond} w_2} \times \mathcal{P}_{\partial_{\diamond} w_3}) // (C = 0), \quad (6.6)$$

where $C = W_1 + W_2 + W_3 - 2\pi i$ is the gluing function for the internal point on the 3-side. There is no canonical final polarization on $\mathcal{P}_{\partial(\text{cluster})}$, but there are two somewhat natural choices: a “longitudinal” polarization (figure 49) and an “equatorial” polarization (figure 50). The two choices give slightly different descriptions of the mirror pairs of theories.

In the longitudinal polarization, the cluster of two octahedra is very simple. We take $\Pi_{\text{long}} = (X_+, X_-; P_+, P_-) := (Z_1, Z_2; -Z'_1, Z''_2)$. Since the final positions X_{\pm} match the initial positions $Z_{1,2}$, the theory is just a product of two octahedron theories, *a.k.a.* an

²⁰The Lagrangian $\mathcal{L}_{\text{SUSY}}$ is the supersymmetric parameter space of a theory on $\mathbb{C} \times S^1$. See [1, 54, 88] and appendix A.6 for a brief review.

$\mathcal{N} = 4$ hypermultiplet:

$$T_{\text{hyper}} : \left\{ \begin{array}{l} \text{Two chirals } q, \tilde{q}, \text{ with } U(1)_+ \times U(1)_- \text{ flavor symmetry} \\ \text{charges:} \end{array} \right. \begin{array}{|c|c|c|} \hline & q & \tilde{q} \\ \hline X_+ & 1 & 0 \\ \hline X_- & 0 & 1 \\ \hline R & 0 & 0 \\ \hline \end{array} \quad \text{CS levels:} \quad \begin{array}{|c|c|c|c|} \hline & X_+ & X_- & R \\ \hline X_+ & \frac{1}{2} & 0 & -\frac{1}{2} \\ \hline X_- & 0 & -\frac{1}{2} & \frac{1}{2} \\ \hline R & -\frac{1}{2} & \frac{1}{2} & * \\ \hline \end{array} \quad (6.7)$$

In the charge matrix, we have used the positions X_{\pm} to label the $U(1)_{\pm}$ symmetries. The background Chern-Simons levels follow from the choice of momenta P_{\pm} . Notice that the cluster of octahedra has three external coordinates Z_1, Z_2 , and $Z_1 + Z_2$ that commute with the positions X_{\pm} ; correspondingly, T_{hyper} has three chiral operators q, \tilde{q} , and $q\tilde{q}$ transforming under the flavor symmetries at these coordinates.

If instead we build the cluster with three octahedra, we now have a non-trivial $\text{Sp}(6, \mathbb{Z})$ transformation acting on the product phase space:

$$\begin{pmatrix} X_+ \\ X_- \\ C \\ P_+ \\ P_- \\ \Gamma \end{pmatrix} = \begin{pmatrix} 0 & 0 & -1 & 0 & 1 & -1 \\ 0 & -1 & 0 & 0 & -1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & -1 & 1 & 0 & -1 \\ 0 & 0 & -1 & 1 & 1 & -1 \end{pmatrix} \begin{pmatrix} W_1 \\ W_2 \\ W_3 \\ W_1'' \\ W_2'' \\ W_3'' \end{pmatrix} - i\pi \begin{pmatrix} -1 \\ -1 \\ 2 \\ 1 \\ -1 \\ 0 \end{pmatrix} \quad (6.8)$$

The fact that there are initial momenta (W_2'', W_3'') mixed into the final positions X_{\pm} means that we must have a non-trivial gauging. Indeed, we find that the gauge theory is now described as $\mathcal{N} = 4$ SQED. In $\mathcal{N} = 2$ language:

$$T_{\text{SQED}}^{\mathcal{N}=4} : \left\{ \begin{array}{l} \text{Dynamical } U(1) \text{ theory with three chirals and a superpotential } W = \tilde{\phi}\phi\phi; \\ \text{flavor symmetry } U(1)_+ \times U(1)_- \text{ corresponding to (axial} \pm \text{topological)} \end{array} \right. \begin{array}{|c|c|c|c|c|c|} \hline & \phi & \tilde{\phi} & \varphi & V_+ & V_- \\ \hline G & 1 & -1 & 0 & 0 & 0 \\ \hline X_+ & -1 & 0 & 1 & 1 & 0 \\ \hline X_- & -1 & 0 & 1 & 0 & 1 \\ \hline R & 2 & 0 & 0 & 0 & 0 \\ \hline \end{array} \quad \text{CS levels:} \quad \begin{array}{|c|c|c|c|c|} \hline & G & X_+ & X_- & R \\ \hline G & 0 & \frac{1}{2} & -\frac{1}{2} & 0 \\ \hline X_+ & \frac{1}{2} & 0 & 0 & 0 \\ \hline X_- & -\frac{1}{2} & 0 & 0 & 0 \\ \hline R & 0 & 0 & 0 & * \\ \hline \end{array} \quad (6.9)$$

Each chiral $\phi, \tilde{\phi}, \varphi$ corresponds to one of the original octahedra, and their product forms an operator \mathcal{O}_C that must be added to the superpotential to enforce the internal gluing constraint. The remaining flavor symmetry is an axial $U(1)_A$ and a topological $U(1)_J$, which are re-grouped into $U(1)_{\pm}$. It can be checked (following, e.g., rules of [4]) that the theory has both a gauge-invariant monopole operator V_+ and an anti-monopole V_- , with charges as shown. The operators V_{\pm} and the neutral chiral φ now sit at the three external




Figure 50. Equatorial polarization on mirror-symmetric clumps of octahedra.

vertices of the clump that commute with X_{\pm} . The duality

$$T_{\text{hyper}} \simeq T_{\text{SQED}}^{\mathcal{N}=4} \quad (6.10)$$

is one of the classic examples of 3d mirror symmetry [89, 90].

In the equatorial polarization, the final positions X_{\pm} now involve initial momenta on the 2-side but not on the 3-side. Thus, the 2-side gives a dynamical $U(1)$ gauge theory while the 3-side has three ungauged chirals with a superpotential.

Explicitly, on the 2-side, we set $(X_+, X_-; P_+, P_-) = (Z'_1 + Z''_2, Z'_1 + Z'_2; Z'_2, Z'_1)$, or

$$\begin{pmatrix} X_+ \\ X_- \\ P_+ \\ P_- \end{pmatrix} = \begin{pmatrix} -1 & 0 & -1 & 1 \\ 0 & -1 & 1 & -1 \\ 0 & -1 & 0 & -1 \\ -1 & 0 & -1 & 0 \end{pmatrix} \begin{pmatrix} Z_1 \\ Z_2 \\ Z''_1 \\ Z''_2 \end{pmatrix} - i\pi \begin{pmatrix} -1 \\ -1 \\ -1 \\ -1 \end{pmatrix}. \quad (6.11)$$

This results in $\mathcal{N} = 2$ SQED,

$$T_{\text{SQED}}^{\mathcal{N}=2} : \left\{ \begin{array}{l} \text{Dynamical } U(1) \text{ theory with two chirals } q, \tilde{q}; \\ \text{flavor symmetry } U(1)_+ \times U(1)_- \text{ corresponding to (axial } \pm \text{ topological)} \\ \\ \text{charges:} \end{array} \right. \begin{array}{c} \begin{array}{c|c|c|c|c} & q & \tilde{q} & v_+ & v_- \\ \hline G & 1 & -1 & 0 & 0 \\ \hline X_+ & -1 & 0 & 1 & 0 \\ \hline X_- & -1 & 0 & 0 & 1 \\ \hline R & 2 & 0 & 0 & 0 \end{array} \\ \\ \text{CS levels:} \end{array} \begin{array}{c} \begin{array}{c|c|c|c|c} & G & X_+ & X_- & R \\ \hline G & 0 & \frac{1}{2} & -\frac{1}{2} & 0 \\ \hline X_+ & \frac{1}{2} & -\frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ \hline X_- & -\frac{1}{2} & \frac{1}{2} & \frac{1}{2} & -\frac{1}{2} \\ \hline R & 0 & \frac{1}{2} & -\frac{1}{2} & * \end{array} \end{array} \quad (6.12)$$

This is nothing but $\mathcal{N} = 4$ SQED without the neutral chiral and the superpotential; indeed, geometrically, if we remove the third octahedron from the right side of figure 49 and rotate the picture by ninety degrees, we recover the left side of figure 50. The theory $T_{\text{SQED}}^{\mathcal{N}=2}$ has axial and topological $U(1)$ symmetries. Its gauge-invariant operators are a monopole/anti-monopole v_{\pm} and a meson $q\tilde{q}$, which sit at three vertices on the boundary of the cluster.

Finally, if we use the equatorial polarization and build the cluster from three octahedra we obtain the “XYZ” model:

$$T_{XYZ} : \left\{ \begin{array}{l} \text{Three chirals } \phi_1, \phi_2, \phi_3 \text{ with a superpotential } W = \phi_1\phi_2\phi_3, \\ \text{unbroken flavor symmetry } U(1)_+ \times U(1)_- \end{array} \right. \quad (6.13)$$

	ϕ_1	ϕ_2	ϕ_3
X_+	-1	1	0
X_-	-1	0	1
R	2	0	0

CS levels:

	X_+	X_-	R
X_+	0	$\frac{1}{2}$	0
X_-	$\frac{1}{2}$	0	0
R	0	0	*

The chiral operators ϕ_1, ϕ_2, ϕ_3 sit at external vertices of the cluster.

The mirror symmetry

$$T_{\text{SQED}}^{\mathcal{N}=2} \simeq T_{XYZ} \quad (6.14)$$

is a classic example of $\mathcal{N} = 2$ mirror symmetry [4, 91]. It is *equivalent* to the $\mathcal{N} = 4$ mirror symmetry (6.10). We can simply act with an $\text{Sp}(4, \mathbb{Z})$ transformation on both sides of (6.10) to change the boundary polarization from longitudinal to equatorial, recovering (6.14). In physical language: if we simultaneously gauge a vector $U(1)$ in the free hyper theory and a corresponding topological $U(1)_J$ in $\mathcal{N} = 4$ SQED, we recover $\mathcal{N} = 2$ SQED and the XYZ model, respectively.²¹

The twisted superpotentials $\widetilde{W}_{(2)}(x, y)$ and $\widetilde{W}_{(3)}(x, y)$ discussed back in section 1.1 of the Introduction are associated to the theories $T_{\text{SQED}}^{\mathcal{N}=2}$ and T_{XYZ} respectively. The variables x and y correspond to real masses (exponentiated and complexified) for the flavor symmetries $U(1)_+$ and $U(1)_-$, respectively.

6.3 Tetrahedra at higher K

We now proceed to describe the ideal-tetrahedron theories $T_K[\Delta]$ for $K > 2$. They were summarized in (6.1), and we want to demonstrate how they are constructed.

6.3.1 $K = 3$

The four octahedra in a $K = 3$ tetrahedron can be assigned triples of parameters (z, w, s, t) as in figure 51. Since there are no internal lattice points, we expect that the phase space $\mathcal{P}_3(\partial\Delta)$ and Lagrangian $\mathcal{L}_3(\Delta)$ are just products of four copies of the octahedron phase space and Lagrangian. Similarly, in suitable polarizations, the theory $T_K[\Delta] \simeq T_{\diamond z} \otimes T_{\diamond w} \otimes T_{\diamond s} \otimes T_{\diamond t}$ should just be a product of four octahedron theories.

Explicitly, the phase space $\mathcal{P}_3(\partial\Delta)$ — an open patch of the space of framed flat $\text{PGL}(3, \mathbb{C})$ -connections on the boundary with unipotent holonomy at the punctures — is parametrized by the twelve edge coordinates and four face coordinates on the boundary of the tetrahedron, subject to eight constraints from the vertices. The vertex constraints,

²¹To see that we recover the XYZ model this way, we must remember that gauging a $U(1)$ and subsequently gauging its $U(1)_J$ is equivalent in the IR to gauging nothing at all — it is the statement that $S^2 = -I$ in the symplectic group [92].




Figure 51. Octahedra and final polarization for $K = 3$.

imposing unipotent holonomy, say that sums of logarithmic coordinates around loops encircling the vertices equal $-2\pi i$, cf. (5.15) and figure 11. A convenient choice of linearly independent boundary coordinates is given by the polarization

$$\Pi = (X_1, X_2, X_3, X_4; P_1, P_2, P_3, P_4) = (Z, W, S, T; Z'', W'', S'', T''), \quad (6.15)$$

a simple product of octahedron polarizations. Then the Lagrangian becomes the product

$$\mathcal{L}_3(\Delta) = \{p_i + x_i^{-1} - 1 = 0\}_{i=1}^4. \quad (6.16)$$

It parametrizes an open patch of the configuration space of four flags in \mathbb{C}^3 — the invariant flags on the tetrahedron’s vertices.

Using the polarization (6.15) we find that the theory is exactly

$$T_3[\Delta, \Pi] = T_{\diamond z} \otimes T_{\diamond w} \otimes T_{\diamond s} \otimes T_{\diamond t}, \quad (6.17)$$

i.e. a theory of four free chirals. It has $U(1)^4$ flavor symmetry. The chiral operators $\phi_z, \phi_w, \phi_s, \phi_t$ are associated with four external points on the $K = 3$ tetrahedron, as indicated in figure 51. Only these external points have coordinates that commute with the positions in Π .

6.3.2 $K = 4$

For $K = 4$, the tetrahedron is decomposed into ten octahedra, as in figure 52. The novelty is that there is an internal lattice point in the middle of the tetrahedron. It leads to a gluing constraint and to a superpotential coupling.

A convenient choice of octahedron labelings is shown in the figure. We arrange that all six octahedron vertices surrounding the internal point are labeled by un-primed coordinates, so that the classical gluing constraint takes the form

$$C = Z_5 + Z_6 + Z_7 + Z_8 + Z_9 + Z_{10} - 2\pi i \rightarrow 0. \quad (6.18)$$

Then in the gauge theory the corresponding operator will easily be built as a product of six fundamental chirals. Moreover, we notice that all ten of the initial positions Z_i appear as coordinates on the edges of the boundary. Therefore, we can choose an independent




Figure 52. Octahedra, final polarization, and operators for $K = 4$. The octahedron in the middle of the back edge is drawn separately; note that its parameter is not part of the polarization. On the right we show the three slices around the top vertex.

set of nine of them (independent modulo the gluing constraint) as positions in the final polarization Π as well; for example, we take

$$X_i = Z_i, \quad i = 1, \dots, 9. \quad (6.19)$$

Then the theory $T_4[\Delta, \Pi]$ involves no gauging at all. It is basically a product of octahedron theories:

$$T_4[\Delta, \Pi] : \begin{cases} \text{ten chirals } \phi_i \text{ with a superpotential } W = \phi_5 \phi_6 \phi_7 \phi_8 \phi_9 \phi_{10} \\ \text{U}(1)^9 \text{ flavor symmetry left unbroken} \end{cases} \quad (6.20)$$

As dictated by this polarization, the first four $\text{U}(1)_i$ flavor symmetries rotate the first four ϕ_i 's; while the for $5 \leq i \leq 9$ the field ϕ_i has charge $+1$ and ϕ_{10} has charge -1 . Moreover, the first nine fields have R-charge zero and ϕ_{10} has R-charge 2. The background Chern-Simons levels can be determined after a precise choice of momenta P_i in Π is made. (In general, changing the momenta only effects background couplings.)

We might notice that the construction of $T_4[\Delta, \Pi]$ is a simple generalization of the construction of the XYZ model in section 6.2.

The phase space $\mathcal{P}_4(\partial\Delta)$ for the boundary of the $K = 4$ tetrahedron is the symplectic quotient of a product of ten octahedron spaces, using the gluing coordinate C as a moment map. Thus it has complex dimension 18. Recall that C should be thought of as a coordinate on the product $\mathcal{P}_\times = \prod_{i=1}^{10} \mathcal{P}_{\partial\circ i}$, with product polarization $\Pi_\times = (Z_i; Z''_i)$. We can describe $\mathcal{P}_4(\partial\Delta)$ more explicitly by choosing new symplectic coordinates on \mathcal{P}_\times that include the 18 positions and momenta (X_i, P_i) of a final boundary polarization, as well as C and a canonical conjugate Γ so that $\{\Gamma, C\} = 1$. For example, it suffices to choose

$$\Gamma = Z''_{10}, \quad (6.21)$$




Figure 53. The central cluster of the $K = 5$ tetrahedron.

together with $X_i = Z_i$ ($1 \leq i \leq 9$) as well as $P_i = Z_i''$ for $1 \leq i \leq 4$ and $P_i = Z_i'' - \Gamma$ for $5 \leq i \leq 9$. Then

$$\mathcal{P}_4(\partial\Delta) = \left(\prod_{i=1}^{10} \mathcal{P}_{\partial\circ i} \right) / (\Gamma \sim \Gamma + t)|_{C=0} = \{X_i, P_i\}_{i=1}^9 \simeq \mathbb{C}^{18}. \quad (6.22)$$

Note that all the P_i here can be written as linear combinations of boundary coordinates on the tetrahedron. This must always be the case, due to the Poisson brackets of section 5.3.

The nine-dimensional Lagrangian submanifold is obtained by eliminating $\gamma = z_{10}''$ from $\mathcal{L}_\times = \{z_i'' + z_i^{-1} - 1 = 0\}_{i=1}^{10}$, setting $c = z_5 z_6 z_7 z_8 z_9 z_{10} = 1$, and re-writing the equations in terms of x_i and p_i :

$$\mathcal{L}_4(\Delta) : \quad \{p_i - x_i^{-1} - 1 = 0\}_{i=1}^4 \cap \{(1 - x_5 x_6 x_7 x_8 x_9) p_i + x_i^{-1} - 1 = 0\}_{i=5}^9. \quad (6.23)$$

In fact, it is also easy to quantize the Lagrangian using the prescription of [9], outlined in section 5.4. We find the left ideal

$$\mathcal{L}_4(\Delta) = (\hat{p}_i - \hat{x}_i^{-1} - 1)_{i=1}^4 \cap ((1 - \hat{x}_5 \hat{x}_6 \hat{x}_7 \hat{x}_8 \hat{x}_9) \hat{p}_i + \hat{x}_i^{-1} - 1)_{i=5}^9. \quad (6.24)$$

These are the line operators that annihilate partition functions of $T_4[\Delta, \Pi]$.

6.3.3 $K = 5$

The $K = 5$ tetrahedron is a good example of the complexity that arises for general K . It gets decomposed into twenty octahedra (figure 54), with four internal lattice points. No matter how we label the octahedra, the gluing functions at these lattice points can never simultaneously be written as sums of initial positions Z_i . Correspondingly, the four operators added to the superpotential of the theory $T_5[\Delta]$ can never simultaneously be written as products of elementary chiral fields — monopole operators are unavoidable.

It is instructive to build the tetrahedron and the corresponding gauge theory in pieces. We start with a cluster of four octahedra, which have a small upright tetrahedron ($\Delta^{0,2}$ simplex) in their center (figure 53). We label octahedra with parameters (z, w, s, t) as shown, and choose a boundary polarization $\Pi_{\text{cluster}} = (X_i; P_i)_{i=1}^4$ for the cluster so that the four light-blue boundary points in the figure commute with the positions of Π_{cluster} . That means that there will be gauge-invariant operators associated with them.

Explicitly, we set $\{X_1 = Z + W, X_2 = S + T, X_3 = Z + S' + T'', X_4 = S + Z' + W''\}$. The positions X_1, X_2 are combinations of initial octahedron positions, while X_3, X_4 involve initial momenta; thus we expect something like a $U(1)^2$ gauge theory. In matrix form,

$$\begin{pmatrix} X_1 \\ X_2 \\ X_3 \\ X_4 \\ P_1 \\ P_2 \\ P_3 \\ P_4 \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 & 0 & 0 & -1 & 1 \\ -1 & 0 & 1 & 0 & -1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} Z \\ W \\ S \\ T \\ Z'' \\ W'' \\ S'' \\ T'' \end{pmatrix} - i\pi \begin{pmatrix} 0 \\ 0 \\ -1 \\ -1 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (6.25)$$

By comparing the form of this matrix with (6.11), we find that the cluster theory is basically two copies of $\mathcal{N} = 2$ SQED. More accurately, it is two copies of SQED whose $U(1)$ gauge fields are coupled by a mixed Chern-Simons term:

$$T_{\text{cluster}} : \left\{ \begin{array}{l} U(1)_1 \times U(1)_2 \text{ gauge theory coupled to four chirals, no superpotential} \\ U(1)_A^2 \times U(1)_J^2 \text{ flavor symmetry (with } (A_1, A_2, J_1, J_2) \sim (X_1, X_2, X_4, X_3) \text{ above)} \end{array} \right. \begin{array}{c} \text{charges:} \\ \begin{array}{|c|cccc|cccc|} \hline & \phi_z & \phi_w & \phi_s & \phi_t & V_1^+ & V_1^- & V_2^+ & V_2^- \\ \hline G_1 & 1 & -1 & 0 & 0 & 0 & 0 & -1 & 1 \\ G_2 & 0 & 0 & 1 & -1 & -1 & 1 & 0 & 0 \\ \hline A_1 & 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 \\ A_2 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & -1 \\ J_1 & 0 & 0 & 0 & 0 & 1 & -1 & 0 & 0 \\ J_2 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & -1 \\ R & 0 & 0 & 0 & 0 & 0 & 2 & 0 & 2 \\ \hline \end{array} \\ \text{CS levels:} \\ \begin{array}{|c|cc|cc|cc|cc|} \hline & G_1 & G_2 & A_1 & A_2 & J_1 & J_2 & R \\ \hline G_1 & 0 & -1 & \frac{1}{2} & 0 & 1 & 0 & -1 \\ G_2 & -1 & 0 & 0 & \frac{1}{2} & 0 & 1 & -1 \\ \hline A_1 & \frac{1}{2} & 0 & -\frac{1}{2} & 0 & 0 & 0 & \frac{1}{2} \\ A_2 & 0 & \frac{1}{2} & 0 & -\frac{1}{2} & 0 & 0 & \frac{1}{2} \\ J_1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ J_2 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ R & -1 & -1 & \frac{1}{2} & \frac{1}{2} & 0 & 0 & -2 \\ \hline \end{array} \end{array} \quad (6.26)$$

Notice that each $U(1)_i$ gauge group has both monopole and anti-monopole operators, denoted by V_i^\pm . However, due to the dynamical mixed Chern-Simons term, the monopoles of $U(1)_1$ are not gauge invariant under $U(1)_2$, and vice versa. We fix this by “dressing” the monopoles with chiral fields. We then find the four gauge-invariant operators

	$V_1^+ \phi_s$	$V_1^- \phi_t$	$V_2^+ \phi_z$	$V_2^- \phi_w$	
A_1	0	-1	0	1	
A_2	0	1	0	-1	
J_1	1	-1	0	0	
J_2	0	0	1	-1	
R	0	2	0	2	

(6.27)

with flavor charges as indicated. These operators correspond to the light-blue dots in figure 53.

To complete the $SL(5)$ tetrahedron, we couple the central cluster to sixteen additional octahedra on the periphery, as shown in figure 54. We label the new octahedra with




Figure 54. Assembling the $SL(5)$ tetrahedron from a central cluster and 16 peripheral octahedra.

parameters $\{y_{ij}\}_{i,j=1}^4$ as shown. The precise gauge theory we obtain depends on the final boundary polarization for the whole ensemble.

The simplest final polarization keeps X_1 and X_2 from Π_{cluster} as position coordinates. We can re-write each of them as a sum of four coordinates on the final boundary; for example $X_1 = Z + W = (Z + Y'_{13} + Y''_{23}) + (W + Y''_{13} + Y'_{23}) + Y_{13} + Y_{23} - 2\pi i$ and $X_2 = S + T = (S + Y'_{33} + Y''_{43}) + (T + Y''_{33} + Y'_{43}) + Y_{33} + Y_{43} - 2\pi i$. For the fourteen remaining positions we take a subset of fourteen Y_{ij} coordinates from octahedra on the periphery (due to the internal gluing constraints, these boundary coordinates are not all independent). Then the final gauge theory is simply the cluster theory (6.26), coupled to 16 T_{\diamond} theories with chirals ϕ_{ij} ($1 \leq i, j \leq 4$) via a superpotential,

$$T_5[\Delta] : \begin{cases} T_{\text{cluster}} \otimes (T_{\diamond})^{16}, \text{ with } U(1)^{16} \text{ flavor symmetry preserved by the superpotential} \\ W = V_1^+ \phi_s \phi_{11} \phi_{12} \phi_{13} + V_1^- \phi_t \phi_{21} \phi_{22} \phi_{23} + V_2^+ \phi_z \phi_{31} \phi_{32} \phi_{33} + V_2^- \phi_w \phi_{41} \phi_{42} \phi_{43}. \end{cases} \quad (6.28)$$

Notice that the four products of chirals $\phi_{i1}\phi_{i2}\phi_{i3}$ are associated with vertices on the groups of peripheral octahedra that become identified with the four distinguished vertices of the central cluster during the gluing. The four superpotential couplings combine the operators on the two sides of the gluing.

6.4 The $K = 3$ bipyramid

We now come to the theory of the bipyramid for $K = 3$. We expect that by splitting the bipyramid into ideal tetrahedra two different ways we will find two different theories $T_3[\text{bip-2}]$ and $T_3[\text{bip-3}]$ that are mirror symmetric, generalizing the $K = 2$ examples of section 6.2. We will check that this is the case by relating $T_3[\text{bip-2}]$ and $T_3[\text{bip-3}]$ via a sequence of fundamental mirror symmetries for octahedra.

Of course, the $K = 3$ bipyramid theories depend on a choice of polarization. There are two somewhat natural options, shown in figure 55, and designated “equatorial” and “tropical”. Some of the main features of the theories in their two polarizations can be




Figure 55. The equatorial (left) and tropical (right) polarizations for the $K = 3$ bipyramid. We only specify the seven position coordinates, indicating them by dark-blue dots on the K -triangulation of the boundary. The light-blue dots are lattice points/coordinates that commute with the positions (and so are assigned operators as well).

summarized as:

pol.	2 tetrahedra, $T_3[\text{bip-2}]$				3 tetrahedra, $T_3[\text{bip-3}]$			
	gauge	flavor	# chirals	# W ops	gauge	flavor	# chirals	# W ops
Π_{eq}	$U(1)^3$	$U(1)^7$	8	1	$U(1)^3$	$U(1)^7$	12	5
Π_{trop}	—	$U(1)^7$	8	1	$U(1)^6$	$U(1)^7$	12	5

(6.29)

Just like in the $K = 2$ case, we switch from one polarization to the other by performing an $\text{Sp}(14, \mathbb{Z})$ transformation simultaneously on both sides — in this case, gauging three $U(1)$ flavor symmetries. In *any* polarization, we will always have mirror symmetry

$$T_3[\text{bip-2}] \simeq T_3[\text{bip-3}]. \quad (6.30)$$

We begin by describing the theories on the two sides, and then explain how the mirror symmetry works.

6.4.1 Two tetrahedra

Let us focus on the tropical polarization, since it makes the two-tetrahedron theory especially simple. The decomposition of the bipyramid into two tetrahedra, and subsequently into eight octahedra, is shown in figure 56. It has a single internal lattice point, lying in the center of the identified faces of the octahedra. Quite conveniently, we can label the octahedra so that their initial positions Z_i coincide with the final positions X_i in the polarization. Specifically, we can set $X_i = Z_i$ for $1 \leq i \leq 7$. Moreover, the central gluing function $C = Z_3 + Z_4 + Z_5 + Z_6 + Z_7 + Z_8 - 2\pi i$ is made of six octahedron positions (with no octahedron momenta).

This situation is exactly analogous to the XYZ model of section 6.2 or to the $K = 4$ tetrahedron theory of section 6.3. We immediately recognize

$$T_3[\text{bip-2}, \Pi_{\text{trop}}] : \begin{cases} \text{Theory of eight chirals } \phi_i \text{ with a superpotential } W = \phi_3\phi_4\phi_5\phi_6\phi_7\phi_8, \\ \text{preserving } U(1)^7 \text{ flavor symmetry.} \end{cases} \quad (6.31)$$




Figure 56. Decomposition into eight octahedra, with Π_{trop} .




Figure 57. Decomposition into 12 octahedra, with Π_{trop} .

The first two flavor symmetries $U(1)_1, U(1)_2$ rotate the decoupled chirals ϕ_1, ϕ_2 ; while the remaining five symmetries $\{U(1)_i\}_{i=3}^7$ rotate ϕ_i and ϕ_8 with charges $(+1, -1)$.

We could proceed further and specify seven momenta P_i for the polarization Π_{trop} . They must be some linear combinations of boundary coordinates. However, the choice only affects background Chern-Simons levels, of which we won't keep track in this example.

If we used the equatorial polarization rather than the tropical polarization, the three pairs of octahedra that are glued together along a face would become three copies of $\mathcal{N} = 2$ SQED, as on the left of figure 50. (In the tropical polarization, we could have said that these pairs of octahedra were three $\mathcal{N} = 4$ hypermultiplets, as on the left of figure 49.) Nothing happens with the two octahedra at the poles. Then we would get

$$T_3[\text{bip-2}, \Pi_{\text{eq}}] : \begin{cases} \text{Three copies of SQED with chirals } \{q_i, \tilde{q}_i\}_{i=1}^3, \\ \text{plus two additional decoupled chirals } \phi_1, \phi_2; \\ W = (q_1 \tilde{q}_1)(q_2 \tilde{q}_2)(q_3 \tilde{q}_3), \text{ preserving } U(1)^7 \text{ flavor sym.} \end{cases} \quad (6.32)$$

The superpotential is a product of the three mesons.

6.4.2 Three tetrahedra

The decomposition of the bipyramid into three tetrahedra, and subsequently into twelve octahedra, is shown in figure 57. Now there are five gluing constraints: three on the glued pairs of tetrahedron faces, and two on the internal edge. The theory can be assembled in a few steps, much like the $K = 5$ tetrahedron theory of section 6.3. We will first build a central “core” out of six octahedra surrounding the central edge. Then we will couple the core to a halo of six octahedra around the equator (see figure 59).

To construct the core, we put together two copies of $\mathcal{N} = 4$ SQED, exactly as in (6.9), corresponding to the right side of figure 49. We call the copies $T_{\text{SQED}}^{\mathcal{N}=4}$ and $T_{\text{SQED}}^{\mathcal{N}=4'}$. Each of them already has a cubic superpotential enforcing a gluing constraint, so there are no additional constraints in assembling them. However, we should change the polarization in order to allow more boundary points on the core to carry gauge-invariant operators, as indicated in figure 58. We choose two new positions $\tilde{X}_1 = X_- + X'_+$ and $\tilde{X}_2 = P_- - P'_+$ (and conjugate momenta $\tilde{P}_1 = P_-$, $\tilde{P}_2 = X'_+$) in addition to the old positions X_+ and X'_- (and conjugate momenta P_+ , P'_-). In matrix form, we have

$$\begin{pmatrix} \tilde{X}_1 \\ \tilde{X}_2 \\ \tilde{P}_1 \\ \tilde{P}_2 \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & -1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix} \begin{pmatrix} X_- \\ X'_+ \\ P_- \\ P'_+ \end{pmatrix} =: g \cdot \begin{pmatrix} X_- \\ X'_+ \\ P_- \\ P'_+ \end{pmatrix}. \quad (6.33)$$

This amounts to gauging the anti-diagonal $U(1)_{\bar{d}}$ of the global symmetry group $U(1)_- \times U(1)'_+$. The operators V_+ and V'_- are clearly uncharged under $U(1)_{\bar{d}}$, as is the product $\eta = V_- V'_+$, so they survive the gauging. In addition, we gain a new pair μ_{\pm} of monopole/anti-monopole operators²² for $U(1)_{\bar{d}}$, charged under a new topological symmetry $U(1)_J$. They are associated with the boundary coordinates shown in figure 58. In summary, the core theory theory is two copies of SQED with a new polarization, $T_{\text{core}} = g \circ (T_{\text{SQED}}^{\mathcal{N}=4} \otimes T_{\text{SQED}}^{\mathcal{N}=4'})$, and explicitly

$$T_{\text{core}} : \left\{ \begin{array}{l} U(1)^3 \text{ gauge theory, with } U(1)_+ \times U(1)'_- \times U(1)_d \times U(1)_J \text{ symmetry,} \\ \text{six chirals, and superpotential } W = \tilde{\phi}\varphi\phi + \tilde{\phi}'\varphi'\phi' \end{array} \right.$$

	ϕ	$\tilde{\phi}$	φ	ϕ'	$\tilde{\phi}'$	φ'	V_+	V'_-	η	μ_+	μ_-
G	1	-1	0	0	0	0	0	0	0	0	0
G'	0	0	0	1	-1	0	0	0	0	0	0
$G_{\bar{d}}$	1	0	-1	-1	0	1	0	0	0	0	0
X_+	-1	0	1	0	0	0	1	0	0	0	-1
X'_-	0	0	0	-1	0	1	0	1	0	0	-1
D	-1	0	1	0	0	0	0	0	1	0	-1
J	0	0	0	0	0	0	0	0	0	1	-1
R	2	0	0	2	0	0	0	0	0	0	2

	G	G'	$G_{\bar{d}}$	X_+	X'_-	D	J	R
G	0	0	$\frac{1}{2}$	$\frac{1}{2}$	0	$-\frac{1}{2}$	0	0
G'	0	0	$\frac{1}{2}$	0	$-\frac{1}{2}$	0	0	0
$G_{\bar{d}}$	$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0	1	0
X_+	$\frac{1}{2}$	0	0	0	0	0	0	0
X'_-	0	$-\frac{1}{2}$	0	0	0	0	0	0
D	$-\frac{1}{2}$	0	0	0	0	0	0	0
J	0	0	1	0	0	0	0	0
R	0	0	0	0	0	0	0	*

(6.34)

²²The anti-monopole μ_- for $U(1)_{\bar{d}}$ turns out not to be invariant under the old gauge symmetries of the two copies of SQED, and must be “dressed” by fundamental chirals to fix this, $\mu_- \rightarrow \mu_- \phi \phi'$. The situation is similar to that for monopoles of the $SL(5)$ tetrahedron theory in section 6.3.3. In the tables below, and in the figure, it is the dressed, gauge-invariant monopole that is denoted μ_- .




Figure 58. The core of the $SL(3)$ bipyramid, built from two copies of $\mathcal{N} = 4$ SQED.




Figure 59. Assembling the $SL(3)$ bipyramid from a core and a halo.

Here we have denoted by $U(1)_d$ (or ‘ D ’) the diagonal of $U(1)_- \times U(1)_+$ that is preserved as a flavor symmetry.

In order to complete the bipyramid, we must add a “halo” of six octahedra to the core of figure 58. This is shown in figure 59. Three new gluing constraints arise, and each of the operators η, μ_+, μ_- from the core participates in one of three new superpotential couplings.

For the tropical polarization on the final boundary, we should represent the halo as three copies of $\mathcal{N} = 2$ SQED (cf. (6.12) and figure 50). The meson operators $q_i \tilde{q}_i$ in the three copies of SQED couple with the three operators η, μ_+, μ_- of the core in a superpotential. The theory becomes

$$T_3[\text{bip-3}] : \begin{cases} T_{\text{core}} \otimes T_{\text{halo}} \simeq [g \circ (T_{\text{SQED}}^{\mathcal{N}=4} \otimes T_{\text{SQED}}^{\mathcal{N}=4})] \otimes (T_{\text{SQED}}^{\mathcal{N}=2})^3, \\ U(1)^6 \text{ gauge theory with 12 chirals and the } U(1)^7 \text{ flavor symmetry preserved by} \\ W = \tilde{\phi} \phi + \tilde{\phi}' \phi' + q_1 \tilde{q}_1 \eta + q_2 \tilde{q}_2 \mu_+ + q_3 \tilde{q}_3 \mu_- . \end{cases} \quad (6.35)$$

The $U(1)^7$ flavor symmetry comes from $U(1)^4$ from the core and $U(1)^6$ from the halo, minus the three symmetries broken by the core-halo couplings. The positions in the final




Figure 60. The first two steps in going from two $SL(3)$ tetrahedra to three.

tropical polarization consist of X_+, X'_- from the core and any five of the six X_i, X'_i from the halo (after gluing, the six halo positions are not all independent). The theory has eight gauge-invariant operators associated with the external light- and dark-blue dots on the bipyramid: the monopoles V_+, V'_- of the core and the monopoles v_i^+, v_i^- ($i = 1, 2, 3$) for the three copies of SQED in the halo.

If instead we wanted to reproduce the equatorial polarization on the final boundary, we would just use a different polarization for the halo. In fact, in the equatorial polarization, the halo just becomes six free chirals (or three $\mathcal{N} = 4$ hypers). Three pairs of these chirals combine with the three core operators η, μ_+, μ_- to form three cubic superpotential couplings. We obtain a $U(1)^3$ gauge theory with a total of 12 fundamental chirals, five superpotential couplings, and a preserved $U(1)^7$ flavor symmetry.

6.4.3 Mirror symmetry

We have found an extremely simple bipyramid theory $T_3[\text{bip-2}]$, and a somewhat more complicated one $T_3[\text{bip-3}]$. To finish, let us proceed through the chain of fundamental mirror symmetries that link these two theories, demonstrating their IR equivalence. We simply follow the chain of 2–3 moves on octahedra shown in figure 47. We will be somewhat loose on the details, since the construction (ultimately) is completely systematic.

We work in the tropical polarization, and start with the theory $T_3[\text{bip-2}]$ (6.31) built from two tetrahedra and eight octahedra. Consider the three pairs of chirals (ϕ_3, ϕ_6) , (ϕ_4, ϕ_7) , (ϕ_5, ϕ_8) corresponding to the three pairs of octahedra that touch across the big internal face. They represent three $\mathcal{N} = 4$ hypers, coupled by an $\mathcal{N} = 2$ superpotential

$$W = \phi_3\phi_4\phi_5\phi_6\phi_7\phi_8. \quad (6.36)$$

The first set of three $2 \rightarrow 3$ moves dualizes the hypers to three copies of $\mathcal{N} = 4$ SQED (figure 60). Let us denote the charged and neutral chirals in the $\mathcal{N} = 4$ SQED's as $(q_i, \tilde{q}_i, \varphi_i)_{i=1}^3$. Each hyper is replaced with a pair of SQED monopole/anti-monopole operators, e.g. $(\phi_3, \phi_6) \rightarrow (v_3^+, v_3^-)$; while the three quadratic operators $(\phi_3\phi_6, \phi_4\phi_7, \phi_5\phi_8)$ are replaced by the neutral chirals $(\varphi_1, \varphi_2, \varphi_3)$. Thus the superpotential becomes

$$W = \varphi_1\varphi_2\varphi_3 + \sum_{i=1}^3 \tilde{q}_i\varphi_iq_i. \quad (6.37)$$

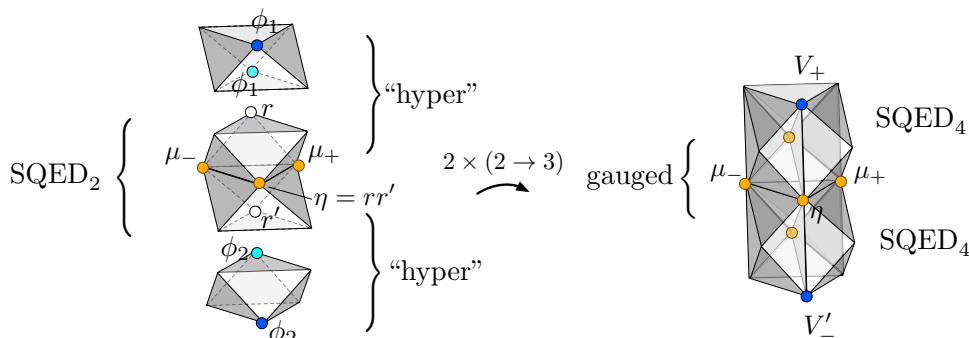


Figure 61. The final step in going from two $SL(3)$ tetrahedra to three: forming the core.

(We have added three $\mathcal{N} = 4$ SQED superpotentials to (6.36).)

Now we notice that there is a copy of the XYZ model embedded in the theory: it involves the three neutral chirals. We dualize it to a copy of $\mathcal{N} = 2$ SQED, performing the next $3 \rightarrow 2$ move (also in figure 60). Let us call the charged chirals of this SQED r and r' . One of the XYZ fields (say φ_3) is replaced by the meson $\eta = rr'$, while φ_1 and φ_2 are replaced by a monopole/anti-monopole pair μ_+, μ_- . The superpotential is now

$$W = \tilde{q}_1 q_1 \eta + \tilde{q}_2 q_2 \mu_+ + \tilde{q}_3 q_3 \mu_- . \quad (6.38)$$

(The XYZ superpotential has been eliminated in the duality.) The gauge group is $U(1)_{\bar{d}}$, acting on (r, r') .

The halo of six octahedra is now complete, and we just need to form the core. It requires the last set of two $2 \rightarrow 3$ moves, shown in figure 61. We bring in the two octahedra at the poles of the bipyramid, and think of each pair (r, ϕ_1) and (r, ϕ_2) as an $\mathcal{N} = 4$ hyper. We want to dualize these pairs to $\mathcal{N} = 4$ SQED. This is subtle: the chirals r, r' are *not* gauge-invariant, and we are effectively assuming that it is ok to interchange an order of limits (of gaugings) as we flow to the IR. With this assumption, the result of the duality is the core theory of (6.34): two copies of $\mathcal{N} = 4$ SQED, for which a combination of axial and topological symmetries — identified with $U(1)_{\bar{d}}$ — has been subsequently gauged.

The operators μ_+, μ_-, η all survive as monopoles of the core theory, as described in section 6.4.2. In addition, there are two new monopoles V_+ and V'_- coming from the copies of $\mathcal{N} = 4$ SQED. The gauge group has become $U(1)^3$, comprising the two $\mathcal{N} = 4$ SQED gauge groups and $U(1)_{\bar{d}}$. The superpotential gets two cubic $\mathcal{N} = 4$ SQED operators added to it:

$$W = \tilde{q}_1 q_1 \eta + \tilde{q}_2 q_2 \mu_+ + \tilde{q}_3 q_3 \mu_- + \tilde{\phi} \varphi \phi + \tilde{\phi}' \varphi' \phi' . \quad (6.39)$$

Altogether, we recover the coupled core-halo system that describes $T_{\text{bip-3}}^{(3)}$.

7 Knot complement theories

We have finally arrived at the theories $T_K[M]$ associated with knot complements $M = S^3 \setminus \mathcal{K}$. More generally, we can take $M = \bar{M} \setminus \mathcal{K}$, where \bar{M} is any closed 3-manifold. Recall from section 5 that in the ideal triangulation of a knot complement, all big faces of

tetrahedra are completely glued, leaving behind a small torus boundary on which all the tetrahedron vertices sit.

Physically, the theory $T_K[M]$ should correspond to wrapping K M5-branes on $\overline{M} \times \mathbb{R}^3$, in the ambient geometry $T^*\overline{M} \times \mathbb{R}^3 \times \mathbb{C}$. We then add K additional “probe” branes along $N^*\mathcal{K} \times \mathbb{R}^3$, where $N^*\mathcal{K} \subset T^*M$ is the conormal bundle of the knot.²³ This is the M-theory version of creating a small torus boundary — specifically one with boundary conditions relevant for parametrizing irreducible flat connections on M . The setup preserves four supercharges and flows to an effective $\mathcal{N} = 2$ theory on \mathbb{R}^3 in the infrared. Alternatively, we may start with the A_{K-1} (2,0) theory in six dimensions and compactify it on \overline{M} (with a suitable topological twist), adding a codimension-two defect along \mathcal{K} . In the language of [45] (see also [51, 95]), the codimension-two defect is of the type corresponding to a “full” or “maximal” puncture. From either construction, it is straightforward to deduce that $T_K[M]$ should be an SCFT with $SU(K)$ flavor symmetry.²⁴

As discussed in section 5 (and appendix A), a knot-complement theory can systematically be defined by extracting symplectic gluing data (g, σ) from an ideal triangulation of M , and the K -decomposition thereof. We will give some examples below in sections 7.3–7.4. The complexity of the theory does grow quite quickly with K . With N tetrahedra, our description of $T_K[M]$ will contain $\frac{1}{6}NK(K^2 - 1)$ chirals, and potentially a very large abelian gauge group. The superpotential will *also* contain $\frac{1}{6}NK(K^2 - 1)$ operators — corresponding to the internal lattice points of the octahedral decomposition — of which $K - 1$ do not break flavor symmetries. Nevertheless, despite this apparent complexity, we will be able to use combinatorics to deduce (and/or verify) several interesting general properties of knot-complement theories. For example, we can see immediately from the dimension of the boundary phase space $\mathcal{P}_K(T^2)$ that $T_K[M]$ has $U(1)^{K-1}$ flavor symmetry, which may be enhanced to $SU(K)$.

One particularly interesting question to ask is how theories $T_K[M]$ depend on K as K becomes large. Simply counting chiral multiplets in a UV description leads to the naive guess that degrees of freedom scale like $K(K^2 - 1)$. We can make this somewhat more precise by (conjecturally) relating degrees of freedom to the volume of flat $PGL(K, \mathbb{C})$ -connections on M — as given by the imaginary part of the Chern-Simons functional, $\text{Im} \int_M \text{Tr} (\mathcal{A}d\mathcal{A} + \frac{2}{3}\mathcal{A}^3)$. In turn, when M is hyperbolic, we will be able to see combinatorially that there is a flat $PGL(K, \mathbb{C})$ -connection whose volume is exactly $\frac{1}{6}K(K^2 - 1)$ times the hyperbolic volume of M . This adds some further evidence that the naive guess for the scaling is correct.

Let us observe that this scaling of degrees of freedom appears to be in very good agreement with results from the AdS/CFT correspondence. If M is a closed manifold, the gravity solutions dual to theories $T_K[M]$ should be given by the AdS_4 compactifications of M-theory that were obtained in [96] by considering the near-horizon limit of M5-branes

²³Constructions of this type, which use intersecting branes to embed a knot-complement geometry in M-theory, are quite common. They were introduced in [93], and have made many recent appearances, including [94] as well as [1, 54].

²⁴The relation between theories $T[M]$, compactifications of the 6d (2,0) theory, and M-theory is further explored in [15].

wrapping a Special Lagrangian 3-cycle. One explicit solution is known, in which the internal 7-manifold is an S^4 -fibration over a closed hyperbolic 3-manifold M (this is in fact the 11d uplift of a 7d supergravity solution originally found in [97]). The free energy $\mathcal{F} = -\log |\mathcal{Z}|$ can be determined from the computation of the effective 4d Newton constant using standard AdS/CFT techniques [32, 98] and is expected to be given essentially by $K^3 \text{Vol}(M)$, where $\text{Vol}(M)$ is the hyperbolic volume. In the particular case of knot complements, one would have to enhance the setup slightly by introducing the additional probe branes that intersect M along the knot \mathcal{K} , much as was done for 4d $\mathcal{N} = 2$ theories of class \mathcal{S} in [99]. However, as long as the knot complement M has nonzero hyperbolic volume, the $K^3 \text{Vol}(M)$ scaling should remain robust. Tantalizingly, a subleading correction to this result modifying $K^3 \rightarrow K(K^2 - 1)$ was predicted by the authors of [33] from the requirement of anomaly cancellation for M5-branes. This appears to match the exact scaling of volumes of flat connections mentioned above.

We remark that all the ideas and constructions in this section apply in a completely straightforward way to “link complements,” or manifolds M with multiple small torus boundaries. We focus on knot complements for simplicity.

7.1 Flavor symmetry and marginal operators of $T_K[M]$

It is expected on general grounds from the 3d-3d correspondence that a knot-complement theory $T_K[M]$ has

- at least $K - 1 = \text{rank}(\text{SU}(K))$ exactly marginal, chiral deformations;
- flavor symmetry $\text{SU}(K)$ at a distinguished point in the marginal parameter space;
- flavor symmetry $\text{U}(1)^{K-1}$ at generic points in the marginal parameter space.

The arguments leading to these properties were presented in [15, section 5.4], and are based on some general relations between marginal deformations and symmetries developed in [100]. The basic idea is as follows.

First consider the M -theory construction of $T_K[M]$ described a few pages above. The K probe M5'-branes that wrap the conormal bundle $N^*\mathcal{K} \subset T^*M$ have an $\mathfrak{su}(K)$ worldvolume symmetry, which translates to an $\text{SU}(K)$ flavor symmetry in the 3d theory $T_K[M]$. (This is directly analogous to the flavor symmetry associated to maximal defects in 4d theories of class \mathcal{S} [22, 45].) The K M5'-branes need not sit directly on top of each other, however. If they are separated in some direction along the plane \mathbb{C} (which is transverse to all the branes in the theory), the flavor symmetry will be broken to a maximal torus $\text{U}(1)^{K-1}$. The $K - 1$ complex separation parameters are the exactly marginal deformations.

Alternatively, one may think of constructing the theory $T_K[M]$ by starting with the four-dimensional theory $T_K[\partial M] = T_K[T^2]$, i.e. 4d $\mathcal{N} = 4$ $\text{SU}(K)$ super-Yang-Mills, and compactifying it on an interval. At one end of the interval, one inserts a half-BPS Dirichlet boundary condition that breaks the $\text{SU}(K)$ gauge symmetry to a flavor symmetry. At the other end of the interval, one can insert a more complicated quarter-BPS boundary condition \mathcal{B}_M that encodes the topology of the 3-manifold M , as mentioned on page 7 in the Introduction. As the interval is squeezed to zero size, one recovers $T_K[M]$.

Now, on the Dirichlet boundary condition, there is a hyperkähler triplet of moment maps associated to the $SU(K)$ flavor symmetry. With respect to the 3d $\mathcal{N} = 2$ symmetry preserved by the other boundary \mathcal{B}_M , two of the three real moment maps form a complex, chiral operator

$$\mu_{\mathbb{C}} \in \mathfrak{su}(K)_{\mathbb{C}} \simeq \mathfrak{sl}(K, \mathbb{C}), \quad (7.1)$$

valued in the Lie algebra of the symmetry. This descends to a chiral operator in the 3d theory $T_K[M]$, with exactly the correct R-charge to be added to the superpotential as a marginal deformation. Not all components of $\mu_{\mathbb{C}}$ are *exactly* marginal, however: this is where the analysis of [100] comes in. Adding a generic component of $\mu_{\mathbb{C}}$ breaks the flavor symmetry to a maximal torus $U(1)^{K-1}$ that commutes with this component. The currents for the broken symmetry generators pair up with and lift all the off-diagonal components of $\mu_{\mathbb{C}}$. The remaining diagonal components remain exactly marginal. They can freely be added to the superpotential without further breaking the flavor symmetry, and their coefficients parameterize the marginal directions discussed above.

The $K - 1$ marginal operators expected to exist in a knot-complement theory $T_K[M]$ can be found very explicitly in our constructions based on triangulations and K -decompositions! The argument is mostly combinatorial. Recall from sections 3.4 and 5.2 that an ideal triangulation of a knot complement M with N big tetrahedra, has exactly N internal edges as well. There are also $2N$ distinct internal faces. Therefore, after decomposing further into octahedra, the total number of internal gluing constraints C_k (i.e. lattice points) will be

$$(K - 1)(\# \text{ edges}) + \binom{K - 1}{2}(\# \text{ faces}) + \binom{K - 1}{3}(\# \text{ faces}) = \frac{1}{6}NK(K^2 - 1). \quad (7.2)$$

This is equal to the number of octahedra. However, not all of these gluing constraints are independent. As discussed in section 5.2.2 (cf. figure 41), $K - 1$ linear combinations of them will vanish automatically. Therefore, the final boundary phase space $\mathcal{P}_K(T^2)$ can be written as a symplectic reduction of the $\frac{1}{3}NK(K^2 - 1)$ -dimensional product octahedron phase space using any independent subset of the C_k as moment maps.

Consider what this means in terms of gauge theory. We build $T_K[M]$ by starting with $\frac{1}{6}NK(K^2 - 1)$ octahedron theories T_{\diamond} , applying a symplectic transformation (g, σ) to obtain a theory \tilde{T} with $U(1)^{\frac{1}{6}NK(K^2 - 1)}$ flavor symmetry, and then adding operators \mathcal{O}_k to the superpotential to break flavor symmetries. We should add one \mathcal{O}_k for every gluing constraint C_k . However, we find that after adding operators \mathcal{O}_k for any independent subset of $\frac{1}{6}NK(K^2 - 1) - (K - 1)$ gluing constraints, we are done. A $U(1)^{K-1}$ flavor symmetry remains. The remaining $(K - 1)$ operators can be added, but they do not break any additional flavor symmetries. Moreover, their R-charge is automatically set equal to two. Thus, they are marginal operators that do not break any flavor symmetry, and by the arguments of [100] they must be exactly marginal.

This field-theory analysis of enhanced symmetry (and related marginal operators) actually applies to $T_K[M]$ for any admissible M , not just to knot complements. The general idea is that *any* $T_K[M]$ can formally be constructed as a compactification of the 4d $\mathcal{N} = 2$

theory $T_K[\partial M]$ on an interval, with one boundary condition defined by the topology of M and the second, Dirichlet-like, boundary condition defined by the chosen polarization of ∂M . (Constructions of this sort are spelled out in more detail in [15].) The flavor symmetry of $T_K[M]$ is expected to match the flavor symmetry of the Dirichlet boundary condition. For example, if ∂M consists of a collection of big 3-holed spheres connected by small annuli, as on the r.h.s. of figure 4, there is a $U(1)^{K-1}$ flavor symmetry for each annulus that could be enhanced to $SU(K)$. In this case, the Dirichlet b.c. preserves the $SU(K)$ symmetry associated to an annulus/tube in the underlying 4d theory $T_K[\partial M]$. In contrast, if ∂M consists entirely of big boundary with holes filled in by small discs, then one only expects to find abelian flavor symmetry in $T_K[M]$, because the Dirichlet b.c. explicitly breaks symmetries of the 4d theory to an abelian subgroup.

We emphasize again that all these flavor symmetries are all associated to the *boundary* of M rather than M itself. This is directly analogous to the behavior of 4d $\mathcal{N} = 2$ theories of class \mathcal{S} , where flavor symmetries are associated to boundaries, or equivalently to punctures, of a Riemann surface rather than to the Riemann surface itself.

7.1.1 Real masses and boundary holonomy

To conclude this circle of ideas, we also comment briefly on the real (as opposed to complex) moment-map operators associated to the generic unbroken flavor symmetry $U(1)^{K-1}$ in a knot-complement theory $T_K[M]$. These operators can be added to the Lagrangian with coefficients m_1, \dots, m_{K-1} , which are the real masses for the $U(1)^{K-1}$ symmetry. In turn, the real masses have a geometric meaning directly related to flat $PGL(K, \mathbb{C})$ (or $SL(K, \mathbb{C})$) connections on the boundary ∂M . It is a direct generalization of the relation between flat connections and masses for 4d $\mathcal{N} = 2$ theories of class \mathcal{S} [45, 51, 95].

To concretely define a knot-complement theory $T_K[M]$ one must (as usual) choose a polarization Π of the boundary. In this case the choice is basically canonical. Indeed, by representing a manifold M with torus boundary as the complement of a knot \mathcal{K} in a specific closed manifold \overline{M} , we automatically identify a “meridian” cycle μ on the torus boundary. It is the cycle that forms an infinitesimally small loop linking \mathcal{K} in \overline{M} , and would be contractible if we removed \mathcal{K} . The eigenvalues of the meridian holonomy of a flat $PGL(K, \mathbb{C})$ connection then become canonical position coordinates for Π . The eigenvalues of the holonomy around a dual “longitude” cycle λ become canonical momentum coordinates, as explained in sections 3.4 and 5.3.

Roughly speaking, the exponentials of real masses $\exp(m_1), \dots, \exp(m_{K-1})$ are the magnitudes of the eigenvalues of the meridian holonomy. More precisely, after compactifying $T_K[M]$ to two dimensions (say, on a circle), the real masses gain a periodic imaginary part, and directly parameterize the complex holonomy eigenvalues.

If we want to keep $T_K[M]$ at a conformal point (that is, on the conformal manifold parametrized by marginal deformations), we should turn off all real mass deformations. In terms of the meridian holonomy, this corresponds to requiring trivial eigenvalues. This can

be done while keeping the actual holonomy matrix *nontrivial*, setting

$$\text{Hol}(\mu) \sim \begin{pmatrix} 1 & 0 & 0 & & 0 & 0 \\ 1 & 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 1 & & 0 & 0 \\ \vdots & & & \ddots & & \\ 0 & 0 & 0 & & 1 & 1 \end{pmatrix} \quad (7.3)$$

up to conjugation. This is compatible with the holonomy representation of a flat $\text{PGL}(K, \mathbb{C})$ connection on M being irreducible, as discussed e.g. in section 3.4. The matrix (7.3) is called the maximal parabolic element of $\text{PGL}(K, \mathbb{C})$, and can be described nicely as $\rho_{[K]} \cdot \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, where $\rho_{[K]} : \text{PGL}(2, \mathbb{C}) \rightarrow \text{PGL}(K, \mathbb{C})$ is the K -dimensional irreducible representation, or principal embedding.

7.2 K^3 scaling in $T_K[M]$

We would like to sketch an argument indicating that the degrees of freedom of $T_K[M]$ at a conformal point scale like $\frac{1}{6}K(K^2 - 1)$ when K is large, and M is hyperbolic. The argument has two parts. First we observe that, when M is hyperbolic, there is a flat $\text{PGL}(K, \mathbb{C})$ -connection on M with boundary conditions (7.3) whose volume is $\frac{1}{6}K(K^2 - 1)$ times the hyperbolic volume $\text{Vol}(M)$. Second, we try to relate the free energy of $T_K[M]$ on a sphere to the volume of this flat $\text{PGL}(K, \mathbb{C})$ -connection. The second part of the argument involves several nontrivial assumptions that we hope to investigate in subsequent work.

7.2.1 $\text{PGL}(K)$ volume from octahedra

The complex volume of a framed flat $\text{PGL}(K, \mathbb{C})$ -connection \mathcal{A} on a knot complement M was discussed in section 4.4. It can be interpreted as the value of the holomorphic Chern-Simons functional,

$$\mathcal{V}(\mathcal{A}) = i(\text{Vol}(\mathcal{A}) + i\text{CS}(\mathcal{A})) \sim \int_M \text{Tr} \left[\mathcal{A}d\mathcal{A} + \frac{2}{3}\mathcal{A}^3 \right], \quad (7.4)$$

with appropriate boundary conditions. When a connection has unipotent boundary holonomy, the real part is defined modulo $(2\pi i)^2\mathbb{Z}$ while the imaginary part is a well-defined real “volume.” When M is a hyperbolic 3-manifold, the unique complete hyperbolic metric can be re-written as a flat $\text{PGL}(2, \mathbb{C})$ -connection $\mathcal{A}_{\text{geom}}$. This geometric flat connection has parabolic holonomy at the meridian, $\text{Hol}(\mu) \sim \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, and its real volume equals the hyperbolic volume of M ,

$$\text{Vol}(\mathcal{A}_{\text{geom}}) = \text{Vol}(M). \quad (7.5)$$

In general there are a finite number of flat $\text{PGL}(2, \mathbb{C})$ -connections with parabolic holonomy, but the geometric one is distinguished as the connection with the *largest* volume.²⁵

On a hyperbolic manifold M there is also a distinguished conjugate connection $\mathcal{A}_{\text{conj}} = \overline{\mathcal{A}_{\text{geom}}}$. Its real volume is clearly $-\text{Vol}(\mathcal{A}_{\text{geom}})$, and it has the *smallest* volume of all connections with parabolic holonomy.

²⁵For physics discussions of the relation between hyperbolic geometry and complex Chern-Simons theory, see [6, 7, 11]. Classic mathematics references are [5, 8, 101].

Given the geometric $\mathrm{PGL}(2, \mathbb{C})$ -connection $\mathcal{A}_{\mathrm{geom}}$, we can always construct an irreducible $\mathrm{PGL}(K, \mathbb{C})$ -connection $\mathcal{A}_{\mathrm{geom}}^{(K)}$ on M by using the irreducible K -dimensional representation $\rho_{[K]}$ to embed $\mathrm{PGL}(2, \mathbb{C})$ into $\mathrm{PGL}(K, \mathbb{C})$. As explained above, $\rho_{[K]}$ simply maps one parabolic element to another; therefore, $\mathcal{A}_{\mathrm{geom}}^{(K)}$ will have parabolic holonomy (7.3) around the meridian. Moreover, the complex volume of $\mathcal{A}_{\mathrm{geom}}^{(K)}$ is exactly $\frac{1}{6}K(K^2 - 1)$ times $\mathcal{V}(\mathcal{A}_{\mathrm{geom}})$; in particular,

$$\mathrm{Vol}(\mathcal{A}_{\mathrm{geom}}^{(K)}) = \frac{1}{6}K(K^2 - 1)\mathrm{Vol}(\mathcal{A}_{\mathrm{geom}}) = \frac{1}{6}K(K^2 - 1)\mathrm{Vol}(M). \quad (7.6)$$

This fact was proved in [25]. It can be intuited from the r.h.s. of (7.4). Momentarily we will give a simple, constructive combinatorial argument for the validity of (7.6), similar to that used by [25].

In general, for a hyperbolic manifold M there may be (finitely) many flat $\mathrm{PGL}(K, \mathbb{C})$ -connections \mathcal{A}_α with parabolic meridian holonomy. They may or may not have anything to do with lower-rank embeddings. Explicitly, they can be described as the solutions to the classical gluing equations of an octahedral decomposition that have meridian eigenvalues $m_a = 1$; or as the points of intersection of the Lagrangian submanifold $\mathcal{L}_K(M)$ with the Lagrangian $\{m_a = 1\}$. Interestingly, however, the geometric connection $\mathcal{A}_{\mathrm{geom}}^{(K)}$ seems to have the largest volume of all the \mathcal{A}_α (and its conjugate $\mathcal{A}_{\mathrm{conj}}^{(K)} = \overline{\mathcal{A}_{\mathrm{geom}}^{(K)}}$ the smallest).²⁶

In order to demonstrate (7.6) using the combinatorics of this paper, we proceed as follows. Let M be a hyperbolic knot complement. Promote $\mathcal{A}_{\mathrm{geom}}$ to a framed flat connection (the choice of invariant flag at the torus boundary is unique); and choose an ideal triangulation $\{\Delta_i\}_{i=1}^N$ of M such that $\mathcal{A}_{\mathrm{geom}}$ induces well-defined cross-ratio coordinates for the octahedra of the 2-decomposition. Almost any triangulation will do. Recall that in the 2-decompositions, there is a single octahedron per tetrahedron; so we can label the octahedron parameters $\{z_i, z'_i, z''_i\}_{i=1}^N$, with $z_i z'_i z''_i = -1$ and $z''_i + z_i^{-1} - 1 = 0$. The hyperbolic volume of M is computed as

$$\mathrm{Vol}(M) = \sum_{i=1}^N \mathrm{Vol}(\diamond_i) = \sum_{i=1}^N \mathcal{L}i(z_i), \quad (7.7)$$

where $\mathcal{L}i(z)$ is the Bloch-Wigner dilogarithm function (4.57).

Now consider the K -decomposition of the same triangulation. Suppose that every octahedron in a big tetrahedron Δ_i is labeled with the same triple of parameters $\{z_i, z'_i, z''_i\}_{i=1}^N$ that encoded $\mathcal{A}_{\mathrm{geom}}$ above. We align all the primes and double-primes to coincide with their former locations in the 2-decomposition. Then every gluing constraint at a lattice point inside a tetrahedron Δ_i takes the form $z_i z'_i z''_i z_i z'_i z''_i = 1$ and is automatically satisfied. Every face gluing constraint takes the form $z_i z'_i z''_i z_j z'_j z''_j = 1$, which is automatic as well; and every edge gluing constraint is a copy of the corresponding edge constraint for $K = 2$, which is already satisfied. Since this uniform assignment of octahedron parameters satisfies all of the octahedron gluing constraints, it defines a framed flat $\mathrm{PGL}(K, \mathbb{C})$ -connection on

²⁶This was recently conjectured by S. Garoufalidis, M. Goerner, and C. Zickert, after computing volumes of boundary-parabolic flat connections for hundreds of hyperbolic knots.

M . Moreover, it is straightforward to check from the snake rules of section 3.3 that the holonomy around any closed cycle in M is simply the image of the $\mathcal{A}_{\text{geom}}$ holonomy under $\rho_{[K]}$. In particular, the meridian holonomy at the boundary is parabolic.

Therefore, this assignment of octahedron parameters at general K simply parametrizes the flat connection $\mathcal{A}_{\text{geom}}^{(K)}$. Its volume is easy to calculate, since every octahedron in a big tetrahedron gives the same contribution:

$$\text{Vol}(\mathcal{A}_{\text{geom}}^{(K)}) = \sum_{\text{all octahedra}} \text{Vol}(\diamond) = \frac{1}{6}K(K^2 - 1) \sum_{i=1}^N \mathcal{L}i_2(z_i) = \frac{1}{6}K(K^2 - 1)\text{Vol}(M). \tag{7.8}$$

7.2.2 From free energy to volumes

Now we come to the more nontrivial part of the scaling argument. A standard measure for the number of degrees of freedom in a three-dimensional theory is the S^3 partition function $\mathcal{Z}[T]$. The free energy $\mathcal{F} = -\log|\mathcal{Z}[T]|$ is extensive in the number of degrees of freedom, and has several extremely nice properties; for example, for an $\mathcal{N} = 2$ theory, it is maximized at the superconformal values of R-charges [102, 103]. Unfortunately, although the S^3 partition function can be computed systematically for theories $T_K[M]$ in class \mathcal{R} [1, 104], it is not immediately clear how it relates to the volumes of flat connections discussed above.

There does exist a closely related measure of degrees of freedom for $\mathcal{N} = 2$ theories that is related to volumes: the partition function $\mathcal{Z}_b[T]$ on an ellipsoid S_b^3 . The ellipsoid deformation stretches out the three-sphere, so that metrically it can be described by the embedding $b^2|z|^2 + b^{-2}|w|^2 = 1$ in flat $\mathbb{C}^2 \simeq \mathbb{R}^4$; for $b = 1$ we recover round S^3 . The partition functions of class- \mathcal{R} theories on S_b^3 can also be systematically calculated [1, 105]. For a 3-manifold theory, $\mathcal{Z}_b[T_K[M]]$ agrees with the analytically continued partition function of $\text{SL}(K, \mathbb{R})$ Chern-Simons theory on M , with appropriate boundary conditions at ∂M . Moreover, it follows from results of [1] (based on [7, 9–11]) that in the $b \rightarrow 0$ limit,

$$\mathcal{F}_b[T_K[M]] = -\log|\mathcal{Z}_b[T_K[M]]| \stackrel{b \rightarrow 0}{\sim} -\frac{1}{2\pi b} \text{Vol}(\mathcal{A}_\alpha) + O(1), \tag{7.9}$$

where \mathcal{A}_α is the volume of one of the flat $\text{PGL}(K, \mathbb{C})$ -connections on M with parabolic meridian holonomy. Note that the holographic free energy on S_b^3 is simply related to the free energy on S^3 by $\mathcal{F}_b = (b + 1/b)^2 \mathcal{F}/4$ [106].

In order to prove the desired large- K scaling from (7.9), we must make two assumptions — both of which are rather interesting and will be elucidated elsewhere. First, we must assume that (7.9) accurately measures the degrees of freedom of $T_K[M]$ at a superconformal fixed point, and in particular with the correct superconformal R-symmetry. Second, we need to know that the flat connection \mathcal{A}_α entering the r.h.s. of (7.9) is the conjugate flat connection $\mathcal{A}_{\text{conj}}^{(K)} = \overline{\mathcal{A}_{\text{geom}}^{(K)}}$ so that the volume becomes $-\text{Vol}(\mathcal{A}_{\text{conj}}^{(K)}) = +\frac{1}{6}K(K^2 - 1)\text{Vol}(M)$. The appearance of $\mathcal{A}_{\text{conj}}^{(K)}$ on the r.h.s. of (7.9) means that this connection dominates the

$SL(K, \mathbb{R})$ Chern-Simons partition function. There have been hints in [59] that this does happen for $K = 2$, though more is needed to understand the general story.²⁷

7.3 Trefoil knot invariants

We conclude this section with two examples of knot complements and the combinatorics of their $K = 3$ triangulations.

To keep things simple, we will use minimal triangulations of the trefoil and the figure-eight knot complements, into two tetrahedra (and eight octahedra) each. These triangulations are not quite sufficiently refined to define theories $T_3[M]$ due to a technical constraint discussed in [1, section 4.1] and appendix A: in order to construct superpotential operators \mathcal{O}_k corresponding to gluing constraints C_k , the constraints must be “easy.” This means that the functions C_k never contain a sum of two non-commuting octahedron parameters, such as $Z_i + 2Z_i''$. For $K \geq 4$ this is never a problem, but for $K = 2, 3$ the triangulations must be further refined. Nevertheless, even the unrefined triangulations are good enough for calculating observables of the theories — such as $\mathcal{L}_{\text{SUSY}} \simeq \mathcal{L}_K(M)$, volumes, wavefunctions, etc. — which is what we focus on here.

7.3.1 $SL(2)$ invariants

Let us briefly review the structure of trefoil invariants for $K = 2$.

The moduli space of flat $SL(2, \mathbb{C})$ -connections on the trefoil knot complement $M = S^3 \setminus \mathbf{3}_1$ has two irreducible components of dimension one, corresponding (respectively) to flat connections whose holonomy is abelian and to flat connections that are genuinely nonabelian and irreducible. Each component projects to a one-dimensional Lagrangian submanifold in the boundary phase space $\mathcal{P}_2(\partial\mathbf{3}_1)$. If we parameterize the meridian and longitude eigenvalues as

$$\mu : \begin{pmatrix} m & \\ & m^{-1} \end{pmatrix}, \quad \lambda : \begin{pmatrix} \ell & \\ & \ell^{-1} \end{pmatrix}, \quad (7.10)$$

so that (in exponentiated coordinates) $\mathcal{P}_2(\partial\mathbf{3}_1) = \{(m, \ell)\} \simeq \mathbb{C}^* \times \mathbb{C}^*$, one finds that the two Lagrangians are

$$\begin{aligned} \mathcal{L}_{\text{abelian}} : \quad \ell - 1 &= 0 \\ \mathcal{L}_2(\mathbf{3}_1) = \mathcal{L}_{\text{non-abelian}} : \quad \ell + m^6 &= 0. \end{aligned} \quad (7.11)$$

The union of these two components is the well-known A-polynomial of the trefoil [16]. Notice that the Lagrangians have Weyl symmetry, i.e. they are invariant under $(\ell, m) \rightarrow (\ell^{-1}, m^{-1})$. However, if we think of $\mathcal{P}_2(\partial\mathbf{3}_1)$ as a moduli space of *framed* flat connections we do not quotient out by the Weyl group.

²⁷The question of which complex critical points dominate a real partition function is a tricky and subtle one, and closely related to the (still open) Volume Conjecture for compact Chern-Simons theory [107, 108]. The Volume Conjecture could be proved physically if it were known that geometric flat connections dominated $SU(2)$ partition functions [7, 109].

Only the non-abelian flat connections can be parameterized by cross-ratio coordinates in an ideal triangulation. Indeed, the Lagrangian $\mathcal{L}_2(\mathbf{3}_1)$ obtained from an ideal triangulation agrees with $\mathcal{L}_{\text{non-abelian}}$.

In order to discuss $\text{PGL}(2, \mathbb{C})$ flat connections rather than $\text{SL}(2, \mathbb{C})$ flat connections, we would simply re-write the above equations in terms of the invariant squared eigenvalues m^2 and ℓ^2 . It is useful to note that the Lagrangians are *already* written in terms of m^2 . Conversely, in order to lift from $\text{PGL}(2, \mathbb{C})$ to $\text{SL}(2, \mathbb{C})$, one only needs to take an appropriate square root of ℓ^2 , and not m^2 . The pattern for general K is that the equations for $\text{SL}(K, \mathbb{C})$ Lagrangians involve the individual longitude eigenvalues but only the K -th powers of the meridian eigenvalues. This is true for any knot complement in a homology sphere.²⁸ Therefore, to lift from $\text{PGL}(K, \mathbb{C})$ to $\text{SL}(K, \mathbb{C})$ one only needs to “take a K -th root of the $\text{PGL}(K, \mathbb{C})$ longitude.”

The full Lagrangian $(\ell - 1)(m^6 \ell + 1)$ has been quantized in the context of $\text{SU}(2)$ Chern-Simons theory [7, 74, 75]. Using combinatorics of ideal triangulations, one can quantize $\mathcal{L}_{\mathbf{3}_1}^{(2)}$, obtaining $\hat{\ell} + q^{\frac{3}{2}} \hat{m}^6 \simeq 0$. One can then proceed to calculate analytically continued Chern-Simons partition functions, *a.k.a.* sphere partition functions and indices of the trefoil-knot theory. However, they are all extremely simple. For example, the index just looks like

$$\mathcal{I}_{\mathbf{3}_1}^{(2)}(m, \zeta; q) = \zeta^{3m}. \tag{7.12}$$

in the conventions of [54]. This ultimately reflects the fact that the trefoil is *not* a hyperbolic knot, and we do not expect $T_2[\mathbf{3}_1]$ (or for that matter $T_K[\mathbf{3}_1]$) to contain many interesting degrees of freedom in the infrared.

7.3.2 The $\text{SL}(3)$ character variety

Now consider the full moduli space of flat $\text{SL}(3, \mathbb{C})$ -connections. A straightforward but extremely lengthy computation²⁹ shows that the moduli space has five irreducible components of dimension two on which the two meridian eigenvalues are allowed to vary freely. If we parametrize the boundary eigenvalues as

$$\mu : \begin{pmatrix} m_1 & & \\ & m_2 & \\ & & \frac{1}{m_1 m_2} \end{pmatrix}, \quad \lambda : \begin{pmatrix} \ell_1 & & \\ & \ell_2 & \\ & & \frac{1}{\ell_1 \ell_2} \end{pmatrix}, \tag{7.13}$$

then the five components project to the Lagrangian submanifolds

$$\begin{aligned} (1) : & \quad \{ \ell_1 - 1 = 0, \quad \ell_2 - 1 = 0 \} \\ (2_1) : & \quad \{ \ell_1 - 1 = 0, \quad \ell_2 + m_1^3 m_2^6 = 0 \} \\ (2_2) : & \quad \{ \ell_1 + m_1^6 m_2^3 = 0, \quad \ell_2 - 1 = 0 \} \\ (2_3) : & \quad \{ m_2^3 \ell_1 + m_1^3 = 0, \quad m_1^3 \ell_2 + m_2^3 = 0 \} \\ (3) : & \quad \{ \ell_1 - m_1^6 = 0, \quad \ell_2 - m_2^6 = 0 \} = \mathcal{L}_3(\mathbf{3}_1). \end{aligned} \tag{7.14}$$

²⁸It is true because for knot complements in homology sphere one can always tensor a flat \mathbb{C}^K bundle with a flat complex line bundle whose holonomy around the meridian is any K -th root of unity.

²⁹A first-principles computation of this moduli space involves looking at representations of $\pi_1(M)$ into $\text{SL}(3, \mathbb{C})$, and imposing the relations in $\pi_1(M)$ as algebraic constraints on elements of $\text{SL}(3, \mathbb{C})$ matrices.




Figure 62. Construction of the trefoil knot complement from two $SL(3)$ tetrahedra. Octahedron edges with matching colors are glued together.

Branch (1) contains the $SL(3, \mathbb{C})$ flat connections with abelian holonomy; branches (2_i) contain the flat connections whose holonomy reduces to an $SL(2, \mathbb{C}) \times U(1)$ subgroup, and branch (3) contains the connections whose holonomy is irreducible. Thus, using the combinatorics of this paper, we expect to be able to reproduce (and associate a 3d theory to) branch (3). Branches (1) and (3) are invariant under the S_3 Weyl symmetry of $SL(3, \mathbb{C})$, while the three branches (2_i) are permuted among themselves.

Note that each of these branches contains embeddings of $SL(2, \mathbb{C})$ flat connections along one-dimensional slices. For example, if we set $\ell_1 = \ell_2$ and $m_1 = m_2$ in branch (3), we obtain the principal embedding (using the representation $\rho_{[3]}$) of the nonabelian $SL(2, \mathbb{C})$ -connections.

We may also observe that, as predicted, the Lagrangian equations only depend on the *cubes* of the meridian eigenvalues. To reduce to $PGL(3, \mathbb{C})$ flat connections we could rewrite the equations in terms of the invariants ℓ_1^3 and ℓ_2^3 as well.

7.3.3 Symplectic gluing data

Let us see how we can extract symplectic gluing data from the trefoil triangulation at $K = 3$, in a (relatively) much more efficient manner.

The complement of the trefoil knot can be decomposed into two tetrahedra, as shown in the figure 62. There are eight gluing functions, corresponding to lattice points on internal faces and edges, which we write in exponentiated form as

$$c_1 = z_3 y_3, \quad c_5 = z_1 z_4' z_3'' y_2 y_3' y_4'', \quad (7.15)$$

$$c_2 = z_4 y_4, \quad c_6 = z_2 z_3' z_4'' y_1 y_4' y_3'', \quad (7.16)$$

$$c_3 = z_1 z_1' z_2'' z_4' z_4'' y_1 y_1' y_2'' y_4' y_4'', \quad c_7 = z_3 z_2' z_1'' y_4 y_1' y_2'', \quad (7.17)$$

$$c_4 = z_1'' z_2 z_2' z_3'' z_3'' y_1'' y_2 y_2' y_3' y_3'', \quad c_8 = z_4 z_1' z_2'' y_3 y_2' y_1''. \quad (7.18)$$

(In logarithmic form, each of these is a sum of octahedron parameters, minus $2\pi i$; e.g., $C_1 = Z_3 + Y_3 - 2\pi i$.) Notice that there are two automatic relations $c_1 c_2 c_3 c_4 = 1$ and $c_5 c_6 c_7 c_8 = 1$. Thus, as we just emphasized in section 7.1, only six gluing constraints are independent.

As explained in section 5.2.4, one can obtain the U_a and V_a coordinates that parametrize meridian and longitude eigenvalues by taking two global slices parallel to the small torus boundary. The slices are drawn in figure 63, along with meridian and longitude

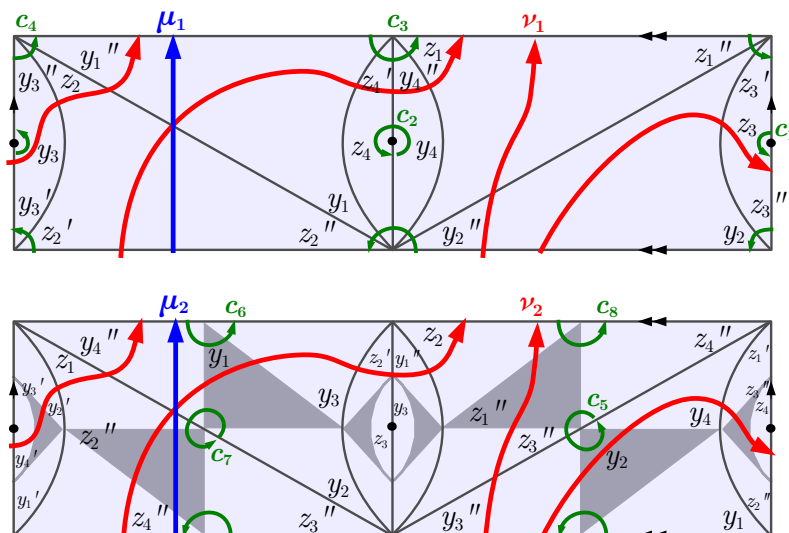


Figure 63. First and second slices of the torus boundary for the trefoil knot. The paths μ_a and ν_a represent the meridian and longitude, while the small circular paths drawn in green correspond to the gluing constraints.

paths on each slice. (We also see the eight gluing functions as coming from small circular paths, and we can check that each of the relations $c_1 c_2 c_3 c_4 = 1$, and $c_1 c_2 c_3 c_4 c_5 c_6 c_7 c_8 = 1$ involve the gluing functions on a single slice.) In exponentiated form, we find the meridian and longitude eigenvalues to be

$$e^{U_1} = \frac{y_1''}{z_2''}, \quad e^{V_1} = \frac{z_1 z_2 z_4' y_1'' y_2'' y_4''}{z_1'' z_2'' z_3'' y_1 y_2 y_3'}, \quad (7.19)$$

$$e^{U_2} = \frac{z_2'' y_4''}{z_1'' z_4''}, \quad e^{V_2} = \frac{z_1 y_2 (y_4'')^2}{z_1'' z_2'' (z_4'')^2 y_1' y_2' y_3^3 y_4}. \quad (7.20)$$

In order to have canonical Poisson brackets, we redefine the longitude coordinates as in section 5.3.2:

$$\tilde{V}_1 := \frac{1}{3}(2V_1 + V_2 + C_1 - C_2 + C_6), \quad (7.21a)$$

$$\tilde{V}_2 := \frac{1}{3}(V_1 + 2V_2 - C_1 + C_2 - C_6). \quad (7.21b)$$

One can check that now $\{\tilde{V}_a, U_b\} = \delta_{ab}$. The division by $K = 3$ here (i.e. by the determinant of the Cartan matrix) is a convenient way of lifting from $\text{PGL}(3, \mathbb{C})$ representations to $\text{SL}(3, \mathbb{C})$ representations. Moreover, the judicious addition of gluing constraints (which are ultimately set to zero) in (7.21) allows \tilde{V}_1, \tilde{V}_2 to be *integer* linear combinations of octahedron parameters after the division.

Note that holonomy eigenvalues are now parametrized in several different ways. For the meridian, the cubes of the eigenvalues in (7.13) are related to the path-coordinates as

$$m_1^3 = \exp(-2U_1 - U_2), \quad m_2^3 = \exp(U_1 - U_2), \quad \frac{1}{m_1^3 m_2^3} = \exp(U_1 + 2U_2); \quad (7.22)$$

while for the longitude the eigenvalues themselves are

$$\ell_1 = \exp(-\tilde{V}_1), \quad \ell_2 = \exp(\tilde{V}_1 - \tilde{V}_2), \quad \frac{1}{\ell_1 \ell_2} = \exp \tilde{V}_2. \quad (7.23)$$

We obtain the classical Lagrangian submanifold corresponding to the space of flat, framed, irreducible $\mathrm{SL}(3, \mathbb{C})$ -connections on the trefoil knot complement by rewriting the product Lagrangian $\mathcal{L}_\times = \{z_i'' + z_i^{-1} - 1 = 0, y_i'' + y_i^{-1} - 1 = 0\}_{i=1}^4$ in terms of the holonomy eigenvalues m_a and ℓ_a and making use of the classical gluing constraints $c_k = 1$. The result is

$$\mathcal{L}_3(\mathbf{3}_1) = \{\ell_1 - m_1^6 = 0, \ell_2 - m_2^6 = 0\}. \quad (7.24)$$

It is precisely branch (3) of (7.14).

Finally, we can summarize the gluing data in the affine-symplectic $\mathrm{ISp}(16, \mathbb{Z})$ transformation (g, σ) that transforms the initial polarization with positions $(Z_1'', Z_2'', Z_3'', Z_4'', Y_1'', Y_2'', Y_3'', Y_4'')$ and momenta $(Z_1', Z_2', Z_3', Z_4', Y_1', Y_2', Y_3', Y_4')$ to the final (canonical!) polarization

$$\Pi = (U_1, U_2, C_1, C_2, C_3, C_5, C_6, C_7; \tilde{V}_1, \tilde{V}_2, \Gamma_1, \Gamma_2, \Gamma_3, \Gamma_5, \Gamma_6, \Gamma_7). \quad (7.25)$$

We obtain

$$g_3^{\mathbf{3}_1} = \begin{pmatrix} 0 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & -1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \\ -1 & 1 & 0 & 1 & -1 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 & 0 & -1 & -1 & 1 & -1 & 0 & 1 & 1 & 0 & -1 & -1 & 0 & 0 & 0 & 0 \\ 0 & -1 & -1 & 1 & -1 & 0 & 0 & 0 & 0 & 0 & -1 & -1 & 0 & -1 & 0 & 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 0 & -1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & -1 & 0 & 0 \\ \hline -2 & -2 & 0 & 0 & 1 & 1 & 0 & 2 & -1 & -1 & -1 & 1 & 0 & 0 & 1 & 1 & -1 & 0 & 0 \\ -2 & -1 & 0 & -2 & 1 & 0 & -2 & 2 & -1 & 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 1 & 0 & 0 & 0 & 0 & -1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & -1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & -1 & 1 & 1 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & -1 & 0 & 1 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad \sigma_3^{\mathbf{3}_1} = \begin{pmatrix} 0 \\ 2 \\ 0 \\ -2 \\ -2 \\ -3 \\ -3 \\ -1 \\ \hline 4 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (7.26)$$

7.4 Figure-eight knot gluing

Since the figure-eight knot is the simplest hyperbolic knot and a fairly standard example in knot theory, we briefly demonstrate the construction of $K = 3$ gluing data for it as well.

Like for the trefoil, the complement of the figure-eight knot can be triangulated into two tetrahedra (figure 64). One of the most interesting properties of the figure-eight knot is that the geometric flat $\mathrm{PGL}(2, \mathbb{C})$ -connection $\mathcal{A}_{\mathrm{geom}}$ with parabolic holonomy is represented by two regular ideal tetrahedra — *i.e.* two octahedra that have all coordinates $z_i = z_i' = z_i'' \equiv \exp \frac{i\pi}{3}$. This value of z_i maximizes the volume function $\mathcal{L}i_2(z_i)$. For general K , the complement of the figure-eight knot can similarly be decomposed into $2 \times \frac{1}{6}K(K^2 - 1)$ “regular” octahedra, again with all $z_i = z_i' = z_i'' \equiv \exp \frac{i\pi}{3}$. (One can check explicitly that this is a solution to the gluing equations below.)




Figure 64. Construction of the figure-eight knot complement from two $SL(3)$ tetrahedra.




Figure 65. First and second slices of the torus boundary for the figure-eight knot.

Using the $K = 3$ decomposition into tetrahedra and octahedra shown in figure 64, we find that the $K = 3$ gluing constraints are

$$c_1 = z_1 z_1' z_3' y_2' y_4 y_4', \quad c_5 = z_2 z_4' z_3'' y_4 y_2' y_1'', \quad (7.27)$$

$$c_2 = z_2 z_2' z_4' y_1 y_3 y_3', \quad c_6 = z_1 z_3' z_4'' y_3 y_1' y_2'', \quad (7.28)$$

$$c_3 = z_2'' z_3 z_3'' y_1 y_1' y_4'', \quad c_7 = z_4 z_2' z_1'' y_1 y_3 y_3' y_4', \quad (7.29)$$

$$c_4 = z_1'' z_4 z_4'' y_2 y_2' y_3'', \quad c_8 = z_3 z_1' z_2'' y_2 y_4 y_4' y_3'', \quad (7.30)$$

and satisfy the relations $c_1 c_2 c_3 c_4 = 1$ and $(c_1 c_2 c_3 c_4) c_5 c_6 c_7 c_8 = 1$, one for each slice.

The coordinates u_a and v_a for the meridian and longitude paths drawn in figure 65 are

$$u_1 = \frac{z_3'}{y_4''}, \quad u_2 = \frac{z_2' y_4''}{y_1'' y_3''}, \quad (7.31)$$

$$v_1 = \frac{z_1' z_2' y_3 y_4}{z_3'' z_4'' y_1 y_2}, \quad v_2 = \frac{z_3' z_4' y_1 y_2}{z_1'' z_2'' y_3 y_4}. \quad (7.32)$$

We redefine the longitude coordinates to have canonical Poisson brackets:

$$\tilde{V}_1 \equiv \frac{1}{3}(2V_1 + V_2 - C_5 - C_6), \quad (7.33)$$

$$\tilde{V}_2 \equiv \frac{1}{3}(V_1 + 2V_2 + C_5 + C_6). \quad (7.34)$$

Taking the initial polarization with positions $\{Z''_1, Z''_2, Z''_3, Z''_4, Y'_1, Y'_2, Y'_3, Y'_4\}$, we find the following symplectic gluing matrix, with all integer coefficients:

$$g_3^{4_1} = \left(\begin{array}{cccccccc|cccccccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 1 & -1 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & -1 & 0 \\ -1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & -1 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 & 0 & -1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & -1 & 1 & 0 & -1 & 1 & 0 & 0 & 0 & -1 & 0 & 1 & -1 & 0 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 & 1 & -1 & 0 & 0 & -1 & 0 & 1 & 0 & 0 & -1 & 1 & 0 & 0 \\ 1 & 0 & 0 & -1 & 0 & 0 & 1 & -1 & 0 & 0 & 1 & 0 & -1 & 1 & 0 & 0 & -1 \\ \hline 0 & 0 & -1 & -1 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & -1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ -1 & 2 & 0 & 1 & -1 & 0 & -1 & 0 & 0 & -1 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & -1 & -1 & 0 & 0 & 0 & 0 & -1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & -1 & -1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & -1 & 0 & 0 & 0 & 0 \\ -1 & 0 & 0 & 1 & 0 & 0 & -1 & 0 & 0 & 0 & -1 & 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 2 & -1 & 2 & 0 & -1 & -2 & 1 & 0 & 0 & -1 & -1 & 0 & -1 & 0 & 0 & 0 \end{array} \right), \quad \sigma_3^{4_1} = \begin{pmatrix} 1 \\ 3 \\ -1 \\ -1 \\ -3 \\ -2 \\ -2 \\ -2 \\ -2 \\ -2 \\ 4 \\ 4 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}. \quad (7.35)$$

Acknowledgments

We would like to thank Christopher Beem, Nicolas Bergeron, Elisha Falbel, Davide Gaiotto, Stavros Garoufalidis, Antonin Guilloux, Sergei Gukov, Lotte Hollands, Daniel Jafferis, Juan Maldacena, Greg Moore, Andrew Neitzke, Nathan Seiberg, Roland van der Veen, Edward Witten, and Don Zagier for helpful comments and discussions. We would especially like to thank Dan Xie and Masahito Yamazaki for collaboration at an early stage of this project.

All the authors would like to thank the Simons Center for Geometry and Physics for hospitality during the completion of this work, both during June 2012 and during the 2012 Summer Workshop in Mathematics and Physics. TD is also grateful to the Max Planck Institute for Mathematics, Bonn, for hospitality during September, 2012; and MG would like to thank the Institute for Advanced Studies for hospitality. The work of TD is supported by a William D. Loughlin Membership at the Institute for Advanced Study, with additional support from DOE grant DE-FG02-90ER-40542. The work of MG is supported by the ERC Starting Independent Researcher Grant 259133. The work of AG was supported by NSF grants DMS-1059129 and DMS-1301776.

A Symplectic data and class \mathcal{R}

Section 5 of this paper was devoted to describing how the K -triangulation of 3-manifolds defines symplectic gluing data $(g, \sigma) \in \text{ISp}(2N, \mathbb{Z})$ or $(g, \sigma) \in \text{ISp}(2N, \mathbb{Q})$ with certain properties, if one assumes the Symplectic Gluing Conjecture 4.1. In this appendix we review how such abstract data (g, σ) — whether it arises from a 3-manifold or not — subsequently defines 3d SCFT's $T_{(g, \sigma)}$ in class \mathcal{R} and various associated invariants. We largely follow [1].

Of course, multiple pairs (g, σ) can lead to the same SCFT $T_{(g, \sigma)}$. Thus it will be natural to define a set \mathcal{R} of equivalence classes of symplectic data whose elements are in one-to-one correspondence with 3d theories. Conjecturally, all nontrivial equivalences $(g, \sigma) \sim (g', \sigma')$ are generated by the algebraic analogues of 2–3 moves and octahedron rotations. (Physically, we would like to say that all non-trivial mirror symmetries of 3d abelian Chern-Simons matter theories are generated by the basic XYZ \leftrightarrow SQED duality.)

A.1 The symplectic data

To be more concrete, we recall that elements of the affine symplectic group³⁰ $\text{ISp}(2N, \mathbb{Z}) \simeq \text{Sp}(2N, \mathbb{Z}) \ltimes \mathbb{Z}^{2N}$ can be represented as pairs (g, σ) , where $g \in \text{Sp}(2N, \mathbb{Z})$ and σ is a vector of $2N$ integers. We define a *decoration* for (g, σ) to be an integer $0 \leq d \leq N$, and distinguish the first d rows of g (and the first d elements of σ). We then let $\widehat{\mathcal{R}}$ be the set of such decorated elements (g, σ) , for all $d, N \geq 0$.

We further represent an element $(g, \sigma) \in \widehat{\mathcal{R}}$ as a transformation on a complex symplectic phase space $\mathcal{P}_\times = \mathbb{C}^{2N}$:

$$\begin{pmatrix} X_1 \\ \vdots \\ X_d \\ C_1 \\ \vdots \\ C_{N-d} \\ P_1 \\ \vdots \\ P_d \\ \Gamma_1 \\ \vdots \\ \Gamma_{N-d} \end{pmatrix} = \underbrace{\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix}}_g \cdot \begin{pmatrix} Z_1 \\ \vdots \\ Z_N \\ Z'_1 \\ \vdots \\ Z''_N \end{pmatrix} - i\pi \underbrace{\begin{pmatrix} \sigma_X \\ \sigma_P \end{pmatrix}}_\sigma, \quad (\text{A.1})$$

where $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$ are $N \times N$ blocks and $\sigma_X = (\sigma_1, \dots, \sigma_N)^T$ and $\sigma_P = (\sigma_{N+1}, \dots, \sigma_{2N})^T$ are N -vectors of integers. The transformation preserves the complex symplectic form

$$\Omega_\times = \sum_{i=1}^N dZ''_i \wedge dZ_i = \sum_{j=1}^d dP_j \wedge dX_j + \sum_{k=1}^{N-d} d\Gamma_k \wedge dC_k, \quad (\text{A.2})$$

and can be thought of as a change of polarization. It sends $\Pi_\times = (Z_i; Z''_i)$ (with Z_i as positions and Z''_i as conjugate momenta) to $\tilde{\Pi} = (X_j, C_k; P_j, \Gamma_k)$, where X_j, C_k are positions and P_j, Γ_k are conjugate momenta. The positions X_j label the d decorated rows.

When (g, σ) comes from a K -triangulation of a 3-manifold M , as described in section 5, the X_j and P_j are coordinates for a boundary phase space $\mathcal{P}_K(\partial M)$; while the C_k are the gluing functions.

³⁰Here we will focus on $\text{ISp}(2N, \mathbb{Z})$ rather than $\text{ISp}(2N, \mathbb{Q})$. The generalization to rational transformations is straightforward, but requires extra technical care.

A.2 Constraints on the data

We impose two constraints on the elements of $\widehat{\mathcal{R}}$, which seem to be necessary for defining 3d theories and other invariants.

The first constraint is satisfied by every decorated (g, σ) that arises from geometry. We call it the *Angle Constraint*. To describe it most symmetrically, we introduce N auxiliary coordinates

$$Z'_i = i\pi - Z_i - Z''_i \tag{A.3}$$

on the phase space \mathcal{P}_\times . The Angle Constraint requires that every undecorated coordinate C_k can be written in the form

$$C_k = \sum_{i=1}^N (\alpha_{ki} Z_i + \alpha'_{ki} Z'_i + \alpha''_{ki} Z''_i) - 2\pi i, \tag{A.4}$$

where $\alpha_{ki}, \alpha'_{ki}, \alpha''_{ki}$ (with $1 \leq k \leq N-d$ and $1 \leq i \leq N$) are three matrices of non-negative integers that satisfy

$$\sum_{j=1}^{N-d} \alpha_{kj} \leq 2, \quad \sum_{j=1}^{N-d} \alpha'_{kj} \leq 2, \quad \sum_{j=1}^{N-d} \alpha''_{kj} \leq 2 \quad \forall 1 \leq k \leq N-d. \tag{A.5}$$

(It is easy to see that geometric gluing functions C_k all satisfy the Angle Constraint, since they are positive sums of octahedron parameters, minus $2\pi i$, and every octahedron parameter occurs at most twice.) The Angle Constraint implies that the undecorated entries of the blocks \mathbf{A} and \mathbf{B} in (A.1) obey $|\mathbf{A}_{d+k,i}| \leq 2$, $|\mathbf{B}_{d+k,i}| \leq 2$, and $|\mathbf{A}_{d+k,i} - \mathbf{B}_{d+k,i}| \leq 2$.

In order to define gauge theories, we need to supplement the Angle Constraint with a stronger *Superpotential Constraint*. It requires that the undecorated coordinates can be written in the form (A.4) where for each i and k only one of $\alpha_{ki}, \alpha'_{ki}, \alpha''_{ki}$ is nonzero. Alternatively, we could say that for any fixed k there exist suitable cyclic rotations of the initial coordinates $Z_i \rightarrow Z'_i \rightarrow Z''_i \rightarrow Z_i$ (independent for each i) so that after these rotations

$$C_k = \sum_{i=1}^N \alpha_{ki} Z_i - 2\pi i. \tag{A.6}$$

This will ultimately allow a construction of chiral operators \mathcal{O}_k associated with the C_k .

A.3 The gauge theory

Given an element $(g, \sigma) \in \widehat{\mathcal{R}}$ that satisfies the Angle and Superpotential constraints, it is possible to define a 3d $\mathcal{N} = 2$ gauge theory $T_{(g,\sigma)}$ in the ultraviolet — which then flows to a desired SCFT in the infrared. The theory preserves a $U(1)_R$ R-symmetry, and the data (g, σ) encodes a specific UV R-charge assignment, which can then be used to compactify $T_{(g,\sigma)}$ on curved manifolds. Moreover, the data (g, σ) encodes a specific choice of Chern-

Simons couplings for the background gauge fields associated with U(1) flavor symmetries and with the R-symmetry.³¹

The gauge theory $T_{(g,\sigma)}$ is built in the following steps.

0. Define $T_{\diamond} \equiv T_2[\diamond, \Pi_Z]$ to be the octahedron theory of section 6.1 (*a.k.a.* the $K = 2$ tetrahedron theory). It is the theory of a free $\mathcal{N} = 2$ chiral multiplet ϕ , with charge +1 under a background U(1) flavor symmetry and R-charge zero, and an additional level $-1/2$ Chern-Simons coupling for difference of flavor and R-charge gauge fields. The Chern-Simons level matrix can be represented as

$$\begin{array}{|c|cc|}
 & F & R \\
 \hline
 F & -1/2 & 1/2 \\
 R & 1/2 & * \\
 \hline
 \end{array} . \tag{A.7}$$

1. Form the product theory $T_{\times} = T_{\Delta_1} \otimes \dots \otimes T_{\Delta_N}$. It is a theory of N chirals ϕ_i , and has maximal abelian flavor symmetry $U(1)^N$. The real masses of the U(1)'s (scalars in the background gauge multiplets) can be associated with the real parts of the coordinates Z_i on \mathcal{P}_{\times} .
2. Apply the $Sp(2N, \mathbb{Z})$ symplectic transformation g to the theory T_{\times} to obtain a theory \tilde{T} with a new $U(1)^N$ flavor symmetry. The $Sp(2N, \mathbb{Z})$ action on 3d $\mathcal{N} = 2$ gauge theories was defined in [92], and can be implemented (e.g.) by factoring g into generators. Then
 - a. “U-type” generators of the form $\begin{pmatrix} \mathbf{U} & 0 \\ 0 & \mathbf{U}^{-1T} \end{pmatrix}$, where $\mathbf{U} \in GL(N, \mathbb{Z})$, act by changing the basis of $U(1)^N$ flavor currents. If $V = (V_1, \dots, V_N)^T$ is a vector of background gauge multiplets corresponding to the $U(1)^N$ symmetry, we send $V \mapsto \mathbf{U}^{-1}V$.
 - b. “T-type” generators of the form $\begin{pmatrix} \mathbf{I} & 0 \\ \mathbf{k} & \mathbf{I} \end{pmatrix}$, where \mathbf{I} is the $N \times N$ identity and \mathbf{k} is $N \times N$ symmetric, act by adding supersymmetric Chern-Simons terms for the background $U(1)^N$ flavor symmetry with (mixed) level matrix \mathbf{k} .
 - c. “S-type” generators of the form $\begin{pmatrix} \mathbf{I} - \mathbf{E}_i & -\mathbf{E}_i \\ \mathbf{E}_i & \mathbf{I} - \mathbf{E}_i \end{pmatrix}$, where E_i is an $N \times N$ matrix with entry 1 at the i -th spot on the diagonal and zeroes everywhere else, act by gauging the i -th $U(1)_i$ symmetry and replacing it with a new topological $U(1)_J$ flavor symmetry. A background gauge multiplet for $U(1)_J$ is coupled to the gauge multiplet of the gauged $U(1)_i$ via a supersymmetric FI term, which is the same thing as a mixed Chern-Simons term with level matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$.

This $Sp(2N, \mathbb{Z})$ action can be extended to rational matrices $Sp(2N, \mathbb{Q})$ if one allows rescaling of the charge lattice of T_{\times} .

³¹More precisely, (g, σ) encodes flavor-flavor and mixed flavor-R Chern-Simons couplings. It does not uniquely specify R-R couplings. Such background couplings and their effect on partition functions and dualities were recently discussed in [86, 87].

- Now the theory \tilde{T} from Step 2 has a $U(1)^N$ flavor symmetry with real masses associated with the top N rows of g , i.e. to the coordinates X_j and C_k .

Change the R-charge assignment of the theory \tilde{T} by adding σ_i ($1 \leq i \leq N$) multiples of each $U(1)_i$ flavor charge to the R-charge.

For each component σ_{N+i} in the bottom half of the shift vector, $1 \leq i \leq N$, add $-\sigma_{N+i}$ units of mixed Chern-Simons coupling between the background gauge field for the $U(1)_i$ flavor symmetry and the background gauge field for the R-symmetry.

This interpretation of affine shifts defines an action of the full affine symplectic group $ISp(2N, \mathbb{Z})$ on 3d SCFT's. In particular, the full relations of $ISp(2N, \mathbb{Z})$ are satisfied (after flowing to the IR), up to an integer ambiguity in the background R-R Chern-Simons level.

- Finally, break $N - d$ flavor symmetries $U(1)_k$ associated with the undecorated coordinates C_k . This is done by adding $N - d$ operators \mathcal{O}_k to the superpotential, charged under each respective $U(1)_k$. We'll review the construction of the operators below. The breaking can be thought of as setting the respective mass parameters $C_k \rightarrow 0$.

In the end, this produces a theory $T_{(g,\sigma)}$ with manifest $U(1)^d$ flavor symmetry. The mass parameters of the d unbroken $U(1)$'s are associated with the coordinates X_j (i.e. with the decorated rows of (g, σ)).

The rank of the gauge group of $T_{(g,\sigma)}$ depends on exactly how the symplectic transformation g is applied in Step 2. Different decompositions into generators lead to slightly different UV descriptions that flow to the same thing in the IR. (In the IR the relations of the symplectic group hold.) It is fairly easy to see that a lower bound on the rank of the UV gauge group is given by $\text{rank}(\mathbf{B})$. Often the bound can be realized — i.e. there exists a decomposition of g that involves exactly $\text{rank}(\mathbf{B})$ S -type generators.

Alternatively, if \mathbf{B} has maximal rank and is unimodular, there exists a way to apply the affine-symplectic transformations of Steps 0–3 all at once, cf. [110]. We define a theory \tilde{T} by starting with N chirals ϕ_i , each charged under an independent dynamical gauge symmetry $U(1)_i$ (and having zero R-charge). Thus the gauge group is $U(1)^N$. The flavor group is also $U(1)^N$, since each dynamical $U(1)_i$ has an associated withtopological $U(1)_{J,i}$. We then specify the full Chern-Simons coupling matrix to be

	G	J	R
G	$\mathbf{B}^{-1}\mathbf{A} - \frac{1}{2}\mathbf{I}_{N \times N}$	$-\mathbf{B}^{-1}$	$\frac{1}{2} - \mathbf{B}^{-1}\sigma_X$
J	$-\mathbf{B}^{-1T}$	$\mathbf{D}\mathbf{B}^{-1}$	$\mathbf{D}\mathbf{B}^{-1}\sigma_X - \sigma_P$
R	$\frac{1}{2} - (\mathbf{B}^{-1}\sigma_X)^T$	$(\mathbf{D}\mathbf{B}^{-1}\sigma_X - \sigma_P)^T$	*

(A.8)

(Here ‘ G ’ stands for the N dynamical gauge fields, ‘ J ’ for the N background flavor gauge fields, and ‘ R ’ for the background $U(1)_R$ field. For instance, the G – J blocks encode FI terms. The G – G blocks encodes dynamical CS terms that couple the otherwise independent gauge fields.) The N flavor symmetries of \tilde{T} correspond directly to the coordinates (X_j, C_k) .

To obtain $T_{(g,\sigma)}$ from \widetilde{T} we “simply” have to apply Step 4, adding appropriate operators to the superpotential to break the C_k symmetries.

A.4 Equivalences

Two elements $(g, \sigma), (g', \sigma') \in \widehat{\mathcal{TR}}$ can sometimes define theories that are equivalent after flowing to the IR, $T_{(g,\sigma)} \simeq T_{(g',\sigma')}$. There are three basic ways in which this can happen. (In the following, we will *assume* that both $T_{(g,\sigma)}$ and $T_{(g',\sigma')}$ possess the chiral operators needed to break flavor symmetries in Step 4 of the gauge theory construction; we return to this in section A.5.)

First, there are some trivial equivalences. We’ll describe them in terms of initial and final coordinates on a phase space \mathcal{P}_\times , as in (A.1). We can permute pairs of initial coordinates $(Z_i, Z''_i) \leftrightarrow (Z_{i'}, Z''_{i'})$ (which corresponds to permuting the N chirals of $T_{(g,\sigma)}$). Similarly, we can permute the final coordinates C_k . Moreover, we can redefine the X_j and P_j coordinates by any integer linear combination of the C_k , as long as we simultaneously redefine the Γ_k so that g remains symplectic. We can also redefine the Γ_k by integer linear combinations of the C_k and any multiples of $i\pi$. In gauge theory, these latter redefinitions of coordinates all correspond to redefining the flavor currents (or background CS levels) for flavor symmetries $U(1)_k$ that will ultimately be broken.

We might also allow ourselves redefine the P_j by integer linear combinations of the X_j and by multiples of $i\pi$, while keeping g symplectic. This changes background Chern-Simons levels for the flavor and R-symmetries of $T_{(g,\sigma)}$, which one is sometimes interested in keeping track of and sometimes not.

The second type of equivalence is an “octahedron rotation.” In terms of phase-space coordinates, it corresponds to a cyclic permutation $Z_i \rightarrow Z'_i \rightarrow Z''_i \rightarrow Z_i$ for any fixed i ; in other words,

$$\begin{pmatrix} Z_i \\ Z''_i \end{pmatrix} \mapsto \begin{pmatrix} i\pi - Z_i - Z''_i \\ Z_i \end{pmatrix}. \tag{A.9}$$

This is an operation on two columns of g , and on σ . In terms of gauge theory, the cyclic rotation corresponds to taking a chiral ϕ_i of $T_{(g,\sigma)}$ and replacing it by a $U(1)$ gauge theory coupled to another chiral ϕ'_i to obtain $T_{(g',\sigma')}$, as discussed in detail in section 6.1. This should be an IR duality.³²

The third type of equivalence is the algebraic version of an elementary “2–3 move” on octahedra. It replaces (g, σ) of rank N and d decorated rows with (g', σ) of rank $N' = N + 1$ and $d' = d$. To implement it in the $2 \rightarrow 3$ direction, we select any two pairs of initial coordinates, say (Z_1, Z''_1) and (Z_2, Z''_2) , and replace them with new initial coordinates $(W_i, W''_i)_{i=1}^3$, where

$$Z_1 = W''_2 + W'_3, \quad Z''_1 = W''_1 + W'_2, \quad Z_2 = W'_2 + W''_3, \quad Z''_2 = W'_3 + W''_1, \tag{A.10}$$

³²As usual, this statement must be taken with a grain of salt. We are assuming that the mirror symmetry between a free chiral T_{\diamond, Π_Z} and a free “vortex” $T_{\diamond, \Pi_{Z'}}$, as in (6.2)–(6.3), implies IR duality between any gauge theories $T_{(g,\sigma)}, T_{(g',\sigma')}$ in which the free chiral and the free vortex (respectively) are embedded. Arguing this carefully may require a subtle interchange of limits of IR flow. The same comments apply to the 2–3 moves below.

as in section 5.5.1, with $W'_i = i\pi - W_i - W''_i$. We also add a new undecorated pair of rows $(C_0, \Gamma_0) = (W_1 + W_2 + W_3 - 2\pi i, W''_1)$. The new (g', σ') can be worked out in a straightforward manner. In the opposite direction, a $3 \rightarrow 2$ move can only be applied to an element (g', σ') that is already the image of a $2 \rightarrow 3$ move. In particular, an undecorated row of (g', σ') must already define a coordinate of the form $C_0 = W_1 + W_2 + W_3 - 2\pi i$.

The interpretation of the 2–3 move in gauge theory was reviewed in section 6.2. To apply a $2 \rightarrow 3$ move, we take any pair of chirals ϕ_1, ϕ_2 , think of them as an embedded hypermultiplet theory (6.7), and replace the hypermultiplet theory with a copy of $\mathcal{N} = 4$ SQED (6.9). In the opposite direction, we isolate any three chirals coupled by a cubic superpotential, treat them as a copy of the XYZ model (6.13), and replace them with a copy of $\mathcal{N} = 2$ SQED (6.12). We must be careful, though: these applications of a 2–3 move do *not* necessarily preserve all of the operators \mathcal{O}_k associated with additional broken symmetries. In order to hope that $T_{(g,\sigma)}$ and its putative image $T_{(g',\sigma')}$ are truly IR dual, both of them must independently contain all the necessary \mathcal{O}_k 's.

Algebraically, we note that the Angle Constraint (A.4) is preserved by all three types equivalences we have just described. In contrast, the Superpotential Constraint is preserved by the first and second types, but not by all 2–3 moves. A preliminary guess for a set that might be in one-to-one correspondence with duality classes of 3d $\mathcal{N} = 2$ SCFT's (that flow from abelian theories in the UV) is

$$\mathcal{R} := \frac{\{(g, \sigma) \in \widehat{\mathcal{R}} \mid (g, \sigma) \text{ satisfy the Angle and Superpotential Constraints}\}}{\text{modulo equivalences}} \quad (\text{A.11})$$

IR ?

$$\mathcal{R} := \frac{\{3\text{d } \mathcal{N} = 2 \text{ abelian Chern-Simons-matter theories with any } U(1)_R\text{-preserving superpotentials}\}}{\text{modulo IR duality}}.$$

In defining \mathcal{R} , we only quotient out by equivalences (including 2–3 moves) that do preserve the Superpotential Constraint. The potential correspondence $\mathcal{R} \simeq \mathcal{R}$ deserves much further study.

A.5 The existence of operators \mathcal{O}_k

For completeness, let us briefly describe the construction of the operators \mathcal{O}_k . We assume that we have followed Steps 0–3 of the gauge theory construction to build a theory \widetilde{T} , associated with symplectic data (g, σ) , and we want to apply Step 4.

If (g, σ) satisfies the Angle and Superpotential Constraints, then for every symmetry $U(1)_k$ that must be broken there is another pair $(g, \sigma)_k$, related to (g, σ) by “octahedron rotations,” such that the corresponding coordinate C_k takes the form (A.6). The pair $(g, \sigma)_k$ defines a theory \widetilde{T}_k after Step 3 that is mirror symmetric to \widetilde{T} . In \widetilde{T}_k one defines an operator

$$\widehat{\mathcal{O}}_k = \phi_1^{\alpha_{k1}} \phi_2^{\alpha_{k2}} \dots \phi_N^{\alpha_{kN}}, \quad (\text{A.12})$$

a simple product of elementary chiral fields of the theory. (Recall that the α_{ki} must be nonnegative integers!) Because of the affine term $-2\pi i$ in (A.6), this operator is guaranteed

to have UV R-charge 2. By using the mirror symmetry between \widetilde{T} and \widetilde{T}_k , one can then “pull back” the operator $\widehat{\mathcal{O}}_k$ to an equivalent chiral operator \mathcal{O}_k in \widetilde{T} , charged precisely under the desired $U(1)_k$ symmetry. This can be repeated for each of the $N - d$ symmetries that must be broken. Then the final superpotential of $T_{g,\sigma}$ takes the form $\sum_{k=1}^{N-d} \mathcal{O}_k$.

A.6 Associated invariants

Finally, let us touch upon some mathematical objects associated with an element $(g, \sigma) \in \widehat{\mathcal{R}}$, and their physical significance.

All of these objects are invariant under the first two types of equivalences presented in section A.4. However, further restrictions are sometimes necessary to ensure invariance under 2–3 moves, such as the non-degeneracy discussed in section 5.5.1.

Phase space and K_2 -Lagrangian. We explained how to build phase spaces and classical Lagrangian submanifolds (and their geometric significance) in section 5. The phase space is a symplectic reduction

$$\mathcal{P}_{(g,\sigma)} = \mathcal{P}_\times // (C_k = 0) \simeq \{(X_j, P_j)_{j=1}^d\}, \tag{A.13}$$

and has the holomorphic symplectic structure

$$\Omega = \sum_{j=1}^d dP_j \wedge dX_j, \tag{A.14}$$

whereas the putative Lagrangian submanifold $\mathcal{L}_{(g,\sigma)} \subset \mathcal{P}_{(g,\sigma)}$ is the image of the product $\mathcal{L}_\times = \{z_i'' + z_i^{-1} - 1 = 0\}_{i=1}^N$ (with $z_i = \exp Z_i, z_i'' = \exp Z_i''$) under the reduction. Explicitly, this means inverting the transformation (A.1) to re-write z_i, z_i'' as Laurent monomials in the final coordinates $x_j = \exp X_j, c_k = \exp C_k, p_j = \exp P_j, \gamma_k = \exp \Gamma_k$; then eliminating γ_k from the equations and setting $c_j = 0$. The result is a Lagrangian submanifold in the exponentiated phase space (with \mathbb{C}^* coordinates x_j and p_j), so long as \mathcal{L}_\times is transverse to the moment maps, and to slices of $\mathcal{P}_{(g,\sigma)}$ at generic constant x_j .

The Lagrangian submanifold is the supersymmetric parameter space for the theory $T_{(g,\sigma)}$ on $\mathbb{C} \times S^1$ [1, 34, 88]. The effective $\mathcal{N} = (2, 2)$ theory obtained from compactifying on S^1 at finite radius is governed by a twisted superpotential $\widetilde{W}(X_1, \dots, X_d; \sigma)$. It is the function of complex twisted masses that correspond to the decorated position coordinates of $\mathcal{P}_{(g,\sigma)}$, as well as dynamical twisted-chiral fields σ . The Lagrangian equations are

$$\mathcal{L}_{(g,\sigma)} \simeq \mathcal{L}_{\text{SUSY}} : \quad \exp \left(\left. \frac{\partial \widetilde{W}}{\partial Z_i} \right|_{\partial \widetilde{W} / \partial \sigma = 0} \right) = p_i. \tag{A.15}$$

We saw in section 5.5.1 that $\mathcal{L}_{(g,\sigma)}$ is invariant under elementary 2–3 moves so long as the transversality condition is preserved. (In particular, it was necessary that at generic values of x_j , the gluing conditions $c_k = 1$ never forced z_i, z_i'' to take values 0, 1, or ∞ .) In terms of gauge theory, transversality means that $T_{(g,\sigma)}$ has isolated massive vacua on $\mathbb{C} \times S^1$. The vacua correspond to the points $p^{(\alpha)}(x)$ on $\mathcal{L}_{(g,\sigma)}$ at fixed mass parameters x .

Volume. Suppose that the product Lagrangian \mathcal{L}_\times is transverse to the gluing constraints and to slices at constant x , as above. Then by using the equations $z_i'' + z_i^{-1} - 1 = 0$ and the exponentiated form of (A.1) we can express z_i and z_i'' as rational functions on $\mathcal{L}_{(g,\sigma)}$ — i.e. rational functions of x, p . (Geometrically, we would be solving for octahedron parameters given fixed boundary coordinates on ∂M .) The real volume function associated with (g, σ) can be defined as

$$\text{Vol}_{(g,\sigma)}(x, p) = \sum_{i=1}^N \mathcal{L}i_2(z_i), \tag{A.16}$$

where $\mathcal{L}i_2(z) = \text{Im Li}_2(z) + \arg(1 - z) \log |z|$ is the Bloch-Wigner dilogarithm function. The volume is a function on $\mathcal{L}_{(g,\sigma)}$, and is invariant under algebraic 2–3 moves (that preserve transversality) due to the 5-term relation for $\mathcal{L}i_2(z)$.

The data $(g, \sigma) \in \text{ISp}(2N, \mathbb{Z})$ actually allows us to promote (A.16) to a complex volume defined modulo $\frac{\pi^2}{6}\mathbb{Z}$ using methods of [70] (or as in [110, Sec 5.2]).

For the theory $T_{(g,\sigma)}$, the volume $\text{Vol}_{(g,\sigma)}(x, p^{(\alpha)}(x))$ is the free energy density (the value of \widetilde{W}) in the vacuum $p^{(\alpha)}(x)$ on $\mathbb{C} \times S^1$. As we already observed in section 7.2.2, the volumes of *particular* solutions $p^{(\alpha)}(x)$ dominate the asymptotics of various partition functions of $T_{(g,\sigma)}$.

Quantum Lagrangian. We reviewed in section 5.4 how to use the data (g, σ) to promote a classical Lagrangian $\mathcal{L}_{(g,\sigma)}$ to a left ideal of quantum operators $\hat{\mathcal{L}}_{(g,\sigma)}$. The basic idea is to promote all logarithmic coordinates to operators satisfying $[\hat{Z}_i'', \hat{Z}_i] = \hbar$, etc., and to promote (A.1) to linear relations among quantum operators with a canonical correction $i\pi\sigma \rightarrow (i\pi + \hbar/2)\sigma$ on the r.h.s. . This then allows the left ideal $(\hat{z}_i'' + \hat{z}_i^{-1} - 1)_{i=1}^N$ to be rewritten in terms of exponentiated operators $\hat{x} = \exp \hat{X}$, $\hat{p} = \exp \hat{P}$, $\hat{c} = \exp \hat{C}$, $\hat{\gamma} = \exp \hat{\Gamma}$. Eliminating $\hat{\gamma}$ and setting $\hat{c} = 1$ should produce $\hat{\mathcal{L}}_{(g,\sigma)}$.³³ The left ideal $\hat{\mathcal{L}}_{(g,\sigma)}$ encodes the Ward identities for line operators satisfied by partition functions of $T_{(g,\sigma)}$.

Quantum partition functions. The element (g, σ) gives rise to several closely related wavefunctions. Geometrically, they are versions of Chern-Simons partition functions on a 3-manifold M ; they are also partition functions of $T_{(g,\sigma)}$ on simple curved backgrounds.

- The supersymmetric index of $T_{(g,\sigma)}$ on $S^2 \times S^1$ was described in [54] following [111–113]. As a wavefunction, it corresponds to quantizing the exponentiated phase space $\mathcal{P}_{(g,\sigma)} \simeq (\mathbb{C}^*)^{2d}$ with a *real* symplectic form $\omega \sim \hbar^{-1} \text{Im } \Omega$. In a particular real polarization, it depends on d integers m_j (magnetic fluxes) and d phases $\zeta_j \in S^1$ (electric fugacities) as well as a formal parameter $q = e^\hbar$. The combinations $q^{m_j/2} \zeta_j$ are identified with the position coordinates x_j . Geometrically, the index matches the partition function of $\text{PGL}(K, \mathbb{C})$ Chern-Simons theory at level $k = 1$ on a manifold M . The index is well-defined and invariant under 2–3 moves so long as the data (g, σ) admits a so-called semi-strict angle structure [61].

³³An alternative method of quantizing one-dimensional Lagrangians involves the topological recursion of [57], adapted to the current setting in [20, 58]. It is expected that the topological recursion is equivalent to the algebraic approach outlined here.

- The ellipsoid (S_b^3) partition function of $T_{(g,\sigma)}$ can be computed by methods of [1, 104, 105]. It corresponds to quantization of $\mathcal{P}_{(g,\sigma)}$ with respect to a real symplectic form $\text{Re}(\Omega) + s \text{Im}(\Omega)$, where $s = i \frac{1-b^2}{1+b^2}$.³⁴ As explained in [55] (after the first version of the current paper appeared), this turns out to be identical to the quantization of the real slice of the phase space $\mathcal{P}_{(g,\sigma)}$ with $X_j, P_j \in \mathbb{R}$. Thus, geometrically, one may identify the ellipsoid partition function with either a $\text{PGL}(K, \mathbb{C})$ Chern-Simons partition function at level $k = 1$ or a $\text{PGL}(K, \mathbb{R})$ Chern-Simons partition function. It depends analytically on d continuous variables X_j (masses). Geometric partition functions of this type (“state integral models”) first appeared in [10], with later developments including [9, 11, 59, 60].
- The partition function of a 3d $\mathcal{N} = 2$ theory on an ellipsoidally deformed lens space $L(k, 1)_b = S_b^3/\mathbb{Z}_k$ was developed and studied in [114–116]. Geometrically, it is expected to agree with a $\text{PGL}(K, \mathbb{C})$ Chern-Simons partition function at level k [55, 65].
- By expanding the formal integrals for ellipsoid partition functions around complex critical points, it is possible to derive a series of perturbative invariants associated with classical solutions $p^{(\alpha)}(x)$ on $\mathcal{L}_{(g,\sigma)}$ (i.e., geometrically, to flat $\text{PGL}(K, \mathbb{C})$ -connections with fixed boundary conditions x) [110]. After the volume (A.16), the first subleading invariant is a Reidemeister-Ray-Singer torsion. In the geometric setting, these perturbative invariants are expected to match the asymptotic expansion of colored HOMFLY polynomials on M .
- The partition functions $B^\alpha(x; q)$ of $T_{(g,\sigma)}$ on a twisted (spinning) geometry $\mathbb{C} \times_q S^1$ were recently discussed in [117]. These “holomorphic blocks” depend locally on a solution $p^{(\alpha)}(x)$ to the classical Lagrangian equations, but are fully nonperturbative quantum objects. It was conjectured that both the index and the ellipsoid partition functions can be written as sums of products of the same holomorphic blocks. Geometrically, the holomorphic blocks seem to come from nonperturbative completions of analytically continued Chern-Simons theory on a manifold M , as in [94, 109].

B The Poisson bracket for eigenvalues

Here we prove that the eigenvalues of holonomies around noncontractible cycles on small boundaries (small tori or annuli) have the expected Atiyah-Bott-Goldman Poisson brackets discussed in sections 3.4 and 5.3.

We consider an admissible 3-manifold M with an ideal triangulation and a subsequent K -decomposition. Let $\mathcal{P}_\times = \mathcal{P}_{\partial\mathcal{O}_i}$ be the linear product phase space of all the octahedra. We work with logarithmic coordinates, as in section 5. For every internal and external point in the K -decomposition we define affine-linear functions C_k, \mathcal{X}_j (respectively) on \mathcal{P}_\times — the sums of octahedron parameters at vertices that touch those points. For any closed path γ on a small boundary, we choose $K - 1$ representative paths γ_a on the slices parallel

³⁴Strictly speaking, the symplectic form is real when $|b| = 1$. To make sense of other values of b , one must analytically continue the answer.

to that small boundary; then we use the rules of section 5.2.1 to define $K - 1$ corresponding affine-linear path-coordinates U_a^γ .

We already argued in section 3.2.3 (using cancellation of arrows on octahedra) that

$$\{C_k, C_{k'}\} = \{C_k, \mathcal{X}_j\} = 0, \tag{B.1}$$

and

$$\{\mathcal{X}_j, \mathcal{X}_{j'}\} = \text{signed sum of boundary arrows connecting points at } \mathcal{X}_j \text{ and } \mathcal{X}_{j'}, \tag{B.2}$$

which is the expected Poisson structure on the boundary. We now want to show that

$$\{C_k, U_a^\gamma\} = \{\mathcal{X}_j, U_a^\gamma\} = 0, \tag{B.3}$$

and

$$\{U_a^\gamma, U_b^\lambda\} = \kappa_{ab} \langle \gamma, \lambda \rangle, \tag{B.4}$$

where $\langle \gamma, \lambda \rangle$ is the signed intersection number of the paths γ, λ on the small boundary, and

$$\kappa_{ab} = \begin{cases} 2 & a = b \\ -1 & |a - b| = 1 \\ 0 & |a - b| \geq 2 \end{cases} \tag{B.5}$$

is the Cartan matrix of $SL(K)$. This justifies all of the brackets summarized in section 5.3, and generalizes a central result of [8] for $K = 2$.

Our proof basically extends that of [8], using the framework of slices and path coordinates. There are two basic steps. First we show that the U_a^γ commute with all other coordinates. The fact that U_a^γ commutes with the C_k immediately implies that the Poisson bracket of U_a^γ only depends on the *homotopy class* of the chosen path γ_a on the a -th slice. In particular, it indicates that $\{U_a^\gamma, U_b^\lambda\}$ should only depend on the intersection number $\langle \gamma, \lambda \rangle$ and some universal function of a, b . The second part of the proof fixes the normalization of $\{U_a^\gamma, U_b^\lambda\}$, through what is unfortunately a brute-force, case-by-case computation.

B.1 $\{C, U\} = \{\mathcal{X}, U\} = 0$

Let p be an internal or external point in the K -decomposition of M . Let \mathcal{X} generically denote the sum of octahedron parameters at p . Let U be the coordinate for any closed path γ (we drop the subscript a in this subsection) on any global slice S in the K -decomposition. We want to show that $\{\mathcal{X}, U\} = 0$.

The approach will be to look at how the octahedron parameters contributing to \mathcal{X} can appear on the slice S . There are two basic cases: either \mathcal{X} contains a sum of parameters in a *crown-shaped* or *wheel-shaped* arrangement of small triangles on S . These two arrangements are shown in figure 66, with the parameters contributing to \mathcal{X} marked by red dots. Any path disjoint from these arrangements automatically commutes with \mathcal{X} . Any path γ that enters and exits these arrangements (without stopping) also has a coordinate U that commutes with the sum of contributions to \mathcal{X} , due to pairwise cancellations in $\{\mathcal{X}, U\}$. For example,




Figure 66. The wheel and crown arrangements that can occur on slices.




Figure 67. Slices in the neighborhood of a point p (blue dot), with octahedron parameters appearing in \mathcal{X} marked by red dots.

in figure 66, \mathcal{X} has contributions $\mathcal{X} = \sum_{i=1}^5 Z_i + \dots$. The coordinate U for the path drawn through the wheel equals $Z_1 - Z_2 + Z_3 + \dots$, and the coordinate for the path drawn through the crown equals $Z_1'' - Z_4'' + \dots$. In each case there is a pair of ± 1 contributions to $\{\mathcal{X}, U\}$ that cancel, coming from the entry and exit points.

Now let us go through the various options for locations of p more carefully, and verify that the crown and wheel arrangements arise.

First suppose that p is a point internal to the K -decomposition of a tetrahedron. Recall that there are four families of $K - 1$ slices through the octahedron (one family is centered around each tetrahedron vertex). In each family there are three slices that contain parameters contributing to p : the slice containing p and the slices directly above and below it. On the slice containing p , the parameters appear in a six-sided wheel; whereas on the slices above and below the parameters appear in three-spiked crowns (figure 67, left).

If p is on a glued face shared by two tetrahedra, then there are five relevant families of slices — one centered around each vertex of the big bipyramid that contains p . In the three families centered around the vertices of the big glued face, the slices containing p and those above and below p have \mathcal{X} parameters, in wheels or crowns (figure 67, center). In the two families of slices parallel to the glued face, only the $(K - 1)$ -st slice has \mathcal{X} parameters, in a three-spiked crown.




Figure 68. Slices in the neighborhood of an external point p (blue dot).

If p is on an internal edge E of the triangulation shared by N tetrahedra (not necessarily distinct) in an N -gonal bipyramid, there are two types of families of slices to consider. The two families centered around the ends of E each have two slices with \mathcal{X} parameters, in an N -sided wheel and an N -spiked crown (figure 67, right). The N families centered around the “equatorial” vertices of the bipyramid each have the $(K - 1)$ -st slice containing \mathcal{X} parameters, in a 2-spiked crown.

Finally, suppose that p is on a big boundary of M . Then the slices containing \mathcal{X} parameters are basically as described above, except some families of slices may get cut in half. For example, if p is on a face of a tetrahedron on the big boundary, the slices centered around the three vertices of that face look like the left side of figure 68. If p is on an external edge, then the slices centered around the endpoints of the edge look like the right side of figure 68. This does not affect the argument for cancellations above. Any *closed* path γ in one of these slices must both enter and exit the “half-crown” or “half-wheel” regions without stopping, and contributions to $\{\mathcal{X}, U\}$ will continue to cancel in pairs.

B.2 $\{U^\gamma, U^\lambda\} = \kappa \cdot \langle \gamma, \lambda \rangle$

In the second part of the proof, we consider any two closed paths γ, λ on slices in the K -decomposition of M . Again we suppress the subscripts (a, b) denoting the slice number. We want to show that the corresponding path-coordinates U^γ, U^λ have Poisson bracket proportional to the signed intersection number of the projection of the paths to the small boundary; and that the proportionality constant equals 2 if the paths are on the same slice, -1 if the paths are on immediately neighboring slices, and 0 otherwise.

From section B.1 we know that $\{U^\gamma, U^\lambda\}$ only depends on the homotopy class of γ and λ (since closed paths commute with all gluing constraints C_k). Therefore, we may “uniformize” the paths, so that 1) they are smooth (have no bounces); 2) whenever they turn counterclockwise (left) inside a big tetrahedron they stay close to the edge they wind around; and 3) whenever they turn clockwise (right) they stay as far away from the edge as possible (figure 69). (It may be useful to recall that slices are always viewed from *above*, from the perspective of a small boundary, in the 3-manifold M . Thus clockwise/right and counterclockwise/left orientations are well defined.)




Figure 69. Uniformization of paths inside a tetrahedron.

The goal now is to show that $\{U^\gamma, U^\lambda\}$ gets contributions from points of intersection of the paths and nowhere else. We fix γ and assume it lies on the a -th global slice (parallel to some small boundary component). We will look at how octahedron parameters that contribute to U^γ appear in other slices of M , and how they interact with other putative paths λ on some b -th slice. We simply go through an exhaustive case-by-case analysis.

Any putative contribution to $\{U^\gamma, U^\lambda\}$ comes from segments of γ and λ that run along a common octahedron \diamond_i — so that potentially non-commuting parameters Z_i, Z'_i, Z''_i from that octahedron occur in both U^γ and U^λ . Moreover, in order for segments of γ and λ to share an octahedron, the segments must run along slices in a common big tetrahedron Δ . There are two basic cases to consider: these slices are in the same family within Δ (i.e. they are parallel to each other); or they are in different families.

B.2.1 Parallel slices

Suppose that segments of γ and λ run on parallel slices within a single tetrahedron Δ . When exiting Δ they may continue running on parallel slices in a second tetrahedron Δ' , and into a third, etc. As long as the paths stay together, they simultaneously turn left or right within each successive tetrahedron. We will see in a second that the contributions to $\{U^\gamma, U^\lambda\}$ from such synchronous segments (where they turn together) vanish identically. Then, if the paths continue on this way for their entire length, it follows that $\{U^\gamma, U^\lambda\} = 0$, consistent with the fact that $\langle \gamma, \lambda \rangle = 0$. Otherwise, there must exist a tetrahedron Δ_0 in which the paths converge, starting their synchronous run, and a tetrahedron Δ_1 in which the paths diverge, ending the synchronous run. We will show that if γ and λ cross during the run then the contribution to $\{U^\gamma, U^\lambda\} = 0$ equals $\pm \kappa_{ab}$, and otherwise vanishes.

Thus, let us first consider slices of a tetrahedron Δ in which the paths turn the same way, modulo orientation (which may be equal or opposite). The path γ , on the a -th slice, either turns left or right. If it turns left (making a small turn) then the octahedron parameter (call it Z_i) contributing to this segment of γ also appears on the $(a + 1)$ -st slice. This is shown on the left of figure 70 with a red ‘+’ sign, indicating that the parameter contributes positively to U^γ . We also draw possible locations for λ , in green. Depending on the precise value of K and whether λ is oriented equal or opposite to γ , the path can appear in several different positions. It can also appear on any slice. However, in any position that allows λ to pass through small triangles containing Z_i (red +’s), the coordinate U^λ




Figure 70. Paths turning together within a tetrahedron. γ and the octahedron parameters contributing to it are drawn in red, while possible options for λ are in green.

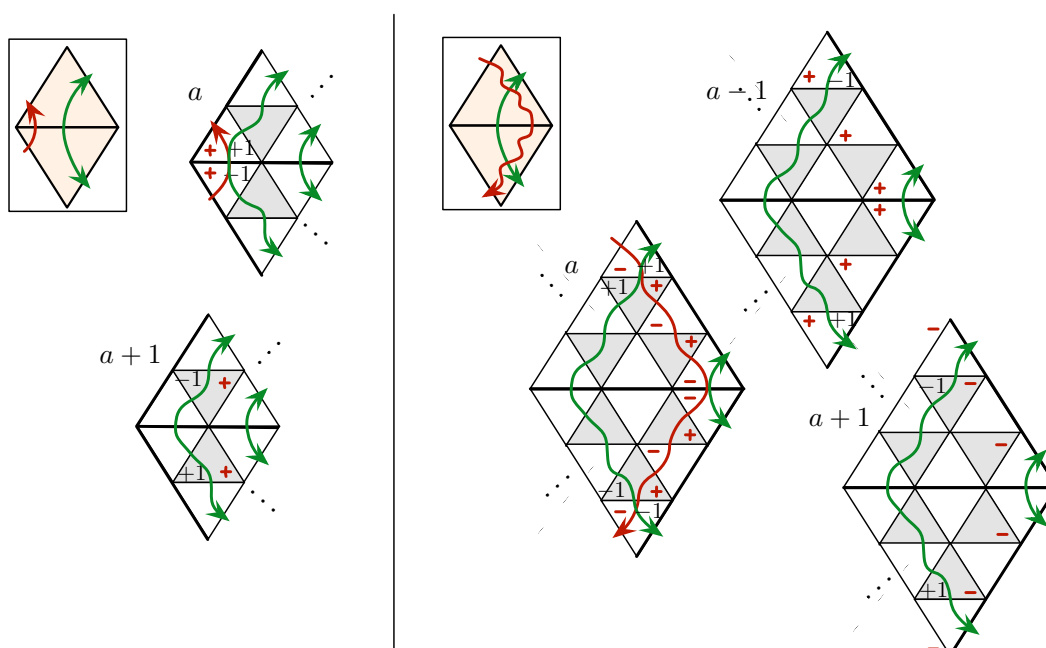


Figure 71. Paths converging and then diverging without a net crossing.

only picks up $\pm Z_i$ — never Z'_i or Z''_i . Thus, in these parallel slices of Δ , the contribution to $\{U^\gamma, U^\lambda\}$ is zero.

Similarly, if γ turns right inside Δ , it picks up a series of (distinct) octahedron parameters with opposite signs, which also appear on the $(a + 1)$ -st and $(a - 1)$ -st slices, shown on the left of figure 70. In any possible configuration, the path for λ at worst picks up these same octahedron parameters. So again the contribution to $\{U^\gamma, U^\lambda\}$ is zero.

Knowing that the intermediate stages of a synchronous run contribute nothing to the Poisson bracket, we can focus on the beginning and end of the run. We might as well assume that the tetrahedra Δ_0, Δ_1 at the beginning and end are immediately adjacent (i.e. there are no intermediate segments).




Figure 72. Paths converging and then diverging with a crossing. $\langle \gamma, \lambda \rangle = \mp 1$ on the left (right).

If the projections of γ and λ to the common small boundary do *not* have a net crossing during the run, there are two possible cases to analyze, depicted in figure 71. On the right, γ turns left (making a small turn), and its octahedron parameters appear on two slices. Call the octahedron parameters Z_0, Z_1 . If λ runs along these slices, it can be in several different positions, depending on K and its orientation. In each possible position, any parameter picked up by U^λ in Δ_0 that does not commute with Z_0 is paired with a parameter in Δ_1 that does not commute with Z_1 , so that the total contribution to $\{Z_0 + Z_1, U^\lambda\}$ cancels. For example, one of the λ paths on the a slice picks up $-Z_0''$ in Δ_0 and $-Z_1'$ in Δ_1 , with $\{Z_0 + Z_1, -Z_0'' - Z_1'\} = 0$.

Similarly, if γ turns to the left, its octahedron parameters appear in slices $a, a + 1$, and $a - 1$. Thus paths λ on these slices might have segments that don't commute with γ . But all non-commuting contributions from Δ_0 cancel with those from Δ_1 .

Altogether, with no net crossing, the total contribution to $\{U^\gamma, U^\lambda\}$ from a full synchronous run of the paths vanishes.

Finally, we consider the most interesting case: a nontrivial crossing of γ, λ going from Δ_0 to Δ_1 . There are two possible arrangements. This time we keep careful track of orientations of both γ and λ . The path γ must have segments turning both right and left, so in both cases its octahedron parameters appear on the $(a \pm 1)$ -st slices as well as the a -th. On the left (resp., right) of figure 72 the intersection number of the projections $\langle \gamma, \lambda \rangle$ equals minus one (resp., plus one).

Correspondingly, on the left of figure 72, we see that if λ is in the a -th slice then each of the two small triangles around the intersection point of γ and λ contributes -1 to $\{U^\gamma, U^\lambda\}$. If λ is on the $(a + 1)$ -st or $(a - 1)$ -st slices then U^λ picks up a single parameter that doesn't commute with part of U^γ , contributing $+1$ to $\{U^\gamma, U^\lambda\}$. On the right of figure 72, the same observations hold, with opposite signs. Altogether, we find that the




Figure 73. Paths entering and exiting a tetrahedron on the same pair of faces.

contribution to the bracket from these crossing segments of path is

$$\{U^\gamma, U^\lambda\} = \kappa_{ab} \langle \gamma, \lambda \rangle, \tag{B.6}$$

when λ lies on the b -th slice.

B.2.2 Slices in different families

To complete the proof, we need to show that whenever γ and λ run through the same tetrahedron (or tetrahedra) but along two *different* families of slices (centered locally around two different tetrahedron vertices) the contribution to the Poisson bracket vanishes.

As for parallel slices, is useful to begin with a single tetrahedron Δ in which the two paths turn the same way — meaning that they enter and exit in the same two tetrahedron faces. This identifies an edge E of Δ that both paths turn around. We depict this situation in figure 73, choosing to place γ on the family of slices centered around the “top” vertex of E , and λ around the “bottom” vertex.

There are two possible cases. If γ (on the a -th slice with respect to the top) turns to the left, then it picks up a single octahedron parameter Z_i . This parameter appears on slices $K - a$ and $K - a + 1$ with respect to the bottom. Any possible paths λ at worst pick up Z_i itself, and so commute with this segment of γ . Alternatively, if γ turns to the right, then its octahedron parameters appear on the $(K - 1)$ -st (i.e. the last) slice with respect to the bottom vertex. They appear in two horizontal rows — the position of these rows depending on a . Again, all possible paths λ at worst pick up the same octahedron parameters as in γ , so that the contribution to $\{U^\gamma, U^\lambda\}$ vanishes.

Now we consider the general situation. The paths γ and λ may run together through a collection of tetrahedra, moving on different families of slices the entire way. In some tetrahedron Δ_0 they must enter on different faces and exit on the same face to start the run. In some other tetrahedron Δ_1 they enter together and exit on different faces to end the run. We may skip intermediate tetrahedra where they move together (because we know the segments in these tetrahedra commute), and simply assume that Δ_0 and Δ_1 are immediately adjacent. There are four basic cases to analyze.

First, it is possible for Δ_0 and Δ_1 to coincide — i.e. γ and λ both enter and exit a single tetrahedron separately. This is depicted in figure 74. If γ turns to the left, then its




Figure 74. Paths entering and exiting a single tetrahedron on four distinct faces.




Figure 75. Paths entering and exiting a pair of glued tetrahedra on four distinct faces: I.

single octahedron parameter appears on the last slice with respect to the bottom vertex (at a distance a away from the edge E). The only possible λ path crossing the small triangle with this parameter just picks up the parameter itself. If γ turns to the right, its $2a - 1$ octahedron parameters appear on the first through the $(K - a)$ -th slices with respect to the bottom vertex. On the first through the $(K - a + 2)$ -th slice the octahedron parameters are arranged in two 2-spiked crowns (as in section B.1, so every closed path commutes with them. On the $(K - a)$ -th and $(K - a + 1)$ -st slices the possible problematic λ paths just pick up the parameters themselves.

Otherwise, we may assume that Δ_0 and Δ_1 are distinct and adjacent. Then there are three remaining arrangements in which γ and λ may occur. They are depicted in figures 75–77. Since γ and λ run along two different families of slices, they distinguish two different vertices common to Δ_0 and Δ_1 , which in turn distinguishes a common edge E . We choose to position γ on slices centered around the “top” vertex of E and λ on slices centered around the bottom.

In each of these three cases, there are two options for orienting γ . It is then shown how its octahedron parameters appear on various slices with respect to the bottom vertex of E (the arrangements of $+$ ’s and $-$ ’s follow by combining the arrangements in figures 73 and 74). Now the potential λ paths do pick up parameters that don’t commute with those in U^γ . However, the non-commuting contributions always cancel in pairs — either within Δ_0 and Δ_1 independently, or between Δ_0 and Δ_1 . We let the figures speak for themselves.

This finally exhausts all possible ways in which the two paths γ and λ may come close enough to share an octahedron — and thus potentially have nontrivial contributions to their Poisson bracket. The *only* contributions that don’t vanish identically or cancel in pairs are those that come from paths that cross on adjacent (or identical) parallel slices,




Figure 76. Paths entering and exiting a pair of glued tetrahedra on four distinct faces: II.

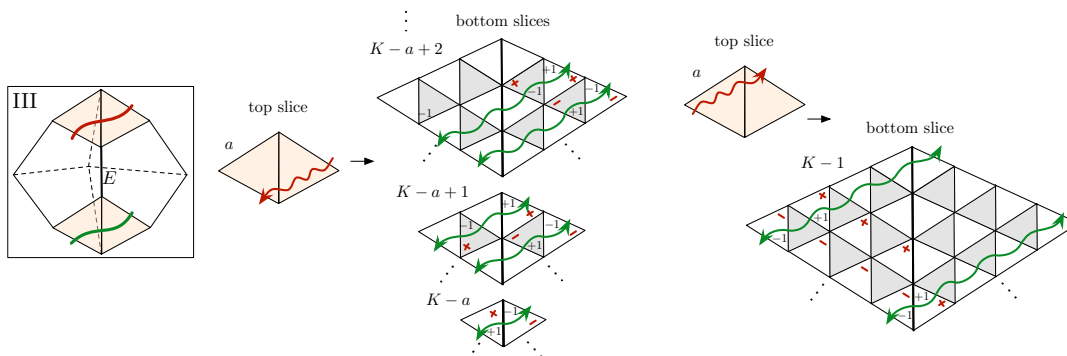


Figure 77. Paths entering and exiting a pair of glued tetrahedra on four distinct faces: III.

as in figure 72. The contribution (B.6) found there is therefore the full Poisson bracket $\{U^\gamma, U^\lambda\}$. This finishes the proof.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] T. Dimofte, D. Gaiotto and S. Gukov, *Gauge theories labelled by three-manifolds*, *Commun. Math. Phys.* **325** (2014) 367 [[arXiv:1108.4389](https://arxiv.org/abs/1108.4389)] [[INSPIRE](#)].
- [2] S.V. Matveev, *Transformations of special spines, and the Zeeman conjecture*, *Izv. Akad. Nauk SSSR Ser. Mat.* **51** (1987) 1104.
- [3] R. Piergallini, *Standard moves for standard polyhedra and spines*, *Rend. Circ. Mat. Palermo (2) Suppl.* (1988) 391.
- [4] O. Aharony, A. Hanany, K.A. Intriligator, N. Seiberg and M.J. Strassler, *Aspects of $N = 2$ supersymmetric gauge theories in three-dimensions*, *Nucl. Phys. B* **499** (1997) 67 [[hep-th/9703110](https://arxiv.org/abs/hep-th/9703110)] [[INSPIRE](#)].
- [5] W. Thurston, *The geometry and topology of three-manifolds*, lecture notes at Princeton University, Princeton, U.S.A. (1980).

- [6] E. Witten, $(2 + 1)$ -dimensional gravity as an exactly soluble system, *Nucl. Phys. B* **311** (1988) 46 [INSPIRE].
- [7] S. Gukov, *Three-dimensional quantum gravity, Chern-Simons theory and the A polynomial*, *Commun. Math. Phys.* **255** (2005) 577 [hep-th/0306165] [INSPIRE].
- [8] W.D. Neumann and D. Zagier, *Volumes of hyperbolic three-manifolds*, *Topology* **24** (1985) 307.
- [9] T. Dimofte, *Quantum riemann surfaces in Chern-Simons theory*, *Adv. Theor. Math. Phys.* **17** (2013) 479 [arXiv:1102.4847] [INSPIRE].
- [10] 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) 1895 [math/0604094] [INSPIRE].
- [11] T. Dimofte, S. Gukov, J. Lenells and D. Zagier, *Exact results for perturbative Chern-Simons theory with complex gauge group*, *Commun. Num. Theor. Phys.* **3** (2009) 363 [arXiv:0903.2472] [INSPIRE].
- [12] V.V. Fock and A.B. Goncharov, *Moduli spaces of local systems and higher teichmuller theory*, *Publ. Math. Inst. Hautes Etudes Sci.* **103** (2006) 1 [math/0311149].
- [13] V.V. Fock and A.B. Goncharov, *The quantum dilogarithm and representations of quantum cluster varieties*, *Invent. Math.* **175** (2008) 223 [math/0702397].
- [14] V. Fock and A. Goncharov, *Symplectic double for moduli spaces of g -local systems on surfaces*, arXiv:1410.3526.
- [15] T. Dimofte, D. Gaiotto and R. van der Veen, *RG domain walls and hybrid triangulations*, *Adv. Theor. Math. Phys.* **19** (2015) 137 [arXiv:1304.6721] [INSPIRE].
- [16] D. Cooper, M. Culler, H. Gillet, D. Long and P. Shalen, *Plane curves associated to character varieties of 3-manifolds*, *Invent. Math.* **118** (1994) 47.
- [17] V.V. Fock and A.B. Goncharov, *Cluster ensembles, quantization and the dilogarithm*, math/0311245 [INSPIRE].
- [18] D.W. Boyd, F. Rodriguez-Villegas and N.M. Dunfield, *Mahler's measure and the dilogarithm (ii)*, *Canad. J. Math.* **54** (2002) 468 [math/0308041].
- [19] A. Champanerkar, *A-polynomial and bloch invariants of hyperbolic 3-manifolds*, Ph.D. Thesis, Columbia University, U.S.A. (2003).
- [20] S. Gukov and P. Sulkowski, *A-polynomial, B-model and quantization*, *JHEP* **02** (2012) 070 [arXiv:1108.0002] [INSPIRE].
- [21] A. Kapustin and E. Witten, *Electric-magnetic duality and the geometric Langlands program*, *Commun. Num. Theor. Phys.* **1** (2007) 1 [hep-th/0604151] [INSPIRE].
- [22] D. Gaiotto, G.W. Moore and A. Neitzke, *Four-dimensional wall-crossing via three-dimensional field theory*, *Commun. Math. Phys.* **299** (2010) 163 [arXiv:0807.4723] [INSPIRE].
- [23] N. Bergeron, E. Falbel, and A. Guilloux, *Tetrahedra of flags, volume and homology of $SL(3)$* , arXiv:1101.2742.
- [24] S. Garoufalidis, M. Goerner and C.K. Zickert, *Gluing equations for $PGL(n, c)$ -representations of 3-manifolds*, arXiv:1207.6711.

- [25] S. Garoufalidis, D.P. Thurston and C.K. Zickert, *The complex volume of $SL(n, c)$ -representations of 3-manifolds*, [arXiv:1111.2828](#).
- [26] C.K. Zickert, *The volume and chern-simons invariant of a representation*, *Duke Math. J.* **150** (2009) 489 [[arXiv:0710.2049](#)].
- [27] C.K. Zickert, *The extended bloch group and algebraic k-theory*, [arXiv:0910.4005](#).
- [28] D. Gaiotto, G.W. Moore and A. Neitzke, *Spectral networks*, *Annales Henri Poincaré* **14** (2013) 1643 [[arXiv:1204.4824](#)] [[INSPIRE](#)].
- [29] D. Gaiotto, G.W. Moore and A. Neitzke, *Spectral networks and snakes*, *Annales Henri Poincaré* **15** (2014) 61 [[arXiv:1209.0866](#)] [[INSPIRE](#)].
- [30] S. Cecotti, C. Cordova and C. Vafa, *Braids, walls and mirrors*, [arXiv:1110.2115](#) [[INSPIRE](#)].
- [31] C. Cordova, S. Espahbodi, B. Haghighat, A. Rastogi and C. Vafa, *Tangles, generalized Reidemeister moves and three-dimensional mirror symmetry*, *JHEP* **05** (2014) 014 [[arXiv:1211.3730](#)] [[INSPIRE](#)].
- [32] M. Henningson and K. Skenderis, *The holographic Weyl anomaly*, *JHEP* **07** (1998) 023 [[hep-th/9806087](#)] [[INSPIRE](#)].
- [33] J.A. Harvey, R. Minasian and G.W. Moore, *Non-Abelian tensor multiplet anomalies*, *JHEP* **09** (1998) 004 [[hep-th/9808060](#)] [[INSPIRE](#)].
- [34] H. Fuji, S. Gukov, M. Stosic and P. Sulkowski, *3D analogs of Argyres-Douglas theories and knot homologies*, *JHEP* **01** (2013) 175 [[arXiv:1209.1416](#)] [[INSPIRE](#)].
- [35] H.-J. Chung, T. Dimofte, S. Gukov and P. Sulkowski, *3D-3D correspondence revisited*, *JHEP* **04** (2016) 140 [[arXiv:1405.3663](#)] [[INSPIRE](#)].
- [36] S. Gukov and D. Pei, *Equivariant Verlinde formula from fivebranes and vortices*, [arXiv:1501.01310](#) [[INSPIRE](#)].
- [37] D. Pei and K. Ye, *A 3d-3d appetizer*, *JHEP* **11** (2016) 008 [[arXiv:1503.04809](#)] [[INSPIRE](#)].
- [38] T. Dimofte, S. Gukov and L. Hollands, *Vortex counting and lagrangian 3-manifolds*, *Lett. Math. Phys.* **98** (2011) 225 [[arXiv:1006.0977](#)] [[INSPIRE](#)].
- [39] Y. Terashima and M. Yamazaki, *$SL(2, \mathbb{R})$ Chern-Simons, Liouville and gauge theory on duality walls*, *JHEP* **08** (2011) 135 [[arXiv:1103.5748](#)] [[INSPIRE](#)].
- [40] C. Cordova and D.L. Jafferis, *Complex Chern-Simons from M5-branes on the squashed three-sphere*, [arXiv:1305.2891](#) [[INSPIRE](#)].
- [41] S. Lee and M. Yamazaki, *3D Chern-Simons theory from M5-branes*, *JHEP* **12** (2013) 035 [[arXiv:1305.2429](#)] [[INSPIRE](#)].
- [42] T. Dimofte, *3d superconformal theories from three-manifolds*, in *New dualities of supersymmetric gauge theories*, J. Teschner ed., Springer, Germany (2016) [[arXiv:1412.7129](#)].
- [43] D. Gang, N. Kim, M. Romo and M. Yamazaki, *Aspects of defects in 3d-3d correspondence*, *JHEP* **10** (2016) 062 [[arXiv:1510.05011](#)] [[INSPIRE](#)].
- [44] E. Witten, *Solutions of four-dimensional field theories via M-theory*, *Nucl. Phys. B* **500** (1997) 3 [[hep-th/9703166](#)] [[INSPIRE](#)].
- [45] D. Gaiotto, *$N = 2$ dualities*, *JHEP* **08** (2012) 034 [[arXiv:0904.2715](#)] [[INSPIRE](#)].

- [46] L.F. Alday, D. Gaiotto and Y. Tachikawa, *Liouville correlation functions from four-dimensional gauge theories*, *Lett. Math. Phys.* **91** (2010) 167 [[arXiv:0906.3219](#)] [[INSPIRE](#)].
- [47] A. Gadde, S. Gukov and P. Putrov, *Fivebranes and 4-manifolds*, [arXiv:1306.4320](#) [[INSPIRE](#)].
- [48] M.T. Anderson, C. Beem, N. Bobev and L. Rastelli, *Holographic uniformization*, *Commun. Math. Phys.* **318** (2013) 429 [[arXiv:1109.3724](#)] [[INSPIRE](#)].
- [49] Y. Tachikawa, *On the 6d origin of discrete additional data of 4d gauge theories*, *JHEP* **05** (2014) 020 [[arXiv:1309.0697](#)] [[INSPIRE](#)].
- [50] D. Gaiotto, G.W. Moore and A. Neitzke, *Framed BPS states*, *Adv. Theor. Math. Phys.* **17** (2013) 241 [[arXiv:1006.0146](#)] [[INSPIRE](#)].
- [51] S. Gukov and E. Witten, *Gauge theory, ramification, and the geometric Langlands program*, *Curr. Devel. Math.* **2006** (2008) 35 [[hep-th/0612073](#)] [[INSPIRE](#)].
- [52] L. Chekhov and V.V. Fock, *Quantum Teichmüller space*, *Theor. Math. Phys.* **120** (1999) 1245 [[math/9908165](#)] [[INSPIRE](#)].
- [53] R.M. Kashaev, *Quantization of Teichmüller spaces and the quantum dilogarithm*, *Lett. Math. Phys.* **43** (1998) 105 [[INSPIRE](#)].
- [54] T. Dimofte, D. Gaiotto and S. Gukov, *3-manifolds and 3d indices*, *Adv. Theor. Math. Phys.* **17** (2013) 975 [[arXiv:1112.5179](#)] [[INSPIRE](#)].
- [55] T. Dimofte, *Complex Chern-Simons theory at level k via the 3d-3d correspondence*, *Commun. Math. Phys.* **339** (2015) 619 [[arXiv:1409.0857](#)] [[INSPIRE](#)].
- [56] N.A. Nekrasov and S.L. Shatashvili, *Supersymmetric vacua and Bethe ansatz*, *Nucl. Phys. Proc. Suppl.* **192-193** (2009) 91 [[arXiv:0901.4744](#)] [[INSPIRE](#)].
- [57] B. Eynard and N. Orantin, *Invariants of algebraic curves and topological expansion*, *Commun. Num. Theor. Phys.* **1** (2007) 347 [[math-ph/0702045](#)] [[INSPIRE](#)].
- [58] G. Borot and B. Eynard, *All-order asymptotics of hyperbolic knot invariants from non-perturbative topological recursion of A -polynomials*, [arXiv:1205.2261](#) [[INSPIRE](#)].
- [59] J. Ellegaard Andersen and R. Kashaev, *A TQFT from quantum Teichmüller theory*, *Commun. Math. Phys.* **330** (2014) 887 [[arXiv:1109.6295](#)] [[INSPIRE](#)].
- [60] J. Ellegaard Andersen and R. Kashaev, *A new formulation of the Teichmüller TQFT*, [arXiv:1305.4291](#) [[INSPIRE](#)].
- [61] S. Garoufalidis, *The 3d index of an ideal triangulation and angle structures*, [arXiv:1208.1663](#).
- [62] S. Garoufalidis, C.D. Hodgson, J.H. Rubinstein and H. Segerman, *1-efficient triangulations and the index of a cusped hyperbolic 3-manifold*, [arXiv:1303.5278](#) [[INSPIRE](#)].
- [63] C. Frohman, R. Gelca and W. Lofaro, *The a -polynomial from the noncommutative viewpoint*, *Trans. Amer. Math. Soc.* **354** (2002) 735 [[math/9812048](#)].
- [64] A. Sikora, *Quantizations of character varieties and quantum knot invariants*, [arXiv:0807.0943](#).
- [65] J.E. Andersen and R. Kashaev, *Complex quantum Chern-Simons*, [arXiv:1409.1208](#) [[INSPIRE](#)].

- [66] A.B. Goncharov, *Explicit construction of characteristic classes*, in *I.M. Gelfand Seminar*, S. Gelfand and S. Gindikin eds., American Mathematical Society, Providence U.S.A. (1993).
- [67] A.B. Goncharov, *Geometry of configurations, polylogarithms, and motivic cohomology*, *Adv. Math.* **114** (1995) 197.
- [68] F. Bonahon, *A Schlafli-type formula for convex cores of hyperbolic 3-manifolds*, *J. Diff. Geom.* **50** (1998) 25 [[dg-ga/9704017v1](#)].
- [69] A.A. Suslin, *Homology of $GL(n)$, characteristic classes, and Milnor's k -theory*, *Trudy Mat. Inst. Steklov* **165** (1984) 188.
- [70] W. Neumann, *Combinatorics of triangulations and the chern-simons invariant for hyperbolic 3-manifolds*, in *Topology '90*, B.N. Apanasov et al. eds., Ohio State University Mathematical Research Institute Publications volume 1, U.S.A. (1992).
- [71] S. Garoufalidis and C.K. Zickert, *The symplectic properties of the $PGL(n, C)$ -gluing equations*, [arXiv:1310.2497](#) [[INSPIRE](#)].
- [72] T. Dimofte and R. van der Veen, *A spectral perspective on Neumann-Zagier*, [arXiv:1403.5215](#) [[INSPIRE](#)].
- [73] Y. Kabaya, *Parametrization of $PSL(2, C)$ -representations of surface groups*, [arXiv:1110.6674](#).
- [74] S. Garoufalidis and T.T. Le, *The colored jones function is q -holonomic*, *Geom. Topol.* **9** (2005) 1253 [[math/0309214](#)].
- [75] S. Garoufalidis, *On the characteristic and deformation varieties of a knot*, *Geom. Topol. Monogr.* **7** (2004) 291 [[math/0306230](#)].
- [76] S. Garoufalidis and C. Koutschan, *The SL_3 Jones polynomial of the trefoil: a case study of q -holonomic sequences*, [arXiv:1011.6329](#) [[INSPIRE](#)].
- [77] H. Fuji, S. Gukov and P. Sulkowski, *Super- A -polynomial for knots and BPS states*, *Nucl. Phys. B* **867** (2013) 506 [[arXiv:1205.1515](#)] [[INSPIRE](#)].
- [78] S. Garoufalidis, *The colored HOMFLY polynomial is q -holonomic*, [arXiv:1211.6388](#) [[INSPIRE](#)].
- [79] A.M. Gabrielov, I. Gelfand and M.V. Losik, *Combinatorial computation of characteristic classes*, *Funct. Anal. Appl.* **9** (1975) 103.
- [80] A. Guilloux, *Representations of 3-manifold groups in $PGL(n, C)$ and their restriction to the boundary*, [arXiv:1310.2907](#).
- [81] A.A. Suslin, *K -theory of a field and the Bloch group*, *Proc. Steklov Inst. Math.* **4** (1991) 217.
- [82] A. Goncharov, *Volumes of hyperbolic manifolds and mixed Tate motives*, *J. Amer. Math. Soc.* **12** (1996) 569 [[alg-geom/9601021](#)].
- [83] A.B. Goncharov, *The classical polylogarithms, algebraic k -theory and $\zeta_f(n)$* , proceedings of the *Gelfand Seminar*, I.M. Gelfand et al. eds., Birkhauser, Germany (1993).
- [84] A.N. Redlich, *Gauge noninvariance and parity violation of three-dimensional fermions*, *Phys. Rev. Lett.* **52** (1984) 18 [[INSPIRE](#)].
- [85] A.N. Redlich, *Parity violation and gauge noninvariance of the effective gauge field action in three dimensions*, *Phys. Rev. D* **29** (1984) 2366.

- [86] C. Closset, T.T. Dumitrescu, G. Festuccia, Z. Komargodski and N. Seiberg, *Contact terms, unitarity and F-maximization in three-dimensional superconformal theories*, *JHEP* **10** (2012) 053 [[arXiv:1205.4142](#)] [[INSPIRE](#)].
- [87] C. Closset, T.T. Dumitrescu, G. Festuccia, Z. Komargodski and N. Seiberg, *Comments on Chern-Simons contact terms in three dimensions*, *JHEP* **09** (2012) 091 [[arXiv:1206.5218](#)] [[INSPIRE](#)].
- [88] T. Dimofte and S. Gukov, *Chern-Simons theory and S-duality*, *JHEP* **05** (2013) 109 [[arXiv:1106.4550](#)] [[INSPIRE](#)].
- [89] K.A. Intriligator and N. Seiberg, *Mirror symmetry in three-dimensional gauge theories*, *Phys. Lett. B* **387** (1996) 513 [[hep-th/9607207](#)] [[INSPIRE](#)].
- [90] J. de Boer, K. Hori, H. Ooguri, Y. Oz and Z. Yin, *Mirror symmetry in three-dimensional theories, $SL(2, \mathbb{Z})$ and D-brane moduli spaces*, *Nucl. Phys. B* **493** (1997) 148 [[hep-th/9612131](#)] [[INSPIRE](#)].
- [91] J. de Boer, K. Hori and Y. Oz, *Dynamics of $N = 2$ supersymmetric gauge theories in three-dimensions*, *Nucl. Phys. B* **500** (1997) 163 [[hep-th/9703100](#)] [[INSPIRE](#)].
- [92] E. Witten, *$SL(2, \mathbb{Z})$ action on three-dimensional conformal field theories with Abelian symmetry*, [hep-th/0307041](#) [[INSPIRE](#)].
- [93] H. Ooguri and C. Vafa, *Knot invariants and topological strings*, *Nucl. Phys. B* **577** (2000) 419 [[hep-th/9912123](#)] [[INSPIRE](#)].
- [94] E. Witten, *Fivebranes and knots*, *Quantum Topol.* **3** (2012) 1 [[arXiv:1101.3216](#)] [[INSPIRE](#)].
- [95] D. Gaiotto and E. Witten, *S-duality of Boundary Conditions In $N = 4$ Super Yang-Mills Theory*, *Adv. Theor. Math. Phys.* **13** (2009) 721 [[arXiv:0807.3720](#)] [[INSPIRE](#)].
- [96] J.P. Gauntlett, N. Kim and D. Waldram, *M five-branes wrapped on supersymmetric cycles*, *Phys. Rev. D* **63** (2001) 126001 [[hep-th/0012195](#)] [[INSPIRE](#)].
- [97] M. Pernici and E. Sezgin, *Spontaneous compactification of seven-dimensional supergravity theories*, *Class. Quant. Grav.* **2** (1985) 673 [[INSPIRE](#)].
- [98] R. Emparan, C.V. Johnson and R.C. Myers, *Surface terms as counterterms in the AdS/CFT correspondence*, *Phys. Rev. D* **60** (1999) 104001 [[hep-th/9903238](#)] [[INSPIRE](#)].
- [99] D. Gaiotto and J. Maldacena, *The gravity duals of $N = 2$ superconformal field theories*, *JHEP* **10** (2012) 189 [[arXiv:0904.4466](#)] [[INSPIRE](#)].
- [100] D. Green, Z. Komargodski, N. Seiberg, Y. Tachikawa and B. Wecht, *Exactly marginal deformations and global symmetries*, *JHEP* **06** (2010) 106 [[arXiv:1005.3546](#)] [[INSPIRE](#)].
- [101] T. Yoshida, *The η -invariant of hyperbolic 3-manifolds*, *Invent. Math.* **81** (1985) 473.
- [102] D.L. Jafferis, *The exact superconformal R-symmetry extremizes Z*, *JHEP* **05** (2012) 159 [[arXiv:1012.3210](#)] [[INSPIRE](#)].
- [103] D.L. Jafferis, I.R. Klebanov, S.S. Pufu and B.R. Safdi, *Towards the F-theorem: $N = 2$ field theories on the three-Sphere*, *JHEP* **06** (2011) 102 [[arXiv:1103.1181](#)] [[INSPIRE](#)].
- [104] A. Kapustin, B. Willett and I. Yaakov, *Exact results for Wilson loops in superconformal Chern-Simons theories with matter*, *JHEP* **03** (2010) 089 [[arXiv:0909.4559](#)] [[INSPIRE](#)].
- [105] N. Hama, K. Hosomichi and S. Lee, *SUSY gauge theories on squashed three-spheres*, *JHEP* **05** (2011) 014 [[arXiv:1102.4716](#)] [[INSPIRE](#)].

- [106] D. Martelli, A. Passias and J. Sparks, *The gravity dual of supersymmetric gauge theories on a squashed three-sphere*, *Nucl. Phys. B* **864** (2012) 840 [[arXiv:1110.6400](#)] [[INSPIRE](#)].
- [107] R. M. Kashaev, *The hyperbolic volume of knots from quantum dilogarithm*, *Lett. Math. Phys.* **39** (1997) 269 [[q-alg/9601025v2](#)].
- [108] H. Murakami and J. Murakami, *The colored Jones polynomials and the simplicial volume of a knot*, *Acta Math.* **186** (2001) 85 [[math/9905075](#)].
- [109] E. Witten, *Analytic continuation of Chern-Simons theory*, *AMS/IP Stud. Adv. Math.* **50** (2011) 347 [[arXiv:1001.2933](#)] [[INSPIRE](#)].
- [110] T.D. Dimofte and S. Garoufalidis, *The quantum content of the gluing equations*, *Geom. Topol.* **17** (2013) 1253 [[arXiv:1202.6268](#)] [[INSPIRE](#)].
- [111] S. Kim, *The Complete superconformal index for $N = 6$ Chern-Simons theory*, *Nucl. Phys. B* **821** (2009) 241 [*Erratum ibid.* **B 864** (2012) 884] [[arXiv:0903.4172](#)] [[INSPIRE](#)].
- [112] Y. Imamura and S. Yokoyama, *Index for three dimensional superconformal field theories with general R-charge assignments*, *JHEP* **04** (2011) 007 [[arXiv:1101.0557](#)] [[INSPIRE](#)].
- [113] A. Kapustin and B. Willett, *Generalized superconformal index for three dimensional field theories*, [arXiv:1106.2484](#) [[INSPIRE](#)].
- [114] F. Benini, T. Nishioka and M. Yamazaki, *4D index to 3D index and 2d TQFT*, *Phys. Rev. D* **86** (2012) 065015 [[arXiv:1109.0283](#)] [[INSPIRE](#)].
- [115] Y. Imamura and D. Yokoyama, *S^3/\mathbb{Z}_n partition function and dualities*, *JHEP* **11** (2012) 122 [[arXiv:1208.1404](#)] [[INSPIRE](#)].
- [116] Y. Imamura, H. Matsuno and D. Yokoyama, *Factorization of the S^3/\mathbb{Z}_n partition function*, *Phys. Rev. D* **89** (2014) 085003 [[arXiv:1311.2371](#)] [[INSPIRE](#)].
- [117] C. Beem, T. Dimofte and S. Pasquetti, *Holomorphic blocks in three dimensions*, *JHEP* **12** (2014) 177 [[arXiv:1211.1986](#)] [[INSPIRE](#)].