Proc. London Math. Soc. (3) 115 (2017) 411–447
doi:10.1112/plms.12045

Spectrally similar incommensurable 3-manifolds

David Futer and Christian Millichap

Abstract

Reid has asked whether hyperbolic manifolds with the same geodesic length spectrum must be commensurable. Building toward a negative answer to this question, we construct examples of hyperbolic 3-manifolds that share an arbitrarily large portion of the length spectrum but are not commensurable. More precisely, for every $n \gg 0$, we construct a pair of incommensurable hyperbolic 3-manifolds N_n and N_n^μ whose volume is approximately n and whose length spectra agree up to length n.

Both N_n and N_n^μ are built by gluing two standard submanifolds along a complicated pseudo-Anosov map, ensuring that these manifolds have a very thick collar about an essential surface. The two gluing maps differ by a hyper-elliptic involution along this surface. Our proof also involves a new commensurability criterion based on pairs of pants.

1. *Introduction*

This paper is devoted to the following question, posed and studied by Reid [**49**, **50**]:

Question 1.1. Let $M_1 = \mathbb{H}^n/\Gamma_1$ and $M_2 = \mathbb{H}^n/\Gamma_2$ be hyperbolic manifolds of finite volume. If the length spectra of the M_i agree, must M_1 and M_2 be commensurable?

All manifolds appearing in this paper are presumed to be orientable. The *length spectrum* of a manifold M, denoted $\mathcal{L}(M)$, is the ordered tuple of all lengths of closed geodesics in M, counted with multiplicity. Meanwhile, M_1 and M_2 are called *commensurable* if some hyperbolic manifold $\hat{M}$ serves as a common finite-sheeted cover of both M_1 and M_2. Commensurability is an equivalence relation, and the equivalence class containing M is called the *commensurability class* of M. Reid's question can be rephrased to ask: does the length spectrum of M determine the commensurability class of M?

It is well-known that the length spectrum $\mathcal{L}(M)$ does not determine the isometry class of M. In the setting of hyperbolic manifolds, the first counterexamples are due to Vignéras [**58**]. Sunada [**54**] gave a general, group-theoretic method that allows one to start with a hyperbolic manifold M_0 and construct finite covers M_1 and M_2 that share the same length spectrum but are not isometric. All examples produced using Sunada's method, as well as the arithmetic examples produced by Vignéras, are commensurable by construction. This common feature was part of the motivation behind Question 1.1.

1.1. *Main results*

In this paper, we provide some evidence toward a negative answer to Question 1.1. We show that a large finite portion of the length spectrum $\mathcal{L}(M)$ does not determine the commensurability class of M.

Received 6 September 2016; published online 11 May 2017.

2010 *Mathematics Subject Classification* 57M50, 30F40, 58J53, 53C22.

Futer was supported in part by NSF grant DMS-1408682 and the Elinor Lunder Founders' Circle Membership at the Institute for Advanced Study.

THEOREM 1.2. *For all sufficiently large $n \in \mathbb{N}$, there exists a pair of non-isometric finite-volume hyperbolic 3-manifolds $\{N_n, N_n^\mu\}$ such that:*

(1) $\mathrm{vol}(N_n) = \mathrm{vol}(N_n^\mu)$, *where this volume grows coarsely linearly with n;*
(2) *the (complex) length spectra of N_n and N_n^μ agree up to length at least n;*
(3) *N_n and N_n^μ have at least e^n/n closed geodesics up to length n;*
(4) *each of N_n and N_n^μ is the unique minimal orbifold in its commensurability class. In particular, N_n and N_n^μ are incommensurable.*

The manifolds N_n and N_n^μ can be taken to be either closed or one-cusped. The statement that volume grows *coarsely linearly* means that there exist absolute constants $A, A' > 0$ such that $An \leqslant \mathrm{vol}(N_n) = \mathrm{vol}(N_n^\mu) \leqslant A'n$.

To place Theorem 1.2 in context, it helps to introduce the notions of commensurators and arithmeticity. The commensurability class of a hyperbolic manifold $M = \mathbb{H}^n/\Gamma$ is entirely determined by its (orientable) *commensurator*, namely

$$C^+(M) = C^+(\Gamma) = \{g \in \mathrm{Isom}^+(\mathbb{H}^n) : \Gamma \cap g\Gamma g^{-1} \text{ is finite index in both } \Gamma \text{ and } g\Gamma g^{-1}\}.$$

The commensurator $C^+(\Gamma)$ is always a group, but is not always discrete or torsion-free. By a fundamental theorem of Margulis, its discreteness obeys a striking dichotomy.

THEOREM 1.3 (Margulis [**38**]). *Let $M = \mathbb{H}^n/\Gamma$ be a hyperbolic n-manifold of finite volume.*

- *If M is arithmetic, the commensurator $C^+(\Gamma)$ is dense in $\mathrm{Isom}^+(\mathbb{H}^n)$. The commensurability class of M contains infinitely many minimal elements.*
- *If M is non-arithmetic, the commensurator $C^+(\Gamma)$ is discrete in $\mathrm{Isom}^+(\mathbb{H}^n)$, hence $\mathcal{O} = \mathbb{H}^n/C^+(\Gamma)$ is the unique minimal orbifold in the commensurability class of M.*

See Maclachlan and Reid [**35**] for the definition of an arithmetic manifold. One way to interpret Margulis's theorem is that it is harder for non-arithmetic manifolds to be commensurable. Indeed, the manifolds N_n and N_n^μ in Theorem 1.2 are non-arithmetic.

By contrast, Question 1.1 has a positive answer for arithmetic manifolds of dimension $n \neq 1 \bmod 4$. In other words, the length spectrum of such a manifold determines its commensurability class. This result is due to Reid [**48**] in dimension 2; to Chinburg, Hamilton, Long, and Reid [**20**] in dimension 3; and to Prasad and Rapinchuk [**46**, **47**] in all remaining dimensions satisfying $n \neq 1 \bmod 4$.

The contrast between arithmetic and non-arithmetic settings is further sharpened by the following counterpart to Theorem 1.2: a finite portion of the length spectrum of an arithmetic 3-manifold *does* determine its commensurability class.

THEOREM 1.4 (Linowitz–McReynolds–Pollack–Thomspon [**33**]). *There exists an absolute constant $c > 0$ such that the following holds. Let M_1, M_2 be closed arithmetic hyperbolic 3-manifolds with volume less than V. If the length spectra $\mathcal{L}(M_1)$ and $\mathcal{L}(M_2)$ agree for all lengths less than $c \cdot \exp(\log V^{\log V})$, then M_1 and M_2 are commensurable.*

Both Theorem 1.2 and Theorem 1.4 have (easier) analogues in dimension 2. See [**33**, Theorem 1.1] and Observation 1.6 below.

The length cutoffs in Theorems 1.2 and 1.4 are qualitatively different, as the function $(\log V)^{\log V}$ grows faster than any polynomial in V. Thus the length cutoff in Theorem 1.4 is considerably higher than the cutoff of Theorem 1.2, which is linear in the volume V.

Outside the realm of hyperbolic manifolds, the analogue of Question 1.1 can have a negative answer for quotients of higher-rank symmetric spaces. For instance, Lubotzky, Samuels, and Vishne have constructed infinite families of closed arithmetic manifolds modeled on

$PGL_n(\mathbb{R})/PO_n(\mathbb{R})$ for $n \geqslant 3$ that have the same length spectrum but are not commensurable [**34**].

To our knowledge, the only prior work on Question 1.1 for non-arithmetic hyperbolic manifolds is due to Millichap [**41**]. He constructed a sequence of knots $K_n, K_n^\mu \subset S^3$ with incommensurable complements, such that their volumes grow linearly with n and such that $S^3 \smallsetminus K_n$ and $S^3 \smallsetminus K_n^\mu$ share the same n shortest geodesics. However, all of the geodesics in his construction have length uniformly bounded by 0.015. Improving upon Millichap's result, we have the following analogue of Theorem 1.2 for knot complements in S^3.

THEOREM 1.5. *For each $n \gg 0$, there exists a pair of non-isometric hyperbolic knot exteriors $E_n = S^3 \smallsetminus K_n$ and $E_n^\mu = S^3 \smallsetminus K_n^\mu$ such that:*

(1) $\mathrm{vol}(E_n) = \mathrm{vol}(E_n^\mu)$, *where this volume grows coarsely linearly with n;*
(2) *the (complex) length spectra of E_n and E_n^μ agree up to length at least $2\log(n)$;*
(3) *E_n and E_n^μ have at least $n^2/(2\log(n))$ closed geodesics up to length $2\log(n)$;*
(4) *E_n is the unique minimal orbifold in its commensurability class, and the only knot complement in its commensurability class. The same statement holds for E_n^μ.*

The knots K_n in Theorem 1.5 can be explicitly described via diagrams (see Figure 7).

The primary difference between Theorem 1.2 and Theorem 1.5 is that in the latter, the length spectrum is only preserved up to a cutoff that grows logarithmically with volume. This contrasts with linear growth in Theorem 1.2 and super-polynomial growth in Theorem 1.4. As a consequence, the number of geodesics that E_n and E_n^μ share is also lower than the corresponding count for N_n and N_n^μ. Although both constructions involve a large product region (whose thickness is essentially n), closed geodesics in E_n and E_n^μ can take shortcuts through the cusp, which results in a logarithmic penalty.

On the other hand, the commensurability statement in Theorem 1.5 is stronger than the one in Theorem 1.2. We are able to conclude that E_n and E_n^μ are the only knot complements in their commensurability classes, by relying on a theorem of Reid and Walsh [**51**, Proposition 5.1] that is specific to hyperbolic knot complements in S^3.

1.2. *Constructing spectrally similar manifolds*

To motivate the construction behind Theorems 1.2 and 1.5, we first sketch a similar result in dimension 2.

OBSERVATION 1.6. Let S be an (orientable) surface of Euler characteristic $\chi(S) < -1$. Then, for all $n > 0$, S admits a pair of complete hyperbolic structures Σ_n and Σ'_n such that:

(1) $\mathrm{area}(\Sigma_n) = \mathrm{area}(\Sigma'_n) = -2\pi\chi(S)$;
(2) the length spectra of Σ_n and Σ'_n agree up to length at least n;
(3) Σ_n and Σ'_n are incommensurable.

Observation 1.6 can be proved as follows. The hypothesis $\chi(S) < -1$ ensures that S contains an essential, separating, simple closed curve γ. Start with a hyperbolic metric Σ_n in which γ is short enough to have a collar of diameter $n/2$. By the collar lemma [**17**], it suffices to have $\coth(\ell(\gamma)/2) > \cosh(n/4)$. Now, build Σ'_n by cutting S along γ, twisting some distance, and then regluing (see Figure 1).

Since γ is separating, and has a collar of diameter $n/2$, any closed geodesic shorter than n must be disjoint from γ. All such geodesics are preserved when we cut and twist along γ, verifying (2). Twisting along γ produces uncountably many distinct metrics on S, only finitely

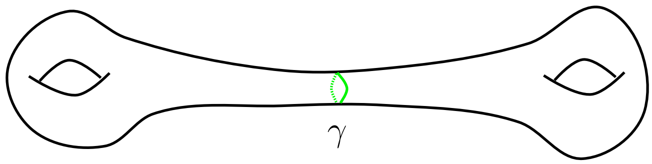

FIGURE 1 (colour online). *A closed genus two surface, whose hyperbolic structure Σ_n has a large collar about a separating geodesic γ. To construct Σ_n', we cut Σ_n along γ, twist by some distance, and re-glue.*

many of which can be commensurable with Σ_n. Thus a generic choice of twisting distance will also satisfy (3).

Conclusion (1), a restatement of the Gauss–Bonnet theorem, is included to illustrate the following point. One can get the length spectra of incommensurable surfaces Σ_n and Σ_n' to agree up to length n while keeping the areas uniformly bounded. This differs from the setting of Theorems 1.2 and 1.5, where our construction requires the volume to grow with the length cutoff. It also implies that for $n \gg 0$, at least one of Σ_n or Σ_n' must be non-arithmetic, as otherwise they would have to be commensurable by [**33**, Theorem 1.1].

Our 3-dimensional construction requires a lot more machinery, but the intuitive idea of Observation 1.6 still applies. First, we produce a 3-manifold N_n with a very thick collar about some separating surface F_n. The modern theory of Kleinian groups (see Section 3) gives a number of ways to do this, the simplest of which is the following. We start with a standard pair of 3-manifolds with boundary, denoted T and B, such that $\partial T \cong \partial B \cong S$. Then, we glue these pieces via some large iterate φ^{2n} of a pseudo-Anosov map $\varphi : S \to S$. A visual summary of the process appears in Figure 6, and the details appear in Section 5. We show in Proposition 5.4 that as $n \to \infty$, both $\mathrm{vol}(N_n)$ and the diameter of a collar about F_n will grow linearly with n.

Once we have built a 3-manifold N_n, with a thick collar about F_n, we obtain N_n^μ by a cut-and-paste operation, just as in Figure 1. However, by Mostow rigidity, most ways of cutting and regluing along F_n will not preserve any of the fine-scale geometry of N_n. To avoid this problem, we cut N_n along F_n and reglue along a hyper-elliptic involution μ, in a procedure called *mutation*. This forces F_n to be a closed surface of genus 2 or one of several small surfaces with punctures; see Definition 5.5 for details. By a theorem of Ruberman [**53**], the mutation process can be realized as a rigid cut-and-paste operation along a minimal surface, preserving the geometry of N_n on either side of this surface. In particular, $\mathrm{vol}(N_n) = \mathrm{vol}(N_n^\mu)$. Furthermore, a theorem of Millichap [**41**, Proposition 4.4] says that any closed geodesic that can be *homotoped* disjoint from F_n will remain unperturbed under mutation.

To get our count of closed geodesics up to length n, we rely on the work of Huber [**31**], Margulis [**37**], and Gangolli and Warner [**25**]. They showed that for all $L \gg 0$, a finite-volume hyperbolic 3-manifold M has approximately $e^{2L}/2L$ geodesics up to length L. See Section 4 and equation (4.1). However, this theorem is not uniform: the length cutoff L_0 after which the asymptotic estimate applies depends on the manifold M. In a manifold such as N_n, which has a 'bottleneck' surface of small area, the Margulis asymptotic may not apply until a length cutoff much larger than n. On the other hand, Theorem 1.2 features a uniform count of closed geodesics that applies for all $n \gg 0$.

The solution to this dilemma is to count closed geodesics in X_n, the cover of N_n (and N_n^μ) corresponding to the surface F_n. These manifolds X_n converge geometrically and algebraically to the infinite cyclic cover of the mapping torus M_φ, permitting a uniform count of their closed geodesics. See Proposition 4.5 for details. Once we have shown that X_n has the desired number of geodesics, we argue that these geodesics project down to closed geodesics in N_n and N_n^μ in a one-to-one fashion.

The knot complements E_n and E_n^μ from Theorem 1.5 are constructed in a very similar fashion, using nearly identical machinery. See Figure 7 for a summary and Section 6 for all the details. As noted earlier, the main difference from Theorem 1.2 occurs in the cutoff length for the length spectra, which results from geodesics traveling through the cusp.

1.3. *Ruling out commensurability*

In Observation 1.6, a simple cardinality argument shows that a generic choice of twisting along γ produces a surface incommensurable with Σ_n. Showing that a pair of 3-manifolds is incommensurable typically requires stronger tools. This can be accomplished using algebraic invariants such as the invariant trace field or Bloch invariant (see, for example, Chesebro and DeBlois [**19**]), or geometric invariants such as the canonical polyhedral decomposition (see Goodman, Heard, and Hodgson [**27**]). All of these invariants use the global geometry of the ambient manifold.

By contrast, Theorems 1.2 and 1.5 use a commensurability criterion that is much more local in nature. The criterion is based around thrice-punctured spheres, or *pairs of pants*. By a theorem of Adams [**1**], every essential pair of pants in a cusped hyperbolic 3-manifold is isotopic to a totally geodesic surface. Furthermore, all totally geodesic pairs of pants are isometric. This rigid geometry provides a lot of structure.

In the following theorem, a choice of cusp neighborhoods in a 3-manifold M induces a choice of cusp neighborhoods in a totally geodesic pair of pants $P \subset M$. We say that P is *pairwise tangent* with respect to a horocusp $\mathcal{C} \subset M$ if the three cusp neighborhoods of $P \cap \mathcal{C}$ are tangent in pairs. The notion of P being *geometrically isolated* is slightly harder to define; see Definition 2.2 for details. In addition, see Figure 5 for a visual description; the pairs of pants P and P' in that figure are each *geometrically isolated on one side*.

THEOREM 1.7. *Let M be a finite-volume hyperbolic 3-manifold with exactly three cusps. Let C_1, C_2, C_3 be embedded horospherical neighborhoods of the cusps of M. Suppose that M contains exactly two pairs of pants P and P' that are pairwise tangent and geometrically isolated on one side, with respect to C_1, C_2, C_3. Suppose that each of P and P' meets every C_i, that P and P' are disjointly embedded, and that $P \cup P'$ cuts N into a pair of submanifolds M_+ and M_-, where M_+ is asymmetric and $\mathrm{vol}(M_+) \neq \mathrm{vol}(M_-)$.*

Let s_i be a Dehn surgery coefficient on ∂C_i. Then, for all choices of s_i that are sufficiently long and sufficiently different, the filled manifold $M(s_1, s_2, s_3)$ is hyperbolic, non-arithmetic, and minimal in its commensurability class. This includes the case where $s_1 = \infty$, that is, the cusp C_1 is left unfilled, and s_2, s_3 are sufficiently long and sufficiently different.

The meaning of *sufficiently different* in the statement of the theorem is that the ratios between the lengths of s_i are very large: $\ell(s_1) \gg \ell(s_2) \gg \ell(s_3)$. This meaning will be quantified and made more precise during the proof of the theorem in Section 2.2.

Despite its somewhat complicated statement, Theorem 1.7 is quite powerful. To apply the theorem, one needs to verify that the local geometry of M near two pairs of pants is as described. However, one needs no global information about M beyond the asymmetry of M_+ and M_- and the knowledge that all pairs of pants in M have been located.

In our application, the manifolds N_n, N_n^μ in Theorem 1.2 can each be described as a Dehn filling of some 3-cusped manifold M satisfying the necessary geometric conditions. See Figure 6 for a schematic. We can verify the hypotheses of Theorem 1.7 in infinitely many examples because the isometry class of M_+ will stay constant as n varies, and the theorem requires very little information about M_-. This will allow us to conclude that each of N_n, N_n^μ, E_n, and E_n^μ is minimal in its commensurability class.

The same type of local argument (involving a single horoball packing picture, in Figure 8) works to show that E_n and E_n^μ are the unique knot complements in their commensurability classes, establishing Theorem 1.5.

1.4. *Organization*

In Section 2, we explain and prove the panted commensurability criterion of Theorem 1.7. We also prove the closely related Theorem 2.3, which describes the full preimage of a pair of pants under certain special covering maps.

In Section 3, we review a number of notions and results from Kleinian groups. These results will be used to describe the large-scale geometry of certain submanifolds and covers of N_n and E_n. In particular, they will be used to show that, just as in Figure 1, each N_n contains a surface F_n with collar of thickness $n/2$.

In Section 4, we review the work of Margulis [**37**] and Gangolli and Warner [**25**] on counting closed geodesics in hyperbolic manifolds. We also use the work of Pollicott and Sharp [**45**] to obtain a uniform count of closed geodesics in a convergent sequence of surfaces (Proposition 4.1) and quasifuchsian 3-manifolds (Proposition 4.5).

In Section 5, we assemble these ingredients to prove Theorem 1.2, following the outline described above. Finally, in Section 6, we describe the similar construction of spectrally similar knots, proving Theorem 1.5.

2. *A panted commensurability criterion*

This section develops conditions under which a hyperbolic 3-manifold N is the minimal orbifold in its commensurability class. The main result, Theorem 1.7, will be applied in Sections 5 and 6 to show that our spectrally similar manifolds are minimal in their respective commensurability classes, hence incommensurable with one another. In order to prove Theorem 1.7, we first prove Theorem 2.3, which describes how pairs of pants with particular geometric properties behave under covering maps. Stating these theorems requires a handful of definitions.

DEFINITION 2.1. Let M be a non-compact hyperbolic 3-manifold with finite volume. Every non-compact end of M is a *cusp*, homeomorphic to $T^2 \times [0, \infty)$. Geometrically, each cusp is a quotient of a horoball in $\mathbb{H}^3$ by a $\mathbb{Z} \times \mathbb{Z}$ group of deck transformations. We call this geometrically standard end a *horospherical cusp neighborhood* or *horocusp*.

A collection of horospherical cusp neighborhoods $\mathcal{C} = \{C_1, \dots, C_k\}$ is called *maximal* if each C_i is embedded in M, but no C_i can be expanded further without overlapping itself or another horocusp. A maximal collection of horocusps in M lifts to a collection of horoballs in $\widetilde{M} = \mathbb{H}^3$, with each horoball tangent to some number of other horoballs. This is called a *horoball packing* of $\mathbb{H}^3$. See Figure 2 for an example in dimension 2.

In the upper half-space model of $\mathbb{H}^3$, a horoball packing can be moved by isometry so that some horoball H_∞ (covering some horocusp C_i of M) consists of all points above Euclidean height 1. Every horoball tangent to H_∞ is called *full-sized*, and has Euclidean diameter 1. Other horoballs will be smaller. The collection of all horoballs other than H_∞, as projected vertically to $\partial H_\infty = \mathbb{R}^2$, is called a *horoball diagram of C_i*. See Figure 5 for an example.

DEFINITION 2.2. Let M be a finite-volume hyperbolic 3-manifold, equipped with a collection of embedded cusp neighborhoods $\mathcal{C} = \{C_1, \dots, C_k\}$. Let $P \subset M$ be a pair of pants. By a theorem of Adams [**1**], P is totally geodesic. We say that P is *pairwise tangent* with respect to $\mathcal{C}$ if the three cusp neighborhoods of $P \cap \mathcal{C}$ are tangent in pairs.

If P is pairwise tangent, every cusp neighborhood in P will lift to a *distinguished line* of pairwise tangent full-sized horoballs in the horoball diagram of the corresponding cusp $C_i \subset M$. This line has a particular slope, determined by the isotopy class of $P \cap \partial C_i$. (See Figure 5, where the slope of $P \cap \partial C_i$ is 0.) Observe that a choice of transverse orientation on P determines a transverse orientation on the line of full-sized horoballs.

We say that P is *geometrically isolated* if the full-sized horoballs corresponding to a component of $P \cap C_i$ are only tangent to full-sized horoballs that lie on the distinguished line. Meanwhile, P is *geometrically isolated on one side* if the full-sized horoballs corresponding to a component of $P \cap C_i$ are only tangent to full-sized horoballs that lie on the distinguished line, or in a fixed transverse direction from the line. In other words, the full-sized horoballs on the distinguished line cannot be tangent to other full-sized horoballs on both sides of the distinguished line. For example, the pairs of pants P and P' depicted in Figure 5 are geometrically isolated to one side.

THEOREM 2.3. *Let M be a cusped hyperbolic 3-manifold, equipped with a collection of maximal horocusps. Let P be a thrice punctured sphere that meets every cusp of M, which is pairwise tangent and geometrically isolated on one side.*

Consider a covering map ψ from M to an orbifold $\mathcal{O}$, where the cusp neighborhoods of M are equivariant with respect to the covering. Then the full preimage $\psi^{-1} \circ \psi(P)$ is a disjoint union of thrice punctured spheres, and furthermore any component $Q \subset \psi^{-1} \circ \psi(P)$ is also pairwise tangent and geometrically isolated on one side.

REMARK 2.4. The hypothesis that the cusp neighborhoods in M are equivariant with respect to the covering projection $\psi : M \to \mathcal{O}$ is somewhat restrictive. Without knowing $\mathcal{O}$ in advance, it is hard to know whether a given choice of cusps in M will be equivariant. In our application (Theorem 1.7), we use a trick involving geodesic lengths to check that this hypothesis is satisfied.

There is an alternate version of Theorem 2.3 that does not require an equivariant choice of cusp neighborhoods. However, this alternate version of the theorem needs a stronger form of geometric isolation. Instead of considering full-sized horoballs that are tangent, as in Definition 2.2, one needs to consider horoballs whose distance (along a lift of ∂C_i) is no greater than $2\sqrt{2}$. If one such horoball lies on the distinguished line, then the other horoball must also lie on the distinguished line or in a fixed transverse direction. With this stronger hypothesis, one can relax the equivariance of cusps and still reach the same conclusion about the preimage $\psi^{-1} \circ \psi(P)$.

The proof of this alternate statement is considerably more involved, as it requires bounding the factor by which the cusps of M must be resized in order to become equivariant. It also requires a detailed case-by-case analysis using the work of Adams on cusp areas of 3-orbifolds with rigid cusps [**3**]. We have decided to omit this alternate statement, as it is not needed for our application.

2.1. *Proving Theorem* 2.3

Before outlining the proof of Theorem 2.3, we recall some facts about Euclidean 2-orbifolds.

DEFINITION 2.5. An (orientable) *Euclidean* 2-*orbifold* is a quotient $\mathbb{R}^2/\Gamma$, where Γ is a discrete subgroup of $\mathrm{Isom}^+(\mathbb{R}^2)$. By the Gauss–Bonnet theorem, such an orbifold must have Euler characteristic 0. It follows that a compact (orientable) Euclidean 2-orbifold is either a torus, a pillowcase, or a turnover. Here, a *pillowcase* is a 2-sphere with 4 cone points of cone angle π. A *turnover* is a 2-sphere with 3 cone points of cone angles $2\pi/p$, $2\pi/q$, and $2\pi/r$, where the triple (p, q, r) is one of $(2, 3, 6)$, $(2, 4, 4)$, or $(3, 3, 3)$. Euclidean turnovers are called *rigid*, because they admit a unique Euclidean structure up to scaling. See, for example, [**22**, Theorem 2.2.3] for background.

The proof of Theorem 2.3 contains the following steps. The first step, which involves ruling out rigid cusps in the quotient orbifold $\mathcal{O}$, also plays a major role in studying commensurability of knot complements. See [**42**, Proposition 9.1], and compare Lemma 6.5.

1. Show that a cusp cross-section of $\mathcal{O}$ cannot be rigid. In other words, the cross-section of every cusp of $\mathcal{O}$ is a torus or pillowcase. We do this in Lemma 2.6.
2. Show that $\psi(P)$ has no transverse self-intersections in $\mathcal{O}$. We do this in Lemma 2.8.
3. It follows from Lemma 2.6 and Lemma 2.8 that all the cusp neighborhoods of $\psi^{-1} \circ \psi(P)$ are parallel to those of P. We may reassemble them to form several pairs of pants.

The following notation will be in use throughout the proof of Theorem 2.3. Let p_1, p_2, p_3 be the three punctures of P. Then puncture p_i is mapped to a maximal horocusp $C_i \subset M$. It may be the case that C_i coincides with C_j for some $i \neq j$, but we maintain the separate names. Let D_i be the component of $P \cap C_i$ corresponding to puncture p_i. By hypothesis of Theorem 2.3, the 2-dimensional horocusps D_1, D_2, D_3 are disjointly embedded and tangent in pairs. This means

$$\text{area}(D_i) = \ell(\partial D_i) = 2 \qquad \text{for all } i.$$

Lemma 2.6. *Let M be a cusped hyperbolic 3-manifold. Suppose that M has a collection of maximal cusps and a thrice-punctured sphere P satisfying the hypotheses of Theorem 2.3.*

Let $\psi : M \to \mathcal{O}$ be a covering map, where the cusp neighborhoods of M cover those of $\mathcal{O}$. Then every cusp cross-section of $\mathcal{O}$ is a torus or pillowcase. Furthermore, curves or arcs of the form $\psi(\partial D_i)$ realize exactly one slope on each cusp of $\mathcal{O}$.

Recall that a *slope* on a torus T^2 is an isotopy class of essential simple closed curves. These isotopy classes on T^2 are in natural 1-to-1 correspondence with $\mathbb{Q} \cup \{\infty\}$. In a covering map $T^2 \to R$, where R is a pillowcase, simple closed curves on T^2 project to closed curves or arcs between singular points on R. We say that curves or arcs on R have the *same slope* if they lift to closed curves on T^2 with the same slope. Equivalently, a closed curve and a curve or arc on R have the same slope if they can be realized disjointly.

Proof of Lemma 2.6. Let $T_i = \partial C_i$ be a cusp torus of M, and let $R_i = \psi(T_i)$. Observe that $\widetilde{R_i} = \widetilde{T_i} = \mathbb{R}^2$ contains a distinguished line of pairwise tangent full-sized horoballs, corresponding to $\widetilde{P}$. This line of horoballs projects to a particular slope on T_i, namely ∂D_i.

Suppose, for a contradiction, that R_i is a (rigid) Euclidean turnover. Then $\pi_1(R_i)$ contains rotations of order 3 or 4. Thus each distinguished line of full-sized horoballs coming from $\widetilde{P}$ must intersect another line of full-sized horoballs at an angle of $2\pi/3$ or $\pi/2$. But the only way for two lines of pairwise tangent full-sized horoballs to intersect is if they share a horoball. This contradicts the hypothesis that P is geometrically isolated on one side.

We have now established that R_i is a torus or pillowcase, hence isotopy classes of curves and arcs on R_i have well-defined slopes. It may very well happen that $\psi(T_i)$ and $\psi(T_j)$ are the same 2-orbifold $R_i = R_j$.

Suppose, for a contradiction, that $\psi(\partial D_i)$ and $\psi(\partial D_j)$ represent distinct slopes on $R_i = R_j$. This means that $\psi(\partial D_i)$ and $\psi(\partial D_j)$ have a transverse intersection, possibly at a cone point of R_i. But then, lifting everything to the universal cover $\widetilde{R_i} = \mathbb{R}^2$, we again find two different intersecting lines of full-sized horoballs. As above, this contradicts the hypothesis that P is geometrically isolated on one side. □

The following lemma will help us rule out self-intersections in $\psi(P)$.

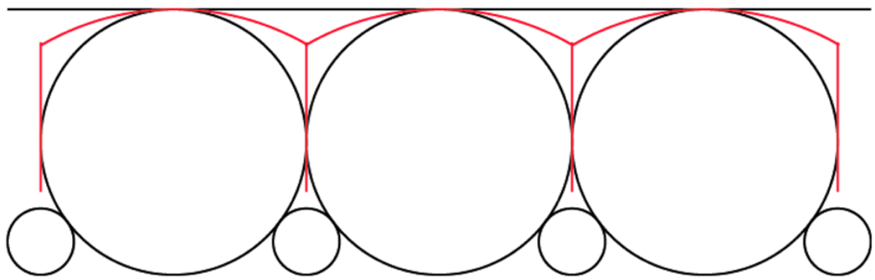

FIGURE 2 (colour online). *The Ford–Voronoi domain for the pairwise tangent cusp neighborhoods in the pair of pants P. Horoballs are shown in black, while the boundaries of the Ford–Voronoi cells are in red (lighter). If the horoball about ∞ lies at Euclidean height 1, the trivalent red vertices lie at height $\sqrt{3}/2$, which means the distance between them is* $\log(2/\sqrt{3})$.

LEMMA 2.7. *Let M and P be as in Theorem 2.3. Let D_1, D_2, D_3 be pairwise tangent cusp neighborhoods in P. Then any point $x \in P$ is within distance $\log(2/\sqrt{3})$ of some D_i.*

Furthermore, suppose that γ is a geodesic path in M, of length $\ell(\gamma) \leqslant \log(2/\sqrt{3})$, whose origin is at $x \in P$ and whose endpoint runs perpendicularly into some torus ∂C_j. Then $\gamma \subset P$.

Proof. We begin by recalling the *Ford–Voronoi domain* of P. For every i, let $E_i \subset P$ be the closure of the set of all points of P that lie closer to D_i than to any other D_j. Because the cusp neighborhoods D_i are pairwise tangent, the collection $\{D_i\}$ is invariant under all the symmetries of P. It follows that every E_i is topologically a once-punctured disk, whose non-ideal boundary consists of two geodesic segments meeting at angles of $2\pi/3$. These vertices of angle $2\pi/3$ are the triple points where E_1, E_2, E_3 all meet in P. A computation in $\mathbb{H}^2$, using Figure 2, shows that

$$\max_{x \in E_i} d(x, D_i) = \log \frac{2}{\sqrt{3}} = 0.1438\ldots, \tag{2.1}$$

with the maximum realized at the vertices. This proves the first assertion of the lemma.

Now, suppose that $\gamma \subset M$ is a geodesic path as in the statement of the lemma. In the universal cover $\widetilde{M} = \mathbb{H}^3$, there is a geodesic segment $\widetilde{\gamma}$ that starts at a point $\widetilde{x}$ covering x and runs perpendicularly to a horoball $\widetilde{C}_j$. The geodesic $\widetilde{\gamma}$ is the unique shortest path from $\widetilde{x}$ to $\partial\widetilde{C}_j$. Furthermore, $\widetilde{x}$ is contained in some Ford–Voronoi cell $\widetilde{E}_i$. Let $\widetilde{C}_i$ be the horoball containing the ideal point of $\widetilde{E}_i$; this will be one of the horoballs in $\mathbb{H}^3$ closest to $\widetilde{x}$.

We may assume without loss of generality that $\widetilde{C}_i$ is a horoball about ∞ in the upper half-space model of $\mathbb{H}^3$, positioned by isometry so that $\partial\widetilde{C}_i$ lies at height 1. This means that the copy of $\widetilde{P} \subset \mathbb{H}^3$ containing $\widetilde{E}_i$ is a vertical plane, with $\widetilde{E}_i$ being the cell about ∞, as in Figure 2.

If $\widetilde{C}_i = \widetilde{C}_j$, then the shortest path from $\widetilde{x}$ to $\partial\widetilde{C}_j$ will be contained in $\widetilde{E}_i$, because $\widetilde{E}_i$ is totally geodesic. Thus $\widetilde{\gamma} \subset \widetilde{E}_i$, which implies $\gamma \subset P$. Similarly, if $\widetilde{C}_j$ is a full-sized horoball that is tangent to $\widetilde{E}_i$, then the shortest path $\widetilde{\gamma}$ from $\widetilde{x}$ to $\widetilde{C}_j$ will be contained in the totally geodesic surface $\widetilde{P}$.

Finally, suppose for a contradiction that $\widetilde{C}_j$ is neither the horoball $\widetilde{C}_i$ about ∞, nor a full-sized horoball that meets our chosen vertical copy of $\widetilde{P}$. We may construct a piecewise geodesic path $\overline{\gamma}$ that follows a vertical geodesic from $\partial\widetilde{C}_i$ to $\widetilde{x}$ and then follows $\widetilde{\gamma}$ from $\widetilde{x}$ to $\partial\widetilde{C}_j$. The hypothesis on $\ell(\gamma)$, combined with (2.1), tells us that $\ell(\overline{\gamma}) \leqslant 2\log(2/\sqrt{3}) < 0.288$. Thus the Euclidean radius of $\widetilde{C}_j$ satisfies

$$\operatorname{rad}(\widetilde{C}_j) = \frac{1}{2\exp(d(\widetilde{C}_i, \widetilde{C}_j))} \geqslant \frac{1}{2\exp(\ell(\overline{\gamma}))} \geqslant \frac{1}{2}\left(\frac{\sqrt{3}}{2}\right)^2 = \frac{3}{8}.$$

As $\widetilde{C_j}$ must be disjoint from the full-sized horoballs that meet $\widetilde{P}$, a calculation using the Pythagorean theorem shows that the center of $\widetilde{C_j}$ must be at (Euclidean) distance at least $1/\sqrt{2}$ from the vertical copy of $\widetilde{P}$, with the lower bound on distance growing faster than $\mathrm{rad}(\widetilde{C_j})$ as the radius grows. Since $\mathrm{rad}(\widetilde{C_j}) \geqslant 3/8$, every point of $\widetilde{C_j}$ is at Euclidean distance at least $1/\sqrt{2} - 3/8 > 0.332$ from $\widetilde{P}$, hence $\ell(\overline{\gamma}) > 0.332$. On the other hand, we have previously seen that $\ell(\overline{\gamma}) < 0.288$. This contradiction shows that $\widetilde{\gamma} \subset \widetilde{P}$. □

LEMMA 2.8. *Under the hypotheses of Theorem 2.3, $\psi(P)$ has no transverse self-intersections.*

We remark that Lemma 2.8 does not require $\psi|_P$ to be 1-to-1. For instance, we may have a non-trivial covering map $\psi : P \to R$, where $R \subset \mathcal{O}$ is a totally geodesic 2-orbifold. See Figure 3 for an example of such a covering map.

Proof. We begin by noting that the 2-dimensional cusp neighborhoods $D_i \subset P$ are disjointly embedded in 3-dimensional cusp neighborhoods $C_i \subset M$. Furthermore, $\psi(C_i)$ and $\psi(C_j)$ are either disjoint or coincide. Hence the only way that $\psi(D_i)$ can have a transverse intersection with $\psi(D_j)$ is if $\psi(C_i) = \psi(C_j)$, and furthermore $\psi(\partial D_i)$ and $\psi(\partial D_j)$ represent distinct slopes. This possibility is ruled out by Lemma 2.6.

Next, we will use the lack of transverse self-intersections between the cusp neighborhoods to rule out transverse self-intersections elsewhere in $\psi(P)$. Recall the Ford–Voronoi decomposition of P into cells E_1, E_2, E_3, as in Lemma 2.7, and suppose for a contradiction that $\psi(P)$ has a transverse self-intersection. Then, for some pair (i, j), a point $x_i \in E_i$ has the same image as $x_j \in E_j$. Furthermore, in P there are disjoint neighborhoods U_i of x_i and U_j of x_j such that $\psi(U_i)$ meets $\psi(U_j)$ at a nonzero angle.

Let $\gamma \subset \psi(E_i)$ be the shortest path in $\mathcal{O}$ from $\psi(x_i)$ to $\psi(D_i)$. By (2.1), this path has length $\ell(\gamma) \leqslant \log(2/\sqrt{3})$. We consider two different lifts of γ to M. Let $\gamma_i \subset M$ be the lift of γ starting at x_i; this lift is contained in E_i by construction. But since $\psi(x_i) = \psi(x_j)$, there is also a lift γ_j of γ that starts at x_j. Since the cusps of M cover cusps of $\mathcal{O}$, each of γ_i and γ_j terminates on the boundary of some cusp, hitting the cusp perpendicularly. Thus, by Lemma 2.7, each of γ_i and γ_j lies in P.

It follows that $\psi(P)$ has a transverse self-intersection along all of $\gamma = \psi(\gamma_i) = \psi(\gamma_j)$. In particular, $\psi(P)$ has a transverse self-intersection at the endpoint of γ, which will have to be contained in $\psi(\partial D_i) \cap \psi(\partial D_j)$. But we have already checked that $\psi(D_i)$ has no transverse intersections with $\psi(D_j)$. □

Proof of Theorem 2.3. Since $\psi(P)$ has no transverse self-intersections by Lemma 2.8, we know that $\psi^{-1} \circ \psi(P)$ is a disjoint union of embedded, totally geodesic surfaces. We need to check that each component is a pair of pants.

Let Q be a component of $\psi^{-1} \circ \psi(P)$. Then the horoball packing and Ford–Voronoi decomposition of $\widetilde{Q}$ are isometric to that of $\widetilde{P}$, as depicted in Figure 2.

For each cusp C_i of M, Lemma 2.6 implies that each component of $Q \cap \partial C_i$ has the same slope as ∂D_i. Thus each component of $Q \cap C_i$ is an isometric copy of D_i, with area exactly 2. Since the density of the cusp neighborhoods in Q is the same as in P, namely $3/\pi$, each Ford–Voronoi cell of Q meeting C_i will have to be an isometric copy of E_i, with area $2\pi/3$ and two vertices of angle $2\pi/3$.

The only way to glue together several copies of E_1 and obtain a complete, connected surface is if three different copies are glued at every vertex. Thus Q is built from three copies of E_1 glued by isometry, hence Q is a pair of pants. In addition, Q is pairwise tangent and geometrically isolated on one side, because horoballs corresponding to $P \cap C_i$ will pull back to horoballs corresponding to $Q \cap C_i$. □

2.2. *Minimal manifolds in a commensurability class*

Our main application of Theorem 2.3 is Theorem 1.7. Before restating and proving the theorem, we need a definition.

DEFINITION 2.9. Let M be a cusped hyperbolic 3-manifold, equipped with a horocusp C. For a slope s on ∂C, define the *normalized length* of s to be

$$\hat{L}(s) = \ell(s)/\sqrt{\operatorname{area}(\partial C)},$$

where $\ell(s)$ is the length of a Euclidean geodesic in the isotopy class of s. The quantity $\hat{L}(s)$ is scale-invariant, hence does not depend on the choice of cusp neighborhood.

THEOREM 1.7. *Let M be a finite-volume hyperbolic 3-manifold with exactly three cusps. Let C_1, C_2, C_3 be embedded horospherical neighborhoods of the cusps of M. Suppose that M contains exactly two pairs of pants P and P' that are pairwise tangent and geometrically isolated on one side with respect to C_1, C_2, C_3. Suppose that each of P and P' meets every C_i, that P and P' are disjointly embedded, and that $P \cup P'$ cuts N into a pair of submanifolds M_+ and M_-, where M_+ is asymmetric and $\operatorname{vol}(M_+) \neq \operatorname{vol}(M_-)$.*

Let s_i be a Dehn surgery coefficients on ∂C_i. Then, for all choices of s_i that are sufficiently long and sufficiently different, the filled manifold $M(s_1, s_2, s_3)$ is hyperbolic, non-arithmetic, and minimal in its commensurability class. This includes the case where $s_1 = \infty$, that is, the cusp C_1 is left unfilled, and s_2, s_3 are sufficiently long and sufficiently different.

Proof. If each Dehn surgery slope $s_i \subset \partial C_i$ is sufficiently long, the filled 3-manifold $M(s_1, s_2, s_3)$ will be hyperbolic. By Gromov's theorem [**56**, Theorem 6.5.6], its volume is less than $\operatorname{vol}(M)$. Since there are only finitely many arithmetic hyperbolic manifolds of bounded volume (see Borel [**9**]), it follows that for sufficiently long s_i, the filled manifold $N = M(s_1, s_2, s_3)$ will be non-arithmetic.

For each i, let $\hat{L}_i = \ell(s_i)/\sqrt{\operatorname{area}(\partial C_i)}$ denote the normalized length of s_i. When each s_i is sufficiently long, the core of the i-th solid torus will be isotopic to a geodesic γ_i. Neumann and Zagier [**43**, Proposition 4.3] showed that the length of this geodesic can be expressed as

$$\ell(\gamma_i) = \frac{2\pi}{\hat{L}_i^2} + O\left(\frac{1}{\hat{L}_i^4}\right). \tag{2.2}$$

See also Hodgson and Kerckhoff [**30**, Theorem 5.12] and Magid [**36**, Theorem 1.2(ii)] for explicit estimates on $\ell(\gamma_i)$.

Let $\mathcal{O}_{\min}$ denote the orientable hyperbolic 3-orbifold of minimal volume. Gehring, Marshall, and Martin [**26**, **39**] have identified this orbifold and showed that $\operatorname{vol}(\mathcal{O}_{\min}) \approx 0.03905$. We set $V = \operatorname{vol}(M)/\operatorname{vol}(\mathcal{O}_{\min})$. The precise meaning of 'sufficiently different' slope lengths is that (in some permutation of the indices) we have $(\hat{L}_1)^2 \gg V(\hat{L}_2)^2 \gg V^2(\hat{L}_3)^2$. By (2.2) and our choice of sufficiently long s_i (to overwhelm the $\hat{L}_i^{-4}$ error term), this implies

$$\ell(\gamma_3) > V\ell(\gamma_2) > V^2\ell(\gamma_1). \tag{2.3}$$

One final application of the 'sufficiently long' hypothesis ensures that $\gamma_1, \gamma_2, \gamma_3$ are the three shortest closed geodesics in N, by a factor of at least V.

Since N is non-arithmetic, Theorem 1.3 says there is a unique minimal orbifold $\mathcal{Q}$ in the commensurability class of N. Since $\operatorname{vol}(\mathcal{Q}) \geqslant \operatorname{vol}(\mathcal{O}_{\min})$ and $\operatorname{vol}(N) < \operatorname{vol}(M)$, the degree of the covering map $\psi : N \to \mathcal{Q}$ is bounded above by V. By (2.3), this means every γ_i is the complete preimage of its image, hence ψ restricts to a covering map $\psi|_M : M \to \mathcal{O}$, where $\mathcal{O}$ is a hyperbolic 3-orbifold homeomorphic to $\mathcal{Q} \smallsetminus \cup_i \psi(\gamma_i)$.

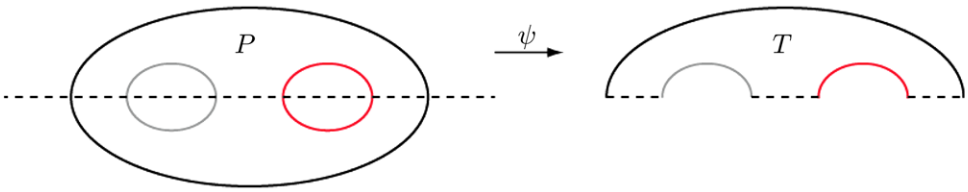

FIGURE 3 (colour online). *A quotient map $\psi : P \to T$, where P is a pair of pants and T is a 2-orbifold. If the three boundary components of P are sent to distinct boundary components of T, then either ψ is* 1*-to-*1 (*this case is not shown*)*, or T is an ideal triangle with mirrored boundary along the dashed arcs shown in the figure.*

From here on, the proof works by restricting the degree of ψ further and further, until we conclude that $\deg(\psi) = 1$, hence $N = \mathcal{Q}$.

Consider what ψ does to P. In the incomplete metric on M whose completion is N, the three boundary components of P are mapped to powers of $\gamma_1, \gamma_2, \gamma_3$. We have already checked that ψ maps these three closed geodesics to three distinct geodesics in $\mathcal{Q}$. Thus, in the restricted map $\psi|_M : M \to \mathcal{O}$, the three boundary components of P are sent to three distinct cusps of $\mathcal{O}$. This means that $\psi|_P$ is either 1-to-1 or a quotient by the hyper-elliptic involution; see Figure 3. In particular, $\psi|_P$ is at most a 2-to-1 map, and similarly for $\psi|_{P'}$.

Since the three cusps of M project to distinct cusps of $\mathcal{O}$, the cusp neighborhoods C_1, C_2, C_3 are automatically equivariant with respect to the cover $\psi : M \to \mathcal{O}$. Thus Theorem 2.3 applies.

Let $R = \psi(P \cup P') \subset \mathcal{O}$. Then we have

$$\psi^{-1}(R) = (\psi^{-1} \circ \psi(P \cup P')) = (\psi^{-1}\psi(P) \cup \psi^{-1}\psi(P')) \subset P \cup P'.$$

The final inclusion comes from Theorem 2.3, combined with the hypothesis that P and P' are the only pairwise-tangent, geometrically isolated pairs of pants in M. Since the opposite inclusion $P \cup P' \subset \psi^{-1}(R)$ is immediate, we conclude that $P \cup P' = \psi^{-1}(R)$. Thus a point of R has at most 4 preimages in M: at most 2 in P and at most 2 in P'. Since $\psi : M \to \mathcal{O}$ is orientation-preserving, the singular locus of $\mathcal{O}$ is at most 1-dimensional, hence some points of $R = \psi(P \cup P')$ must be non-singular. We conclude that $\deg(\psi) \leqslant 4$.

Let $\eta(P \cup P')$ be a regular neighborhood of $P \cup P'$, chosen to be equivariant with respect to ψ. Then $\partial\eta(P \cup P') \cap M_+ = P_+ \cup P'_+$, where P_+ and P'_+ are pairs of pants parallel to P and P', respectively. Similarly, $\partial\eta(P \cup P') \cap M_- = P_- \cup P'_-$. Since the neighborhood was chosen to be equivariant, $\eta(P \cup P')$ maps to a regular neighborhood $\eta(R)$. Note that if $\psi(P)$ is a pair of pants, then $\psi(P_+)$ will be a pair of pants parallel to $\psi(P)$. Meanwhile, if $\psi(P) = T$ is an ideal triangle with mirrored boundary, as in Figure 3, then $\psi(P_+)$ will still be a pair of pants that forms the entire boundary of $\eta(T)$. Thus $\psi|_{P_+}$ is 1-to-1, hence $\psi|_{P_+ \cup P'_+}$ has degree at most 2, and this degree is the same as that of $\psi|_{P_- \cup P'_-}$.

Now, consider the restriction of ψ to M_+ and M_-. Since $\partial(M_+ \setminus \eta(P \cup P')) = (P_+ \cup P'_+)$, and similarly for M_-, we have an equality of degrees:

$$\deg(\psi|_{M_+}) = \deg(\psi|_{P_+ \cup P'_+}) = \deg(\psi|_{P_- \cup P'_-}) = \deg(\psi|_{M_-}) \in \{1, 2\}. \tag{2.4}$$

We learn two things from (2.4). First, $\mathcal{O} \setminus R$ must have two connected components, namely $\psi(M_+)$ and $\psi(M_-)$. Otherwise, if $\mathcal{O} \setminus R$ is connected, each of M_+ and M_- would have to cover it. But then the equality of degrees $\deg(\psi|_{M_+}) = \deg(\psi|_{M_-})$ would imply $\mathrm{vol}(M_+) = \mathrm{vol}(M_-)$, contradicting our hypotheses. Thus $\psi(M_+)$ and $\psi(M_-)$ are disjoint subsets of $\mathcal{O}$, hence

$$\deg(\psi) = \deg(\psi|_M) = \deg(\psi|_{M_+}) \in \{1, 2\}.$$

Second, observe that a covering map of degree 2 must be regular, hence given by a symmetry of M. Any non-trivial covering transformation of M must respect $P \cup P'$, hence must either interchange the components $M_\pm$ or stabilize $M_\pm$. But M_+ cannot be interchanged with M_- because they have different volumes, and a symmetry cannot stabilize M_+ because we have assumed M_+ is asymmetric. It follows that $\deg(\psi) = 1$, hence $N = \mathcal{Q}$ is minimal in its commensurability class.

We close the proof by observing that the entire argument goes through if one cusp of M, say C_1, is left unfilled. If we choose long Dehn filling slopes s_2 and s_3 such that $(\hat{L}_2)^2 \gg V(\hat{L}_3)^2$, equation (2.3) would still hold with the convention that $\ell(\gamma_1) = 0$. Thus any covering map $\psi : N \to \mathcal{Q}$ would restrict to a covering map $\psi : M \to \mathcal{O}$, and the rest of the argument applies verbatim. □

3. *The geometry of pared convex cores*

Most of this section is a review of standard notions and results in the theory of Kleinian groups, particularly Kleinian surface groups. In Proposition 3.2, we describe a sequence of geometrically finite surface groups whose convex core boundaries are separated by distance approximately n. This construction will be used to build the spectrally similar manifolds in Section 5 and the spectrally similar knot complements in Section 6.

None of the results in this section are original. Our line of exposition is mostly modeled on that of Brock and Dunfield [**15**], as their construction has many similarities to ours. Another excellent reference for this material is Canary, Epstein, and Green [**18**].

3.1. *Kleinian surface groups and convergence*

Let $S = S_{g,k}$ be a compact (orientable) surface of genus g with k boundary components. We define the complexity $\xi(S) = 3g + k - 3$, and require that $\xi(S) > 0$. A representation $\rho : \pi_1(S) \to PSL(2,\mathbb{C})$ is called *type-preserving* if ρ maps peripheral loops in S to parabolic elements. Define the space

$$AH(S) = \{\rho : \pi_1(S) \to PSL(2,\mathbb{C}) \mid \rho \text{ discrete, faithful, type-preserving}\}/\sim,$$

where the equivalence relation $\sim$ is conjugation of the image in $PSL(2,\mathbb{C})$. The image $\rho(\pi_1 S)$ is called a *Kleinian surface group*. Every element of $AH(S)$ naturally corresponds to a hyperbolic 3-manifold $M = \mathbb{H}^3/\rho(\pi_1 S)$, homeomorphic to $S \times \mathbb{R}$ and equipped with a homotopy equivalence to S. This homotopy equivalence $S \to S \times \{*\} \subset M$, which is well-defined up to homotopy, is called a *marking*.

A hyperbolic 3-manifold $M = \mathbb{H}^3/\Gamma$ is determined up to isometry by the conjugacy class of Γ in $PSL(2,\mathbb{C})$. One way to pin down a particular conjugacy representative is to fix a *baseframe* ω in M (that is, a basepoint $x \in M$ together with an orthonormal frame in $T_x M$), and require that ω must lift to a fixed baseframe $\widetilde{\omega}$ at the origin in $\mathbb{H}^3$. Thus pairs (M, ω) consisting of a hyperbolic 3-manifold and a baseframe are in 1-to-1 correspondence with Kleinian groups $\Gamma \subset PSL(2,\mathbb{C})$, whereas un-framed manifolds are in 1-to-1 correspondence with conjugacy classes.

We will sometimes abuse notation by keeping baseframes (or equivalently, conjugacy representatives) implicit. Thus we will refer to either $\rho : \pi_1(S) \to PSL(2,\mathbb{C})$ or $M = \mathbb{H}^3/\rho(\pi_1 S)$ as elements of $AH(S)$.

The topology on $AH(S)$ is that of *algebraic convergence*. In this topology, a sequence of representations ρ_n converges to ρ_∞ if, for a generating set $\gamma_1, \dots, \gamma_k$ of $\pi_1(S)$, the sequence of matrices $\rho_n(\gamma_i) \in PSL(2,\mathbb{C})$ converges to $\rho_\infty(\gamma_i) \in PSL(2,\mathbb{C})$ for every i.

There is a finer topology on the space of (framed) hyperbolic manifolds, called the *Chabauty topology*. In this topology, a sequence of Kleinian groups $\Gamma_n \subset PSL(2,\mathbb{C})$ *converges geometrically* to $\Gamma \subset PSL(2,\mathbb{C})$ if the following hold:

(a) for each $\gamma \in \Gamma$, there are $\gamma_n \in \Gamma_n$ so that $\gamma_n \to \gamma$;

(b) if a sequence $\{\gamma_n \in \Gamma_n\}$ converges in $PSL(2, \mathbb{C})$, then the limit lies in Γ.

The Chabauty topology is metrizable [**18**, Proposition 3.1.2]. See Biringer [**5**] for an explicit metric.

In a sequence of representations $\rho_n : \pi_1 S \to PSL(2, \mathbb{C})$, with image $\Gamma_n = \rho_n(\pi_1 S)$, algebraic convergence only requires the ρ_n-images of the generators to converge, while geometric convergence requires the images of all group elements to converge. A sequence Γ_n converges *strongly* if it converges both algebraically and geometrically, and furthermore the algebraic and geometric limits coincide. Equivalently, a sequence of framed manifolds (M_n, ω_n) converges strongly to $(M_\infty, \omega_\infty)$ if the algebraic and geometric limits coincide.

There is an intrinsic description of geometric convergence, as follows. Let M_n be a sequence of complete hyperbolic manifolds, each equipped with a baseframe ω_n. Let M_∞ be another manifold with a baseframe ω_∞. For every $R > 0$, let $B_R(x_\infty) \subset M_\infty$ be the metric R-ball about the basepoint x_∞ of ω_∞. We say that $(M_\infty, \omega_\infty)$ is the *Gromov–Hausdorff limit* of (M_n, ω_n) if, for every $R > 0$, we have embeddings

$$\psi_{n,R} : (B_R(x_\infty), \omega_\infty) \hookrightarrow (M_n, \omega_n), \tag{3.1}$$

for all n sufficiently large, which converge to isometries in the C^∞ topology as $n \to \infty$. See [**4**, Theorem E.1.13] or [**18**, Theorem 3.2.9] for the equivalence between Gromov–Hausdorff convergence and convergence in the Chabauty topology.

3.2. *Limit sets and cores*

Given a hyperbolic 3-manifold $N = \mathbb{H}^3/\Gamma$, the *limit set* $\Lambda(\Gamma) \subset \partial\mathbb{H}^3$ is the set of accumulation points of any orbit Γx. The *convex core*, denoted $\mathrm{core}(N)$, is the smallest geodesically convex subset of N. Alternately, $\mathrm{core}(N)$ is the quotient by Γ of the convex hull of the limit set $\Lambda(\Gamma)$. If $\mathrm{core}(N)$ has finite volume, N is called *geometrically finite*.

Let $C^\epsilon(N) \subset N^{<\epsilon}$ be the disjoint union of ϵ-thin horospherical neighborhoods about the cusps of N (if any), for ϵ equal to the Margulis constant. Then the *pared convex core* is defined to be $\mathrm{core}^0(N) = \mathrm{core}(N) \smallsetminus C^\epsilon(N)$.

The set $QF(S) \subset AH(S)$ consists of the representations ρ for which $\Lambda(\rho(\pi_1 S))$ is a Jordan curve. In this case, both $\rho \in QF(S)$ and the associated manifold $M = \mathbb{H}^3/\rho(\pi_1 S)$ are called *quasifuchsian*. By a theorem of Bers, $QF(S) \cong \mathcal{T}(S) \times \mathcal{T}(S)$, where $\mathcal{T}(S)$ is the Teichmüller space of S and the two coordinates correspond to conformal structures on the two ends of $M \cong S \times \mathbb{R}$.

3.3. *Pared acylindrical manifolds*

We will study certain geometrically finite manifolds in $AH(S)$ whose hyperbolic structures correspond to special points of $\partial\mathcal{T}(S) \times \partial\mathcal{T}(S)$.

Let Q be a *pants decomposition* of S, namely a collection of disjointly embedded, essential simple closed curves cutting S into pairs of pants; then $|Q| = \xi(S) > 0$. A pants decomposition Q defines a point of $\partial\mathcal{T}(S)$, in which the lengths of curves in Q have been pinched all the way to 0; see Brock [**12**, p. 502] for more detail. A pair of pants decompositions (Q, Q') determines the pair

$$M_S(Q, Q') = \left(S \times [0,1],\ \ Q \times \{0\} \cup Q' \times \{1\} \cup \partial S \times \{\tfrac{1}{2}\}\right).$$

This 3-manifold is homeomorphic to $S \times [0, 1]$, with a designated *paring locus* along $Q \times \{0\} \cup Q' \times \{1\} \cup S \times \{\frac{1}{2}\}$. If no simple closed curve is isotopic into both Q and Q', the resulting 3-manifold will be *pared acylindrical*, meaning that it does not contain any essential annuli with boundary in the paring locus.

By Thurston's hyperbolization theorem [57], every pared acylindrical 3-manifold admits a finite-volume hyperbolic metric, with rank-1 cusps along every closed curve in the paring locus, and totally geodesic boundary along every pair of pants. Mostow–Prasad rigidity applied to the double of $M_S(Q,Q')$ verifies that this hyperbolic metric is unique up to isometry. The inclusion $S \to S \times \{\frac{1}{2}\} \subset M_S(Q,Q')$ is a homotopy equivalence, which specifies a marking of $M_S(Q,Q')$ by the surface S. Thus $M_S(Q,Q') \in AH(S)$.

The pared acylindrical manifold $M_S(Q,Q')$ constructed above is the convex core of a complete hyperbolic manifold $\hat{M}_S(Q,Q')$, which contains a flaring end for each totally geodesic pair of pants in $\partial M_S(Q,Q')$. The pared convex core, $\mathrm{core}^0 M_S(Q,Q')$, is a compact 3-manifold whose boundary is the union of horospherical annuli along the paring locus (including the annuli of $\partial S \times [0,1]$) and the ϵ-thick subsurfaces of the pants. As those pants live either in $S \times \{0\}$ or $S \times \{1\}$, we get a natural decomposition of the non-parabolic portion of $\partial \mathrm{core}^0 M_S(Q,Q')$ into *lower* and *upper* surfaces, denoted $\partial_- \mathrm{core}^0 M_S(Q,Q')$ and $\partial_+ \mathrm{core}^0 M_S(Q,Q')$. The lower surface can be identified as

$$\partial_- \mathrm{core}^0 M_S(Q,Q') = \partial\, \mathrm{core}^0 M_S(Q,Q') \cap (S \times \{0\}),$$

and similarly for the upper surface.

3.4. *Pseudo-Anosov double limits*

Given any pseudo-Anosov homeomorphism $\varphi : S \to S$, we can construct the mapping torus $M_\varphi = S \times [0,1]/(x,1) \sim (\varphi(x),0)$. By a theorem of Thurston [55], M_φ has a unique hyperbolic metric. Define $\widetilde{M}_\varphi$ to be the infinite cyclic cover of M_φ corresponding to $\pi_1(S)$. Then $S \times \mathbb{R} \cong \widetilde{M}_\varphi \in AH(S)$.

For a fixed pants decomposition Q on S, we consider the pared manifold

$$W_{\varphi,n} = M_S(\varphi^{-n}Q,\, \varphi^n Q).$$

For all n sufficiently large, all the curves of $\varphi^{-n}(Q)$ will be distinct from those of $\varphi^n(Q)$, hence $W_{\varphi,n}$ is pared acylindrical. See [15, Figure 3.7] for a helpful visual depiction of $M_S(\varphi^{-n}Q,\, \varphi^n Q)$, as well as $\widetilde{M}_\varphi$.

PROPOSITION 3.1. *For every baseframe ω_∞ on $\widetilde{M}_\varphi$, there is a sequence of baseframes ω_n on $W_{\varphi,n} = M_S(\varphi^{-n}Q, \varphi^n Q)$, such that $(W_{\varphi,n}, \omega_n)$ converges strongly to $(\widetilde{M}_\varphi, \omega_\infty)$.*

This result is due to Thurston [55], and forms the key step in his proof that M_φ is hyperbolic. We refer to [14, Theorem 1.2] for an alternate proof, and to [15, Proposition 3.4] for the version stated here.

We also learn that, in two different senses, the geometry of $W_{\varphi,n}$ grows linearly with n.

PROPOSITION 3.2. *There are positive constants $A_\varphi, A'_\varphi, B_\varphi, B'_\varphi$, depending on φ, such that the following holds. For all $n \gg 0$,*

$$A_\varphi \leqslant \frac{\mathrm{vol}(W_{\varphi,n})}{n} \leqslant A'_\varphi \tag{3.2}$$

and

$$B_\varphi \leqslant \frac{d_0(\partial_- \mathrm{core}^0 W_{\varphi,n},\, \partial_+ \mathrm{core}^0 W_{\varphi,n})}{n} \leqslant B'_\varphi \tag{3.3}$$

Here, $d_0(\cdot,\cdot)$ is the shortest length of a path from the lower boundary of $\mathrm{core}^0 W_{\varphi,n}$ to the upper boundary of $\mathrm{core}^0 W_{\varphi,n}$, among paths that remain inside $\mathrm{core}^0 W_{\varphi,n}$.

Proof. This follows as a consequence of theorems that relate the geometry of $W_{\varphi,n}$ to the coarse geometry of certain combinatorial graphs associated to the surface S.

The *curve graph*, denoted $\mathcal{C}(S)$, has vertices corresponding to isotopy classes of essential simple closed curves on S and edges corresponding to curves that can be realized disjointly. (When $\xi(S) = 1$, the definition must be modified slightly. Edges in $\mathcal{C}(S_{1,1})$ correspond to curves that intersect once, while edges in $\mathcal{C}(S_{0,4})$ correspond to curves that intersect twice.) The *pants graph*, denoted $\mathcal{P}(S)$, has vertices corresponding to isotopy classes of pants decompositions and edges corresponding to moves where a curve is removed, liberating a copy of $S_{1,1}$ or $S_{0,4}$ inside S, and replaced by another curve that intersects it a minimal number of times on that subsurface. See [**12**, Figure 3].

Each of $\mathcal{C}(S)$ and $\mathcal{P}(S)$ is endowed with the *graph metric*, in which every edge has length 1. The mapping class group $MCG(S)$ acts by isometries on these metric graphs. For any $\varphi \in MCG(S)$, there is a well-defined *stable translation length*

$$t_{\mathcal{C}}(\varphi) = \lim_{n\to\infty} \frac{d_{\mathcal{C}(S)}(v, \varphi^n(v))}{n}, \quad t_{\mathcal{P}}(\varphi) = \lim_{n\to\infty} \frac{d_{\mathcal{P}(S)}(v, \varphi^n(v))}{n}, \tag{3.4}$$

where the limit is independent of the base vertex v. If φ is a pseudo-Anosov element, we have $t_{\mathcal{P}}(\varphi) > 0$ and $t_{\mathcal{C}}(\varphi) > 0$; for $\mathcal{C}(S)$, the converse also holds.

With this background, equation (3.2) follows from a theorem of Brock [**12**]. He showed that there are positive constants K and K', depending only on S, such that

$$K\, d_{\mathcal{P}(S)}(\varphi^{-n}Q, \varphi^n Q) - K \leqslant \mathrm{vol}(W_{\varphi,n}) \leqslant K'\, d_{\mathcal{P}(S)}(\varphi^{-n}Q, \varphi^n Q).$$

(To interpret Brock's theorem in our setting, one needs to regard the pants decomposition Q as a point of the Weil–Petersson completion of $\mathcal{T}(S)$, denoted $\overline{\mathcal{T}(S)}$, which allows us to consider the pair $(\varphi^{-n}Q, \varphi^n Q) \in \overline{\mathcal{T}(S)} \times \overline{\mathcal{T}(S)}$. See [**12**, pp. 499 and 502].) After dividing by n and taking a limit as in (3.4), we obtain (3.2) with any values A_φ and A'_φ satisfying $0 < A_\varphi < 2Kt_{\mathcal{P}}(\varphi)$ and $A'_\varphi > 2K't_{\mathcal{P}}(\varphi)$.

Working toward (3.3), recall that $\partial \mathrm{core}^0\, W_{\varphi,n}$ is the union of ϵ-thick subsurfaces of the pants in $S \times \{0, 1\}$, plus horospherical annuli that correspond to the boundary of the ϵ-thin part about the paring locus. Here, as above, ϵ is the Margulis constant. Let γ be a component of Q, and let $T^\epsilon(\varphi^n\gamma)$ be the horospherical annulus of $\partial \mathrm{core}^0\, W_{\varphi,n}$ corresponding to the thin part about $\varphi^n(\gamma)$. Similarly, let $T^\epsilon(\varphi^{-n}\gamma)$ be the horospherical annulus corresponding to $\varphi^{-n}(\gamma)$.

With this setup, the lower bound of (3.3) follows from a theorem proved independently by Bowditch [**10**, Theorem 5.4] and Brock and Bromberg [**14**, Theorem 7.16]. They show that there exists a positive constant E, depending only on S, such that

$$E\, d_{\mathcal{C}(S)}(\varphi^{-n}\gamma,\, \varphi^n\gamma) - E \leqslant d_0(T^\epsilon(\varphi^{-n}\gamma),\, T^\epsilon(\varphi^n\gamma)). \tag{3.5}$$

When S is not a closed surface, this inequality crucially relies on measuring only the lengths of paths that stay inside the pared convex core, $\mathrm{core}^0\, W_{\varphi,n}$. See the remark following [**14**, Theorem 7.16]. Since $\partial_+ \mathrm{core}^0\, W_{\varphi,n}$ has bounded diameter (this fact also requires the core to be pared), every point of $\partial_+ \mathrm{core}^0\, W_{\varphi,n}$ lies at uniformly bounded distance from $T^\epsilon(\varphi^n\gamma)$. Thus, for some constant $E' > E$, we have

$$E\, d_{\mathcal{C}(S)}(\varphi^{-n}\gamma,\, \varphi^n\gamma) - E' \leqslant d_0(\partial_- \mathrm{core}^0\, W_{\varphi,n},\, \partial_+ \mathrm{core}^0\, W_{\varphi,n}).$$

Compare [**14**, Corollary 7.18] for a very similar statement. Now, dividing by n and taking a limit as in (3.4), we obtain the lower bound of (3.3) with any $B_\varphi \in (0, 2Et_{\mathcal{C}}(\varphi))$.

The upper bound of (3.3) also follows by considering $d_{\mathcal{C}(S)}(\varphi^{-n}\gamma,\, \varphi^n\gamma)$. A straighforward argument using the bounded diameter lemma for surfaces (see, for example, Biringer and Souto [**6**, Theorem 4.1]) shows that there is a positive function $F(\delta)$ depending on S such that if the injectivity radius of $\mathrm{core}^0\, W_{\varphi,n}$ is bounded below by δ, then

$$d_0(\partial_- \mathrm{core}^0\, W_{\varphi,n},\, \partial_+ \mathrm{core}^0\, W_{\varphi,n}) \leqslant F(\delta)\, d_{\mathcal{C}(S)}(\varphi^{-n}\gamma,\, \varphi^n\gamma) + F(\delta). \tag{3.6}$$

But Brock and Dunfield show in [**15**, Theorem 3.5] that for $n \gg 0$, $\operatorname{core}^0 W_{\varphi,n}$ satisfies

$$\operatorname{injrad}(\operatorname{core}^0 W_{\varphi,n}) \geqslant \operatorname{injrad}(\operatorname{core}^0 M_\varphi)/2 > 0,$$

everywhere outside a bounded-diameter collar of $\partial_\pm \operatorname{core}^0 W_{\varphi,n}$ whose geometry converges with n. Since this bounded collar must itself have injectivity radius bounded below, it follows that there is a value $\delta = \delta_\varphi$ that serves as a lower bound on $\operatorname{injrad}(\operatorname{core}^0 W_{\varphi,n})$ for all $n \gg 0$. Now, dividing (3.6) by n and taking a limit as in (3.4) yields the upper bound of (3.3). □

Remark 3.3. We make two observations about the statement and proof of Proposition 3.2. First, the upper bound of (3.3) depends rather inefficiently on φ. One could obtain a much tighter estimate by using the so-called *electric distance* between the upper and lower boundary of $\operatorname{core}^0 W_{\varphi,n}$, as in Bowditch [**10**, Theorem 5.4] and Biringer–Souto [**6**, Theorem 4.1]. Since the upper bound of (3.3) will not be needed in the sequel, and is included mainly for completeness, we chose the formulation that is easiest to state.

Second, the lower bound of equation (3.3) will be our main point of entry for the cutoff up to which the length spectra of N_n and N_n^μ agree in Theorem 1.2; see Proposition 5.4 and Lemma 5.9. To make this cutoff as simple as possible, it will be convenient to choose a pseudo-Anosov φ for which $2B_\varphi > 1$. As the above proof illustrates, it suffices to choose a φ for which $4Et_{\mathcal{C}}(\varphi) > 1$, where E comes from equation (3.5) and $t_{\mathcal{C}}(\varphi)$ is as in (3.4). For instance, by (3.4), a sufficiently high power of any pseudo-Anosov suffices.

4. *Counting geodesics, uniformly*

In the statement of Theorem 1.2, we claim that the manifolds N_n and N_n^μ contain a certain number of closed geodesics up to length n. For any finite-volume manifold $M = \mathbb{H}^d/\Gamma$, let $\mathcal{L}(M) = (\ell_1, \ell_2, \ldots)$ be the length spectrum of M, and define

$$\pi_M(L) = \max\{i : \ell_i \leqslant L\} = \big(\text{the number of closed geodesics in } M \text{ of length } \leqslant L\big).$$

The study of the asymptotic behavior of $\pi_M(L)$ has a long and distinguished history, starting with Huber [**31**] and Margulis [**37**]. They proved that for any closed hyperbolic manifold M of dimension d,

$$\pi_M(L) \sim \frac{e^{hL}}{hL}, \quad \text{where} \quad h = d-1. \tag{4.1}$$

Here and below, the notation $f(x) \sim g(x)$ means that $\lim_{x\to\infty} f(x)/g(x) = 1$. In fact, Margulis worked out a version of (4.1) in variable negative curvature for an appropriate definition of h. Gangolli and Warner showed that (4.1) also holds for non-compact hyperbolic d-manifolds of finite volume [**25**]. Among many other works in the subject, Roblin [**52**] proved an analogous result for geometrically finite hyperbolic 3-manifolds such as $M_S(Q, Q')$; see Remark 4.6.

Although equation (4.1) is beautiful in its simplicity, we cannot use it directly because the rate of convergence depends on the manifold M. On the other hand, we are looking for a uniform statement that will hold in all the manifolds N_n or all the $W_{\varphi,n}$, for $n \gg 0$. To obtain the desired result, we start with a uniform statement in bounded regions of Teichmüller space.

The Techmüller space $\mathcal{T}(S)$ can be interpreted as the space of discrete, faithful, type-preserving representations $\rho : \pi_1(S) \to PSL(2,\mathbb{R})$, up to conjugation in $PSL(2,\mathbb{R})$. Closed geodesics in S are in 1-to-1 correspondence with non-peripheral conjugacy classes in $\pi_1(S)$. Thus, given a hyperbolic structure $\Sigma \in \mathcal{T}(S)$, every conjugacy class $[\gamma]$ has an associated Σ-length $\ell(\gamma) = \ell_\Sigma(\gamma)$, which can be computed via

$$\cosh \frac{\ell(\gamma)}{2} = \frac{\operatorname{tr} \rho_\Sigma(\gamma)}{2},$$

where ρ_Σ is the representation corresponding to Σ. This makes it clear that $\ell_\Sigma(\gamma)$ is a conjugacy invariant, which varies continuously with Σ. The gist of the following proposition is that $\pi_\Sigma(L)$ also varies in a well-behaved fashion over $\mathcal{T}(S)$.

PROPOSITION 4.1. *Fix a topological surface $S = S_{g,k}$, and let $R \subset \mathcal{T}(S)$ be a compact region of the Teichmüller space of S. Then there is a constant L_0 depending on R, such that for every hyperbolic structure $\Sigma \in R$ and for every $L \geqslant L_0$, we have*

$$\pi_\Sigma(L) \geqslant \frac{e^L}{L}. \tag{4.2}$$

Proof. This follows from the work of Pollicott and Sharp [**45**], who worked out the error term in the asymptotic formula (4.1). They did this by studying the zeta function

$$\zeta(z) = \zeta_\Sigma(z) = \prod_{[\gamma]} (1 - e^{-z\ell(\gamma)})^{-1},$$

where $z \in \mathbb{C}$ and the product ranges over all the non-peripheral conjugacy classes in $\pi_1(\Sigma)$. By [**45**, Proposition 5], there is a constant $c_0(\Sigma) < 1$ such that $\zeta_\Sigma(z)$ converges to an analytic function on the half-plane $Re(z) > c_0(\Sigma)$, except for a simple pole at $z = 1$. In this region of convergence (which is uniform on compact sets), the analytic function $\zeta_\Sigma(z)$ depends continuously on Σ, hence $c_0(\Sigma)$ also depends continuously on Σ.

The continuous dependence on Σ means that on a compact set $R \subset \mathcal{T}(S)$, we may take a uniform constant

$$c_0 = c_0(R) = \max\{c_0(\Sigma) : \Sigma \in R\} < 1,$$

such that $\zeta_\Sigma(z)$ converges on $Re(z) > c_0(R)$ for every $\Sigma \in R$, except for a simple pole at $z = 1$. Making $c_0(R)$ larger if necessary, we may assume $c_0(R) \geqslant 0$.

Pollicott and Sharp then show that for every

$$c = c(\Sigma) \in \left(\frac{c_0(\Sigma) + 1}{2}, 1 \right)$$

the number of closed geodesics up to length L satisfies

$$\pi_\Sigma(L) = \mathrm{li}(e^L) + O(e^{cL}).$$

Here, $\mathrm{li}(y)$ is the logarithmic integral

$$\mathrm{li}(y) = \int_2^y \frac{du}{\log u}, \qquad \text{which satisfies} \quad \mathrm{li}(y) \sim \frac{y}{\log y} \quad \text{as} \quad y \to \infty. \tag{4.3}$$

See [**45**, Proposition 6 and p. 1033] for the definition of $c(\Sigma)$ in terms of $c_0(\Sigma)$.

Since there is a uniform value $c_0(R) < 1$ that works for every $\Sigma \in R$, we may also take a uniform value

$$c = c(R) \in \left(\frac{c_0(R) + 1}{2}, 1 \right).$$

The proofs of [**45**, Propositions 6 and 7] show that the multiplicative constant implicit in $O(\cdot)$ can also be taken uniformly for every $\Sigma \in R$. Thus, for positive constants $c = c(R) < 1$ and $A = A(R)$, we have

$$\pi_\Sigma(L) \geqslant \mathrm{li}(e^L) - Ae^{cL} \quad \forall\, \Sigma \in R. \tag{4.4}$$

It remains to derive (4.2) from (4.4). We do this via the following

CLAIM 4.2. *For every $A \in \mathbb{R}$ and every $c < 1$, we have* $\lim_{x\to\infty}(\mathrm{li}(e^x) - Ae^{cx} - e^x/x) = \infty$.

To prove the claim, we compute the derivative, using (4.3):

$$\frac{d}{dx}\left(\mathrm{li}(e^x) - Ae^{cx} - \frac{e^x}{x}\right) = \left(\frac{1}{\log(e^x)}\cdot e^x - Ac\,e^{cx}\right) - \left(\frac{e^x}{x} - \frac{e^x}{x^2}\right) = -Ac\,e^{cx} + \frac{e^x}{x^2},$$

which blows up as $x \to \infty$ whenever $c < 1$. Since the derivative is eventually very large, the function $(\mathrm{li}(e^x) - Ae^{cx} - e^x/x)$ will also blow up as $x \to \infty$.

Combining Claim 4.2 and (4.4) shows that there is a number $L_0 = L_0(R) > 0$, such that

$$\pi_\Sigma(L) \geqslant \mathrm{li}(e^L) - Ae^{cL} \geqslant \frac{e^L}{L} \quad \text{for} \quad \Sigma \in R, \quad L \geqslant L_0. \qquad \square$$

REMARK 4.3. In Proposition 4.1, it suffices to let $R \subset \mathcal{T}(S)$ be any subset whose image in the moduli space $\mathcal{M}(S) = \mathcal{T}(S)/MCG(S)$ is compact. For instance, one could let R be the entire *thick part* of $\mathcal{T}(S)$, consisting of all (marked) hyperbolic structures whose shortest geodesic has length at least $\epsilon > 0$. Since markings of S play no role in counting closed geodesics up to length L, compactness is only needed in the image of R in $\mathcal{M}(S)$.

The next result uses Proposition 4.1 to show that geometrically finite manifolds such as $W_{\varphi,n}$ contain the appropriate number of closed geodesics shorter than L. We move up to the 3-dimensional setting using the machinery of pleated surfaces, which we now recall.

DEFINITION 4.4. A *lamination* λ on a surface S is a 1-dimensional foliation of a closed subset of S. Every point of λ lies on a unique 1-dimensional leaf. The lamination λ is *finite* if it has finitely many leaves, and *filling* if every complementary region of $S \smallsetminus \lambda$ is an ideal triangle. If S is a closed surface, every finite lamination contains at least one closed leaf, together with finitely many lines asymptotic toward the closed leaves. If S has punctures, one example of a finite, filling lamination is an ideal triangulation.

Given a lamination λ on S, and a hyperbolic 3-manifold M, a *pleating map realizing* λ is a proper map $f : S \to M$, such that every leaf of λ is mapped to a geodesic and every component of $S \smallsetminus \lambda$ is mapped to a (possibly self-intersecting) totally geodesic surface. The image $f(S)$ is called a *pleated surface*. For $M \in AH(S)$, Thurston showed every finite, filling lamination λ is realized by a pleating map in the homotopy class of the marking, provided that no closed leaf of λ is homotopic to a parabolic in M. See [**56**, Proposition 9.7.1; **18**, Theorem 5.3.6].

PROPOSITION 4.5. *Fix a hyperbolic surface S with complexity $\xi(S) > 0$. Let $X_n \in AH(S)$ be a sequence of marked, geometrically finite hyperbolic 3-manifolds converging strongly to $X_\infty \in AH(S)$. Then there is a constant L_0 depending only on X_∞, and a constant n_0 depending on the sequence X_n, such that for all $L \geqslant L_0$ and $n \geqslant n_0$,*

$$\pi_{X_n}(L) \geqslant \frac{e^L}{L}.$$

Proof. Fix a finite filling lamination λ on S, such that the marking map $S \to X_\infty$ does not map any closed leaf of λ to a parabolic. Thurston's realization theorem [**56**, Proposition 9.7.1; **18**, Theorem 5.3.6] says that there is a pleating map $f_{\lambda,\infty} : S \to X_\infty$ realizing λ. Choose a basepoint $x_\infty \in X_\infty$ that lies in a totally geodesic triangle in the image of $f_{\lambda,\infty}$, together with an orthonormal frame ω_∞ at x_∞.

Since X_n converges to X_∞, a closed leaf of λ can only be parabolic in finitely many of the X_n. Thus, for $n \gg 0$, there is also a pleating map $f_{\lambda,n} : S \to X_n$ realizing λ. By pulling back the path-metric from X_n via $f_{\lambda,n}$, each one of these pleating maps endows S with a complete (marked) hyperbolic metric, denoted Σ_n.

Since X_n converges geometrically to X_∞, equation (3.1) says that the compact set $f_{\lambda,\infty}(S) \cap \mathrm{core}^0 X_\infty$ lies in a metric ball about x_∞ that is almost-isometric to a region of X_n for every large n. In particular, the 1-dimensional set $f_{\lambda,\infty}(\lambda)$, which is a union of geodesics in the metric on X_∞, must be nearly geodesic in the metric on X_n for n large. As $n \to \infty$, this union of finitely many almost-geodesic curves can be perturbed to be geodesic in X_n by a smaller and smaller homotopy.

By the strong convergence of X_n, the marked hyperbolic structures Σ_n converge in $\mathcal{T}(S)$ to a hyperbolic structure Σ_∞, which the pleating map $f_{\lambda,\infty}$ induces on S. (Compare [**18**, Theorem 5.2.2; **44**, Lemma 1.3].) Thus we may choose an ϵ-ball R_ϵ about $\Sigma_\infty \in \mathcal{T}(S)$ (say, in the Teichmüller metric) such that $\Sigma_n \in R_\epsilon$ for all $n \geqslant n_0$.

Since each pleating map $f_{\lambda,n} : S \to X_n$ is 1-Lipschitz with respect to the metric Σ_n, every closed geodesic on Σ_n tightens to a shorter geodesic in X_n. Thus we may apply Proposition 4.1 to every $\Sigma_n \in R_\epsilon$ and learn that there is a length cutoff L_0 such that for all $L \geqslant L_0$ and $n \geqslant n_0$,

$$\pi_{X_n}(L) \geqslant \pi_{\Sigma_n}(L) \geqslant \frac{e^L}{L}. \qquad \square$$

Remark 4.6. The reason why Proposition 4.5 assumes geometric finiteness is that geometrically infinite manifolds in $AH(S)$ actually contain an *infinite* set of closed geodesics shorter than some $L_0 = L_0(S)$.

Even for geometrically finite manifolds, the statement of Proposition 4.5 is surely an underestimate. Roblin [**52**] showed that a geometrically finite hyperbolic 3-manifold $M = H^3/\Gamma$ satisfies

$$\pi_M(L) \sim \frac{e^{hL}}{hL},$$

where $h = h(M)$ is the Hausdorff dimension of the limit set $\Lambda(\Gamma)$. By a theorem of Bowen [**11**], if M is not Fuchsian, then $h(M) > 1$. Thus $\pi_M(L)$ grows strictly faster than e^L/L.

However, because the rate of convergence of $\pi_M(L)$ to e^{hL}/hL may a priori depend on M, Roblin's result does not seem to imply the uniform statement that we need.

5. *Spectrally similar* 3-*manifolds*

This section is devoted to proving Theorem 1.2. To that end, Section 5.2 describes how to build the spectrally similar hyperbolic 3-manifolds N_n and N_n^μ. Each N_n comes from performing sufficiently long Dehn fillings on a manifold J_n that is constructed by gluing a pair of *caps*, denoted T and B, to the ends of a large product region $W_{\varphi,n}$. In Section 5.3, we describe a cut-and-paste operation known as *mutation*. Mutating N_n produces its partner N_n^μ, which shares much of the geometry of N_n. We will show that for $n \gg 0$, both N_n and N_n^μ have the properties that are claimed in Theorem 1.2.

5.1. *A general recipe*

We will use certain *caps* with the following properties.

Definition 5.1. A compact, orientable 3-manifold M is called

(1) *simple* if M does not contain any essential spheres, tori, disks, or annuli;
(2) *asymmetric* if every self-homeomorphism of M is isotopic to the identity.

We say that M is a *simple, asymmetric cap* for a surface S if $\partial M \cong S$ and both properties (1) and (2) hold for M.

By Thurston's hyperbolization theorem [**57**], a 3-manifold with boundary of genus at least 2 is simple if and only if it admits a hyperbolic metric with totally geodesic boundary.

We extend Definition 5.1 to pared 3-manifolds, as in Section 3.3, by designating a collection of curves $Q \subset \partial M$ as the paring locus. In this setting, (2) only prohibits symmetries that respect Q (and similarly for (1)). It follows that if M is asymmetric, then (M, Q) will be asymmetric.

We can now give a recipe for constructing N_n and N_n^μ. Let S be the closed surface of genus 2. We start with a pair of simple, asymmetric caps for S, denoted T and B (for *top* and *bottom*), which are required to be distinct up to homeomorphism. In addition, choose a pseudo-Anosov homeomorphism $\varphi : S \to S$ and a pants decomposition Q. Then, let

$$N_n = B \cup_{\tau_n \circ (\varphi^{2n}) \circ \tau_n'} T \quad \text{and} \quad N_n^\mu = B \cup_{\tau_n \circ (\varphi^{2n}) \circ \tau_n' \circ \mu} T. \tag{5.1}$$

Here, the subscript in (5.1) denotes the gluing map from ∂B to ∂T. Each of τ_n and τ_n' is a product of large Dehn twists about the curves of Q, where the notion of *large* grows with n. Meanwhile μ is the (unique) *hyper-elliptic involution* on S, which is central in $MCG(S)$. See Section 5.3 for more on hyper-elliptic involutions.

To emphasize the flexibility of the construction, we assert that essentially *any* sequence of manifolds constructed using the recipe (5.1) will satisfy the claims of Theorem 1.2, up to a linear factor depending on φ. For $n \gg 0$, the middle portion of N_n (between the curves being twisted) will become almost-isometric to the pared acylindrical manifold $W_{\varphi,n}$, which in turn converges to $\widetilde{M}_\varphi$ as $n \to \infty$. Thus all the results of Sections 3 and 4 will apply to the geometry of N_n and N_n^μ.

In Section 5.2, we give a very explicit construction of the caps T and B, along with slightly restricted choices for the pants decomposition Q and the gluing map φ. The point of the explicit choices is to make it easier to verify that T and B are indeed simple and asymmetric. In addition, we will drill out the curves of the pants decomposition in order to make the gluing of caps more rigid and to plug into the commensurability criterion of Theorem 1.7.

5.2. *An explicit construction*

Now, we give explicit choices for all the pieces used to build N_n. We begin with an explicit construction of T and B.

Let H be the handlebody of genus 2, and let $K \subset H$ be the knot depicted in Figure 4. We let T be the result of $5/3$ surgery on K, while B is the result of $7/4$ surgery on K. For general p/q surgery, we have the following result.

LEMMA 5.2. *Let H be the handlebody of genus 2, let $S = \partial H$, and let $K \subset H$ be the knot depicted in Figure 4. Then, for every slope p/q with $q > 2$, the 3-manifold $H(p/q)$ obtained by p/q surgery on K is a simple cap for S.*

In particular, T and B are simple caps for S.

Proof. We begin by checking that $H \smallsetminus K$ is a simple 3-manifold. This can be accomplished by doubling $H \smallsetminus K$ along $\partial H = S$ to obtain the complement of a 2-component link in $(S^2 \times S^1) \# (S^2 \times S^1)$. SnapPy verifies that this link complement is hyperbolic, and does not have any symmetries apart from the reflection along a totally geodesic copy of ∂H. By the easy direction of Thurston's hyperbolization, it follows that $H \smallsetminus K$ is simple. We remark that this argument also shows $H \setminus K$ is asymmetric.

Next, observe that $1/0$ surgery along K yields the handlebody H, which definitely contains essential disks. Then a combination of theorems by Gordon, Luecke, Scharlemann, and Wu

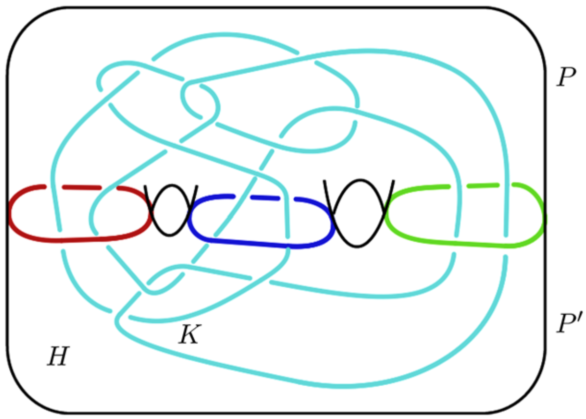

FIGURE 4 (colour online). *To construct the top cap T, start with the knot K in a genus 2 handlebody H. Then perform 5/3 Dehn surgery on K to get a simple, asymmetric 3-manifold. Then, add a paring locus along three simple closed curves on ∂H, as shown. The resulting pared manifold (T, Q_T) has totally geodesic boundary along $P \cup P'$.*

(see [**28**, Table 2.1]) says that any p/q surgery with $q > 2$ will produce a 3-manifold without any essential spheres, tori, disks, or annuli. Thus all these manifolds are simple. □

Next, we fix a pants decomposition $Q \subset \partial H \cong S$ consisting of the three darker curves in Figure 4. The construction of T via 5/3 surgery on $K \subset H$ identifies ∂T with ∂H, allowing us to designate Q to be the paring locus of ∂T. Similarly, the construction of B via 7/4 surgery on $K \subset H$ allows us to designate Q as the paring locus of ∂B. To avoid confusion on issues of marking, we refer to the resulting pared manifolds as (T, Q_T) and (B, Q_B). Since T and B are simple by Lemma 5.2, this construction makes (T, Q_T) and (B, Q_B) into pared, acylindrical 3-manifolds with totally geodesic boundary consisting of pairs of pants. Using SnapPy, we compute that

$$\mathrm{vol}(T, Q_T) = 36.4979\ldots \quad \text{and} \quad \mathrm{vol}(B, Q_B) = 36.5377\ldots. \tag{5.2}$$

LEMMA 5.3. *The pared manifolds (T, Q_T) and (B, Q_B) are simple, asymmetric caps for the pair (S, Q). Furthermore, these caps are not isometric.*

Proof. This follows from rigorous verification routines included in SnapPy [**23**]. The program can rigorously verify the canonical cell decomposition of a hyperbolic 3-manifold, which enables it to find all symmetries and certify that two 3-manifolds are not isometric.

To assist SnapPy in this endeavor, it is convenient to isometrically embed (T, Q_T) in a finite volume cusped 3-manifold, as follows. If the handlebody H is embedded in S^3 as shown in Figure 4, the pair of pants $P \subset \partial H$ becomes isotopic to $P' \subset \partial H$ through $S^3 \setminus H$. Thus we may glue P to P', realizing $K \cup Q$ as the 4-component link shown in Figure 4. Performing $(5, 3)$ surgery on $K \subset S^3$ results in a 3-cusped manifold, which becomes isometric to (T, Q_T) after cutting along a pair of pants to separate P from P'. An identical construction works for (B, Q_B). □

Now, we perform the following sequence of steps to build the 3-manifold N_n. We refer the reader to Figure 6 for a visual description of this process.

1. *Construct $W_{\varphi,n}$:* Let $Q \subset S$ be the pants decomposition specified above. Choose a pseudo-Anosov $\varphi : S \to S$. As in Sections 3.3 and 3.4, the triple (Q, φ, n) specifies a pared acylindrical

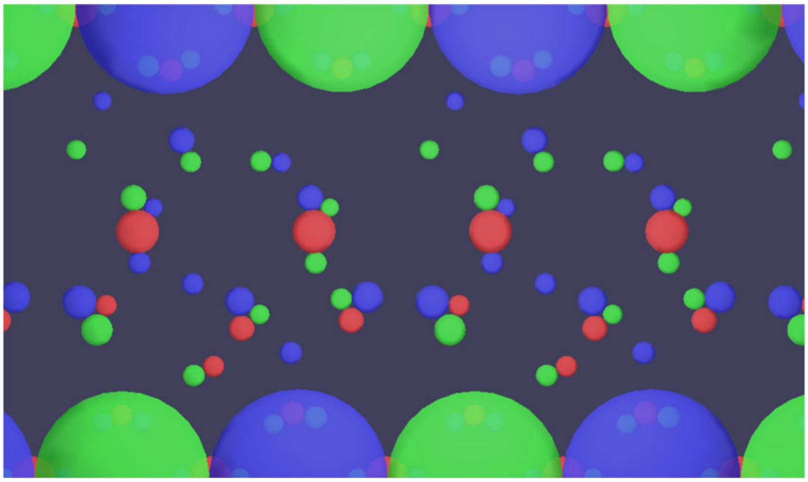

FIGURE 5 (colour online). *A horoball diagram for the top cap* (T, Q_T). *The view is from the red cusp in Figure 4. The distinguished line of pairwise tangent full-sized horoballs corresponding to* P *is at the top, while that of* P' *is at the bottom. Note that the full-sized horoballs along the top and bottom of the figure are not tangent to any other full-sized horoballs in the interior of the figure. In the language of Definition 2.2,* P *and* P' *are geometrically isolated on one side. The pattern of smaller horoballs shows that there are no symmetries interchanging* P *with* P'*; this is also confirmed with a rigorous computation.*

manifold $W_{\varphi,n} = M_S(\varphi^{-n}Q, \varphi^n Q)$. Recall that $M_S(\varphi^{-n}Q, \varphi^n Q)$ is equipped with a marking by S, which comes by including S as $S \times \{1/2\}$. In this marking, the bottom paring locus is $\varphi^{-n}Q$, and the top paring locus is $\varphi^n Q$.

As discussed in Remark 3.3, we choose a pseudo-Anosov φ that satisfies $4Et_{\mathcal{C}}(\varphi) > 1$. Here, E is the constant of (3.5), while $t_{\mathcal{C}}(\varphi)$ is the translation distance of φ in the curve complex $\mathcal{C}(S)$, as in (3.4). Our choice of φ means we have $B_\varphi > 1/2$ in Proposition 3.2.

2. *Construct* J_n*:* We build a finite-volume hyperbolic 3-manifold J_n as follows. Glue the lower boundary of $W_{\varphi,n}$ to the bottom cap (B, Q_B). This joins the rank-1 cusps of $\varphi^{-n}(Q) \subset \partial W_{\varphi,n}$ with the rank-1 cusps of $Q_B \subset \partial B$, turning them into rank-2 cusps with a torus cross-section. In a similar manner, glue the upper boundary of $W_{\varphi,n}$ to the top cap (T, Q_T), joining rank-1 cusps to form rank-2 cusps. All of the gluing occurs by isometry along rigid, totally geodesic pairs of pants.

This construction endows each cusp of J_n with a canonical *surface-framed longitude*. This is the slope along which an annulus parallel to Q_T is joined to an annulus parallel to $\varphi^n Q$ (and similarly for Q_B and $\varphi^{-n}Q$). We also choose a *meridian* to be a slope intersecting the longitude once.

3. *Construct* M_n*:* This 3-manifold is obtained from J_n by performing sufficiently long $1/k_i$ Dehn fillings along the three cusps that meet ∂B. The surgery coefficient $1/k_i$, with respect to the meridian-longitude framing chosen above, ensures that the filling slope intersects the surface-framed longitude exactly once. Equivalently, $1/k_i$ Dehn filling means that we are gluing the un-pared cap B to $S \times \{0\}$ via k_i Dehn twists along the i-th curve of $\varphi^{-n}Q$. The integers k_i must be sufficiently large so that the curves of $\varphi^{-n}Q$ become the three shortest geodesics in M_n.

By construction, M_n contains an isometrically embedded copy of (T, Q_T), whose boundary in M_n consists of totally geodesic pants P and P'. We call this submanifold $M_+ \subset M_n$, and note that the geometry of $M_+ \cong (T, Q_T)$ is independent of n. By contrast, the remainder of M_n is a submanifold M_- whose geometry changes with n.

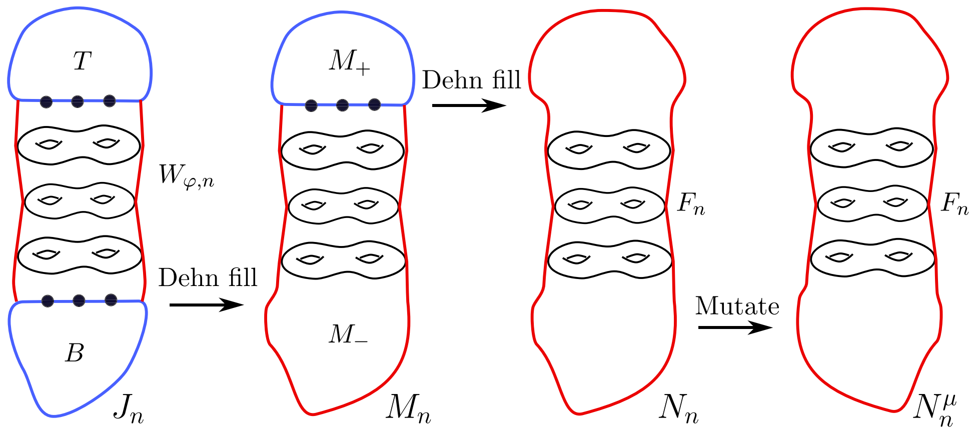

FIGURE 6 (colour online). *A schematic summary of the construction of N_n and N_n^μ. Start with J_n (left), which is built by gluing caps T and B to the ends of the product region $W_{\varphi,n}$. Next, Dehn fill J_n along the bottom 3 cusps to obtain M_n. Then, Dehn fill along either two or all three of the top cusps to obtain N_n. Finally, we can obtain N_n^μ by mutating along a genus two surface F_n. In our diagram, the bottom caps of N_n and N_n^μ differ by a reflection.*

4. *Construct N_n:* This 3-manifold is obtained from M_n by performing sufficiently long and sufficiently different $1/k'_i$ Dehn filling on either two or all three cusps of M_n. Here, the surgery coefficient has the same meaning as above: the filling slopes realize some number of Dehn twists along the curves of $\varphi^n Q$. The meaning of *sufficiently long* and *sufficiently different* comes from Theorem 1.7, and more specifically from equation (2.3): we need the curves of $\varphi^n Q$ to be the three shortest geodesics in N_n, and their lengths to have large ratios.

We could have combined the last two steps, by Dehn filling J_n to obtain N_n directly. However, building M_n (with its submanifolds M_+ and M_-) will allow us to apply Theorem 1.7 and show that N_n is minimal in its commensurability class.

Because of our choice of Dehn filling coefficients (which realize Dehn twists along S), N_n decomposes as the union of caps homeomorphic to B and T, connected via a product region homeomorphic to $S \times I$. This product region is glued to B along all of $S \times \{0\}$, and to T along either all of $S \times \{1\}$ or the complement of one closed curve in this surface. See Figure 6. We will continue to think of B and T as submanifolds of N_n. Note as well that this construction of N_n satisfies the recipe of (5.1).

Given J_n and N_n, we construct two auxiliary objects that will be useful in our arguments.

5. *Construct $\Sigma_\pm \subset N_n$:* Recall that $\partial(B, Q_B)$ consists of two pairs of pants, which are totally geodesic in J_n. Choose an orientation on the three curves of Q_B and an ideal triangulation of the complementary pairs of pants. Each endpoint of an ideal edge can be spun about a curve of Q_B, following the prescribed orientation of that curve. This gives a finite, filling lamination on S, which we call λ_-.

After J_n is filled to obtain N_n, Thurston's theorem [**56**, Proposition 9.7.1] says that we may pleat S along λ_- to obtain a pleated surface Σ_-. (Recall Definition 4.4.) For long Dehn fillings, the geometry of Σ_- closely approximates the geometry of $\partial(B, Q_B) \subset J_n$.

We perform the same construction on $\partial(T, Q_T)$ to get a finite, filling lamination λ_+. Once J_n is filled to obtain N_n, this gives a pleated surface $\Sigma_+ \subset N_n$. (If all three curves of Q_T are filled in N_n, the pleated surface Σ_+ will be homotopic to S; otherwise, it will be the complement of a curve in S.)

6. *Construct X_n:* Recall that $W_{\varphi,n} \in AH(S)$ is marked by an embedded copy of the surface S. This copy of S survives as an essential surface in J_n and then in N_n. (This is because our choice of $1/k_i$ Dehn filling is realized by Dehn twists along S. A compression disk for S in N_n

would have to contain a compression disk in either B or T. But B and T have no compression disks, by Lemma 5.2.) Furthermore, the topology of B and T prevents S from being a fiber or virtual fiber.

Let X_n be the cover of M_n corresponding to S. This manifold is marked by S, placing $X_n \in AH(S)$. By the work of Thurston and Bonahon [**7**], X_n is geometrically finite. Its convex core contains an isometric copy of Σ_+ close to $\partial_+ \mathrm{core} X_n$, and an isometric copy of Σ_- close to $\partial_- \mathrm{core} X_n$.

PROPOSITION 5.4. *For each n, there is a length cutoff z_n, such that if all slopes in the Dehn filling $J_n \to N_n$ are longer than z_n, the following will hold.*

(1) *There are positive constants A_φ and A'_φ, depending on φ, such that for all $n \gg 0$,*

$$A_\varphi \leqslant \frac{\mathrm{vol}(N_n)}{n} \leqslant A'_\varphi.$$

(2) *The cores of the Dehn filling solid tori are the shortest closed geodesics in N_n.*
(3) *The pleated surfaces $\Sigma_\pm \subset N_n$ satisfy $d(\Sigma_-, \Sigma_+) > n/2$.*
(4) *For an appropriate choice of baseframes, X_n converges strongly to $\widetilde{M}_\varphi$.*

Proof. For (1), recall that J_n was constructed by gluing three pieces by isometry along their boundaries. Thus, by (5.2),

$$\mathrm{vol}(J_n) = \mathrm{vol}(W_{\varphi,n}) + \mathrm{vol}(B, Q_B) + \mathrm{vol}(T, Q_T) = \mathrm{vol}(W_{\varphi,n}) + 73.035\ldots.$$

Let $s_1^n, \ldots, s_6^n$ be the Dehn filling slopes on the cusps of J_n that produce N_n. Each s_i^n has a normalized length $\hat{L}_i^n$, as in Definition 2.9. Neumann and Zagier showed that

$$\begin{aligned}\mathrm{vol}(N_n) &= \mathrm{vol}(J_n) - \pi^2 \left(\sum \left(\hat{L}_i^n\right)^{-2}\right) + O\left(\sum \left(\hat{L}_i^n\right)^{-4}\right)\\ &= \mathrm{vol}(W_{\varphi,n}) + O(1) - \pi^2 \left(\sum \left(\hat{L}_i^n\right)^{-2}\right) + O\left(\sum \left(\hat{L}_i^n\right)^{-4}\right).\end{aligned}$$

See [**43**, Theorem 1A]. Thus, if every normalized length $\hat{L}_i^n$ is sufficiently large, the difference $|\mathrm{vol}(N_n) - \mathrm{vol}(W_{\varphi,n})|$ will be uniformly bounded. Since $\mathrm{vol}(W_{\varphi,n})/n$ is bounded above and below in equation (3.2) of Proposition 3.2, it follows that $\mathrm{vol}(N_n)/n$ satisfies the same bounds with slightly modified values of the constants A_φ and A'_φ. This proves (1).

Conclusion (2) also follows from the work of Neumann and Zagier [**43**, Proposition 4.3]. This is because each normalized length $\hat{L}_i^n$ predicts the length of the corresponding closed curve in N_n, via equation (2.2). Thus choosing $\hat{L}_i^n$ sufficiently long forces the cores of the surgery solid tori to be very short.

Conclusion (3) follows from the K-bilipschitz theorem of Brock and Bromberg [**13**, Theorem 1.3]; see also Magid [**36**, Theorem 1.2]. They show that, for every bilipschitz constant $K > 1$, there is a length cutoff $z = z(K)$ depending only on K, such that if $\hat{L}_i > z(K)$ for all i, there will be a K-bilipschitz embedding $\psi_n : \mathrm{core}^0 W_{\varphi,n} \hookrightarrow N_n$. Furthermore, by [**13**, Theorem 6.10], a geodesic on $\partial W_{\varphi,n}$ is mapped by ψ_n to a curve whose geodesic curvature is vanishingly small. This is a strong, quantified version of geometric convergence.

Now, recall the finite, geodesic laminations $\lambda_\pm \subset \partial W_{\varphi,n}$. By the above paragraph, their images $\psi_n(\lambda_\pm)$ are nearly geodesic in N_n. Thus ψ_n maps the geodesic pairs of pants comprising $\partial_+ W_{\varphi,n}$ vanishingly close to Σ_+, and similarly for Σ_-.

In constructing $W_{\varphi,n}$, we chose a mapping class φ to have large translation distance in the curve complex $\mathcal{C}(S)$, ensuring that $B_\varphi > 1/2$ in Proposition 3.2. Thus, by Proposition 3.2, for $n \gg 0$ we have

$$d(\partial_- \mathrm{core}^0 W_{\varphi,n}, \partial_+ \mathrm{core}^0 W_{\varphi,n}) > n/2.$$

(In this section, $d_0(\cdot,\cdot)$ is the same as ordinary distance because S is a closed surface.) By choosing a bilipschitz constant $K = K_n$ sufficiently close to 1, and taking $z_n = z(K_n)$, we ensure that this lower bound is preserved in N_n, hence $d(\Sigma_-, \Sigma_+) > n/2$.

For (4), recall from Proposition 3.1 that there is a choice of basepoints $x_n \in W_{\varphi,n}$, with baseframes ω_n at x_n, such that we have a geometric limit $(W_{\varphi,n}, \omega_n) \to (\widetilde{M}_\varphi, \omega_\infty)$. The distance from x_n to the totally geodesic boundary of $W_{\varphi,n}$ must go to ∞ with n, for otherwise, this boundary would appear in the geometric limit. So long as $d(x_n, \partial W_{\varphi,n}) > R$, an R-ball about $x_n \in J_n$ is contained entirely in $W_{\varphi,n}$. Thus, by the Gromov–Hausdorff characterization of geometric limits in (3.1), we have a geometric limit $(J_n, \omega_n) \to (\widetilde{M}_\varphi, \omega_\infty)$. Similarly, for very long Dehn fillings, an R-ball about $x_n \in J_n$ is nearly isometric to a metric R-ball about $\psi_n(x_n) \in N_n$, where $\psi_n : \mathrm{core}^0\, W_{\varphi,n} \hookrightarrow N_n$ is the bilipschitz diffeomorphism described above. This will enable us to conclude that N_n converges geometrically to $\widetilde{M}_\varphi$. Furthermore, since all of the topology of this R-ball is captured by $\pi_1(X_n)$, we will in fact have a geometric and algebraic limit $X_n \to \widetilde{M}_\varphi$.

To make this summary precise, recall that the Chabauty topology is given by a metric d_{Chab} [**5**]. Thus we may write

$$d_{\mathrm{Chab}}((J_n, \omega_n),\, (\widetilde{M}_\varphi, \omega_\infty)) = \epsilon_n \to 0.$$

By Thurston's Dehn surgery theorem (see, for example, [**4**, Theorem E.5.1]), any sequence of Dehn fillings of J_n in which the lengths of the filling slopes approach ∞ will converge to J_n in the Chabauty topology. Let $\omega_n \in W_{\varphi,n} \subset J_n$ be the same set of baseframes as before. Then there is a length cutoff z_n such that if $\widehat{L}_i^n \geqslant z_n$ for all i, we will have baseframes ν_n on N_n such that

$$d_{\mathrm{Chab}}((J_n, \omega_n),\, (N_n, \nu_n)) \leqslant \epsilon_n.$$

Combining the last two equations gives

$$d_{\mathrm{Chab}}((N_n, \nu_n),\, (\widetilde{M}_\varphi, \omega_\infty)) \leqslant 2\epsilon_n \to 0,$$

hence there is a geometric limit $(N_n, \nu_n) \to (\widetilde{M}_\varphi, \omega_\infty)$, as claimed.

For every $R > 0$, we know that there is an integer $n(R)$ such that $d(x_n, \partial W_{\varphi,n}) > 2R$ for all $n \geqslant n(R)$. By the bilipschitz theorem of Brock and Bromberg [**13**, Theorem 1.3], and the above argument, we have $d(\psi_n(x_n), \Sigma_\pm) > R$. Since the entire region between Σ_- and Σ_+ has an isometric lift to $\mathrm{core}(X_n)$, we know that the metric R-ball about $\psi_n(x_n)$ also lifts isometrically to $\mathrm{core}(X_n)$ for $n \geqslant n(R)$. Since this works for every $R > 0$, we have a geometric limit $(X_n, \nu_n) \to (\widetilde{M}_\varphi, \omega_\infty)$.

Finally, observe that each of X_n, $W_{\varphi,n}$, and $\widetilde{M}_\varphi$ is marked by an inclusion of S. For long Dehn fillings, the representations $\pi_1(S) \to \pi_1(X_n) \subset \pi_1(N_n)$ converge algebraically in $AH(S)$ to the limiting representation given by the marking of $W_{\varphi,n}$. Since $W_{\varphi,n}$ converges strongly to $\widetilde{M}_\varphi$, we have a strong limit $(X_n, \nu_n) \to (\widetilde{M}_\varphi, \omega_\infty)$. □

5.3. *Mutant partners*

We now describe a procedure, called *mutation*, that will preserve much of the geometry of N_n while changing its commensurability class.

Given a surface S, a *hyper-elliptic involution* is a self-homeomorphism $\mu : S \to S$ that preserves the isotopy class of every simple closed curve. A hyperbolizable surface S admits hyper-elliptic involutions if and only if $\chi(S) = -1$ or -2. Since a marked hyperbolic structure on S is determined by the lengths of finitely many closed curves, and μ preserves all of these lengths, it acts trivially on $\mathcal{T}(S)$, the compactification $\overline{\mathcal{T}(S)}$, and the product $\overline{\mathcal{T}(S)} \times \overline{\mathcal{T}(S)}$. Consequently, hyper-elliptic involutions of S act trivially on $AH(S)$.

Definition 5.5. Let M be a 3-manifold and $F \subset M$ an incompressible, boundary-incompressible surface admitting a hyper-elliptic involution μ. We may cut M along F and reglue via μ. This creates a 3-manifold M^μ, called the *mutant partner* of M. The process is called *mutation along* F.

We can now describe how to create mutant partners for our manifolds. First, let $F_n \subset W_{\varphi,n} \subset J_n$ denote the copy of S marking $W_{\varphi,n}$, that is, $F_n = S \times \{\frac{1}{2}\} \subset S \times I$. Since S is the closed surface of genus 2, there is a unique hyper-elliptic element $\mu \in MCG(S)$. We obtain J_n^μ from J_n by mutating along F_n. Note that the hyper-elliptic involution μ extends over $W_{\varphi,n} \in AH(S)$. Thus J_n^μ is built from the same pieces as J_n, with a modified gluing map. To build J_n^μ, we glue $W_{\varphi,n}$ to (T, Q_T) by the same map as before, and glue it to (B, Q_B) by the previous map composed with μ.

By definition, the hyper-elliptic involution μ preserves the curves of $\varphi^{\pm n}Q$, which correspond to the cusps of J_n. We construct M_n^μ and N_n^μ by Dehn filling J_n^μ along exactly the same filling slopes used to produce M_n and then N_n. After the Dehn filling, N_n still contains an embedded, incompressible copy of F_n. (See Figure 6.) Thus we could also obtain N_n^μ by first filling J_n to obtain N_n, and then mutating along F_n to obtain N_n^μ. Because the filling coefficients are unchanged, the operations of Dehn filling and mutation commute.

The covers of N_n and N_n^μ corresponding to F_n are both isometric to X_n. In fact, the hyper-elliptic involution $\mu : S \to S$ extends to $X_n \cong S \times \mathbb{R} \in AH(S)$. This can be rephrased to say that although N_n and N_n^μ are incommensurable (as we will see in Lemma 5.8), they do share a common infinite-sheeted cover, which is quasifuchsian if N_n is closed.

Before checking the properties claimed in Theorem 1.2, we establish the following fact.

Lemma 5.6. *When $n \gg 0$, the only essential pairs of pants in J_n are isotopic into $\partial W_{\varphi,n}$. The only essential pairs of pants in M_n are isotopic into $\partial M_+ = \partial M_-$. Furthermore, the same statements hold in J_n^μ and M_n^μ.*

Proof. First, we claim that the only essential pairs of pants in (T, Q_T) or (B, Q_B) are isotopic to the boundary. This is verified using Regina [**16**], by enumerating all orientable normal surfaces with Euler characteristic -1 in these manifolds.

Next, consider an essential pair of pants $P \subset W_{\varphi,n}$, with boundary mapped to the parabolic locus. If boundary components of P are mapped to both $\varphi^{-n}Q$ and $\varphi^n Q$, the ϵ-thick part of P (where ϵ is the Margulis constant) will necessarily contain a path from $\partial_-\mathrm{core}^0\, W_{\varphi,n}$ to $\partial_+\mathrm{core}^0\, W_{\varphi,n}$. But the ϵ-thick part of P has bounded diameter (compare Lemma 2.7), whereas the distance between the upper and lower boundary of $\mathrm{core}^0\, W_{\varphi,n}$ must grow linearly with n by Proposition 3.2. For $n \gg 0$, this is a contradiction. Thus, for large n, all three boundary components of P are mapped to curves of $\varphi^n Q$, or all three to $\varphi^{-n}Q$. In either case, P is isotopic to one of the pairs of pants comprising $\partial W_{\varphi,n}$.

Now, consider an essential pair of pants in P in J_n or J_n^μ. The intersection $P \cap \partial W_{\varphi,n}$ must consist of simple closed curves that are essential in P. But the only essential simple curves in a pair of pants are peripheral, hence P can be isotoped to be disjoint from $\partial W_{\varphi,n}$. After this isotopy, P is entirely contained in one of the pieces T, B, or $W_{\varphi,n}$. But we have already checked that any pair of pants in these submanifolds is isotopic into $\partial W_{\varphi,n}$.

Finally, consider an essential pair of pants in $P \subset M_n$. By the same intersection argument as above, P can be isotoped to be disjoint from $\partial M_+ = \partial M_-$. If $P \subset M_+ = (T, Q_T)$, then we have already checked that P is isotopic into ∂M_+. If $P \subset M_-$, observe again that $P \cap \mathrm{core}^0 M_n$ has bounded diameter. Thus all of P is contained in a collar neighborhood of ∂M_-, homeomorphic to $S \times I$. Now P must be isotopic into ∂M_-, by the same argument as for $W_{\varphi,n}$. □

LEMMA 5.7. *For $n \gg 0$, the mutant partners N_n and N_n^μ are non-isometric hyperbolic 3-manifolds. Furthermore,* $\mathrm{vol}(N_n) = \mathrm{vol}(N_n^\mu)$.

Proof. Suppose, for a contradiction, that there is an isometry $\psi : N_n \to N_n^\mu$. By Proposition 5.4, the shortest closed geodesics in N_n are the cores of the (five or six) solid tori that were added when we filled J_n. The isometry ψ must respect these shortest geodesics, hence ψ restricts to an isometry $\psi : J_n \to J_n^\mu$.

Let $P \subset \partial W_{\varphi,n}$ be a pair of pants. Then $\psi(P)$ is a totally geodesic pair of pants in J_n^μ. By Lemma 5.6, all pairs of pants in J_n^μ occur along $\partial W_{\varphi,n}$. Thus ψ must respect the decomposition of J_n^μ into top and bottom caps, joined along the product region $W_{\varphi,n}$. In particular, $\psi(T) \subset J_n^\mu$ must be a cap isometric to (T, Q_T). Since the top and bottom caps are not isometric by Lemma 5.3, it follows that $\psi(T)$ is the top cap of J_n^μ, which we may identify with (T, Q_T) in a unique way.

By Lemma 5.3, the isometry $\psi|_T$ must be the identity map. By restricting to the boundary, $\psi|_{\partial_+ W_{\varphi,n}}$ is also the identity map. Since we have performed a mutation along $F_n \subset W_{\varphi,n}$, it follows that $\psi|_{\partial_- W_{\varphi,n}}$ is the hyper-elliptic involution μ. But, by Lemma 5.3, μ cannot extend over the bottom cap (B, Q_B), giving a contradiction, Thus N_n and N_n^μ are not isometric.

Finally, $\mathrm{vol}(N_n) = \mathrm{vol}(N_n^\mu)$ by a theorem of Ruberman [**53**]. □

We remark that the argument of Lemma 5.7 also shows that J_n and N_n have no symmetries. For, any isometry $\psi : J_n \to J_n$ must respect the decomposition of J_n into top and bottom caps, joined along $W_{\varphi,n}$. But then the restriction $\psi|_T$ must be the identity map, hence ψ is the identity. It follows that N_n cannot be a regular cover of any orbifold (other than itself). In fact, we can show more.

LEMMA 5.8. *For $n \gg 0$, each of N_n and N_n^μ is non-arithmetic and minimal in its commensurability class. In particular, N_n and N_n^μ are incommensurable.*

Proof. This will follow as a consequence of Theorem 1.7. We need to check the hypotheses of that theorem.

By construction, M_n decomposes into submanifolds M_+ and M_-, where $M_+ \cong (T, Q_T)$. These submanifolds are glued along $P \cup P'$, which are the only pairs of pants in M_n by Lemma 5.6. The volume of M_+ is independent of n (it is given in (5.2)), whereas $\mathrm{vol}(M_n) > \mathrm{vol}(N_n)$ grows linearly with n by Proposition 5.4. Thus $n \gg 0$ implies $\mathrm{vol}(M_-) \gg \mathrm{vol}(M_+)$. Furthermore, M_+ is asymmetric by Lemma 5.3.

Next, we check that there is a choice of cusp neighborhoods $\{C_1, C_2, C_3\} \subset M_n$ such that the pairs of pants P and P' are pairwise tangent and geometrically isolated on one side (see Definition 2.2). We choose a transverse orientation on both P and P', pointing away from M_- and toward M_+. Then, we verify using SnapPy that in the horoball diagram of M_+, the full-sized horoballs of $(P \cup P') \cap C_i$ are not tangent to any other full-sized horoballs in the direction of M_+. See Figure 5, and recall that the isometry class of $M_+ \cong (T, Q_T)$ does not depend on n.

Thus, by Theorem 1.7, our choice of sufficiently long and sufficiently different Dehn filling slopes in the definition of N_n ensures that N_n is the unique minimal orbifold in its (non-arithmetic) commensurability class. The same argument applies to M_n^μ and N_n^μ, because M_n^μ can be assembled from the same pieces M_+ and M_- by modifying the gluing map by μ. Thus N_n^μ is also non-arithmetic and minimal. By Lemma 5.7, N_n and N_n^μ are not isometric, hence they are incommensurable. □

5.4. *Geodesics under mutation*

To complete the proof of Theorem 1.2, it remains to check that N_n and N_n^μ share the same set of closed geodesics up to length n, and that there are at least e^n/n such geodesics. We verify this in the following two lemmas.

LEMMA 5.9. *For $n \gg 0$, any closed geodesic $\gamma \subset N_n$ whose length is at most n can be homotoped to be disjoint from the mutation surface F_n. Consequently, there is a bijection between the complex length spectra of N_n and N_n^μ up to length n.*

Proof. Recall the pleated surfaces $\Sigma_\pm \subset N_n$. Since $d(\Sigma_-, \Sigma_+) > n/2$ by Proposition 5.4, any closed geodesic $\gamma \subset N_n$ of length at most n must be disjoint from one of these surfaces. If γ is disjoint from Σ_-, we recall that Σ_- is homotopic to F_n, hence γ can be homotoped to be disjoint from F_n. If γ is disjoint from Σ_+, then it is homotopic into B, which is disjoint from F_n. (The asymmetry between Σ_- and Σ_+ only arises if we leave one cusp of J_n unfilled, and Σ_+ meets that cusp.)

Since all geodesics of length at most n can be homotoped to be disjoint from F_n, the conclusion about length spectra follows by a theorem of Millichap [**41**, Proposition 4.4]. □

LEMMA 5.10. *For $n \gg 0$, each of N_n and N_n^μ contains at least e^n/n geodesics up to length n. That is,*

$$\pi_{N_n}(n) = \pi_{N_n^\mu}(n) \geqslant \frac{e^n}{n}.$$

Proof. We begin with the hyperbolic manifold X_n, which is a covering space of both N_n and N_n^μ corresponding to F_n. By Proposition 5.4, there is a choice of baseframes for which X_n converges strongly to $\widetilde{M_\varphi}$. Applying Proposition 4.5 with $L = n$, we see that for $n \gg 0$,

$$\pi_{X_n}(n) \geqslant \frac{e^n}{n}. \tag{5.3}$$

Every closed geodesic $\gamma \subset X_n$ projects to a closed geodesic in N_n. To complete the proof, we need to check that distinct geodesics in X_n – that is, distinct free homotopy classes of loops in X_n – project to distinct free homotopy classes in N_n.

Suppose, for a contradiction, that γ and γ' are geodesic loops in X_n that are not freely homotopic in X_n, but become freely homotopic in N_n. We may suppose after a homotopy that γ and γ' lie in the marking surface S, which is embedded in N_n. In N_n, the free homotopy from γ to γ' is realized by a (singular) annulus A, which must have an essential intersection with T or B. By Jaco's annulus theorem [**32**, Theorem VIII.13], T or B must contain an essential embedded annulus A', which contradicts Lemma 5.2. Thus no annulus can exist, hence γ and γ' represent distinct geodesics in N_n.

By the same argument, distinct geodesics in X_n project to distinct geodesics in N_n^μ. In group-theoretic terms, we have checked that $\pi_1(X_n)$ is malnormal in both $\pi_1(N_n)$ and $\pi_1(N_n^\mu)$. Combining this result with (5.3) and Lemma 5.9 gives

$$\pi_{N_n}(n) = \pi_{N_n^\mu}(n) \geqslant \pi_{X_n}(n) \geqslant \frac{e^n}{n}.$$ □

Proof of Theorem 1.2. In the last several lemmas, we have checked that N_n and N_n^μ satisfy all the criteria required by Theorem 1.2. To recap: Conclusion (1) of the theorem holds by Proposition 5.4 and Lemma 5.7. Conclusion (2) holds by Lemma 5.9. Conclusion (3) holds by Lemma 5.10. Finally, conclusion (4) holds by Lemma 5.8. □

6. *Spectrally Similar Knots*

In this section, we will describe how to construct pairs of spectrally similar, incommensurable, hyperbolic knot complements that differ by mutation. In particular, we will outline the steps required to prove the following theorem.

THEOREM 1.5. *For each $n \gg 0$, there exists a pair of non-isometric mutant hyperbolic knot exteriors $E_n = S^3 \smallsetminus K_n$ and $E_n^\mu = S^3 \smallsetminus K_n^\mu$ such that:*

(1) $\mathrm{vol}(E_n) = \mathrm{vol}(E_n^\mu)$*, where this volume grows coarsely linearly with n;*
(2) *the (complex) length spectra of E_n and E_n^μ agree up to length at least $2\log(n)$;*
(3) *E_n and E_n^μ have at least $n^2/(2\log(n))$ closed geodesics up to length $2\log(n)$;*
(4) *E_n is the unique minimal orbifold in its commensurability class, and the only knot complement in its commensurability class. The same statement holds for E_n^μ.*

The construction behind Theorem 1.5 is extremely similar in spirit to the construction in Section 5 that proves Theorem 1.2. We will still use the recipe (5.1), except this time the surface S will be a 4-holed sphere. The knot complement $E_n = S^3 \smallsetminus K_n$ will have two caps, denoted T and B, connected by a product region $S \times I$ whose diameter grows linearly with n. In this setting, the caps will be tangles with boundary S, and the product region corresponds to a large power φ^{2n} of a pseudo-Anosov braid φ. Just as in Section 5, $E_n^\mu = S^3 \smallsetminus K_n^\mu$ will differ from E_n by mutation. See Figure 7 for a visual summary.

Because of the high degree of similarity to Section 5, most of this section will be a sketch. We will focus more attention on the few places where the argument differs.

6.1. *Our construction*

For the rest of the section, fix a surface $S = S_{0,4}$. We begin by describing the analogue of caps for S.

DEFINITION 6.1. A *tangle* is a pair (D, L), where D is the 3-ball, and $L \subset D$ is a properly embedded 1-manifold such that $\partial L \subset \partial D$ consists of exactly 4 points. We identify $\partial D \smallsetminus \partial L$ with $S_{0,4}$. The tangle (D, L) is called *rational* if L consists of two arcs that are simultaneously boundary-parallel in D.

Conway defined the operations of addition and multiplication on tangles, both of which involve joining a pair of tangles along half of their boundaries. A tangle is called *arborescent* if it can be obtained from rational tangles via these operations. See, for example, Wu [**59**] for a clear summary of these operations. The relevant fact in our setting is that the tangles B and T depicted in Figure 7 are arborescent but not rational.

We now perform the following sequence of constructions.

1. *Construct T and B:* Let $T = (D_T, L_T)$ and $B = (D_B, L_B)$ be exactly the tangles depicted in Figure 7. In particular, each of T and B is an arborescent tangle built out of 3 rational tangles. Note that in Figure 7, the four points of ∂L_T are directly above the 4 points of ∂L_B. We obtain a model copy of $S = S_{0,4} \subset S^3 \smallsetminus (B \cup T)$ by taking a horizontal plane in $\mathbb{R}^3$ between B and T, puncturing it at the 4 points below ∂L_T, and compactifying via the point at ∞.

2. *Construct Q, τ, and φ:* Let $Q \subset S$ be the simple closed curve encircling the two left-most punctures in the model copy of $S_{0,4}$. Translating this model surface up and down in S^3 gives an unknotted circle $Q_T \subset \partial T$ and an unknotted circle $Q_B \subset \partial B$, as shown. Let τ be a (left-handed) Dehn twist along $Q \subset S$, or equivalently a left-handed full twist about the left-most pair of strands passing through the model copy of S.

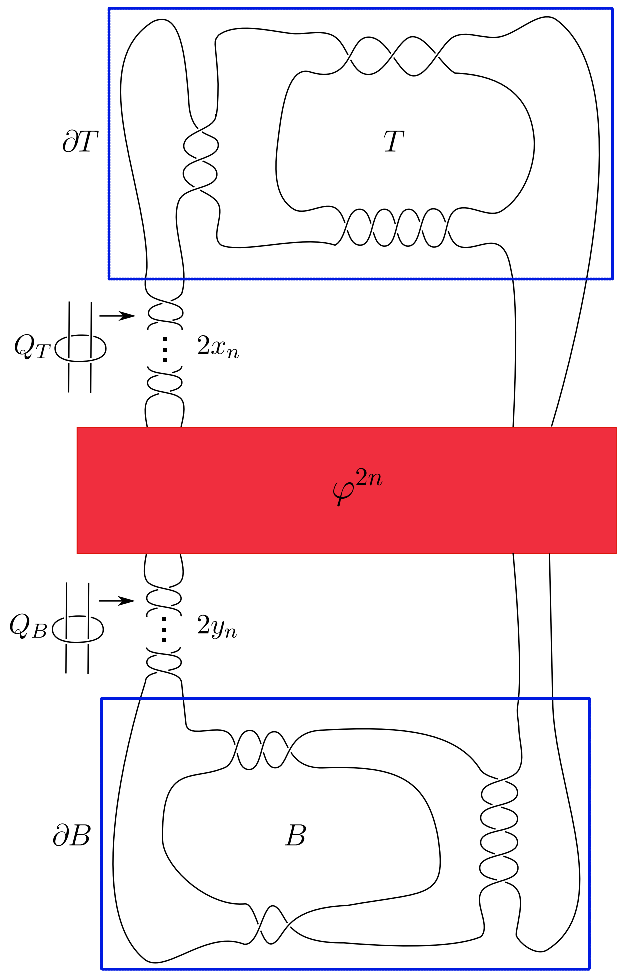

FIGURE 7 (colour online). *A schematic diagram of the arborescent knot K_n. The tangles T and B are exactly as shown. The twist regions below T and above B have $2x_n$ and $2y_n$ crossings, respectively. The braid $\tau^{x_n} \circ \varphi^{2n} \circ \tau^{y_n}$ that connects B to T is required to be alternating. We obtain L_n by replacing the twist regions τ^{x_n} and τ^{y_n} by unknotted circles, Q_T and Q_B, as shown. Then K_n can be recovered by Dehn filling along those circles.*

Let φ be a pure braid on four strands in S^2, chosen to satisfy the following properties. First, choose φ so that $\tau \circ \varphi \circ \tau$ is an alternating 4-braid. In addition, choose φ so that $4Et_{\mathcal{C}}(\varphi) > 1$. Here, E is the constant of (3.5), while $t_{\mathcal{C}}(\varphi)$ is the translation distance of φ in the curve complex $\mathcal{C}(S_{0,4})$, as in (3.4). As discussed in Remark 3.3, this criterion can always be achieved by taking an (alternating, pure) pseudo-Anosov φ_0 and letting φ be some power of φ_0.

3. *Construct K_n and $E_n = S^3 \smallsetminus K_n$:* For every $n \geqslant 1$, we build a knot $K_n \subset S^3$ as follows. Choose positive integers x_n and y_n, so that $x_n \gg y_n \gg 0$. Then, join the 1-manifold $L_B \subset D_B$ to $L_T \subset D_T$ via the alternating pure braid $\tau^{x_n} \circ \varphi^{2n} \circ \tau^{y_n}$. (If the braid word is read right to left, as a function, the crossings should be read bottom to top.)

The resulting link K_n, depicted in Figure 7, will be a knot because the pure braid $\tau^{x_n} \circ \varphi^{2n} \circ \tau^{y_n}$ connects the punctures of S in exactly the same way as the empty braid. By construction, K_n is alternating and arborescent. In addition, $E_n = S^3 \smallsetminus K_n$ satisfies the recipe (5.1).

4. *Construct L_n, J_n and $W_{\varphi,n}$:* Consider the 3-component link $L'_n = K_n \cup Q_B \cup Q_T$, and let $J_n = S^3 \setminus L'_n$. Equivalently, $J_n \cong S^3 \setminus L_n$, where L_n is the link obtained from L'_n by removing the full twists τ^{x_n} and τ^{y_n}.

By construction, L_n is an augmented alternating link, hence $J_n = S^3 \setminus L_n$ is hyperbolic by a theorem of Adams [**2**]. In the hyperbolic metric on J_n, the surface $\partial T \setminus Q_T$ is the union of two totally geodesic pairs of pants, and similarly for $\partial B \setminus Q_B$. These totally geodesic pairs of pants decompose J_n into caps corresponding to pared tangles $(T, L_T \cup Q_T)$ and $(B, L_B \cup Q_B)$, along with a product region $W_{\varphi,n}$ between the caps. By construction, the paring locus of $\partial_+ W_{\varphi,n}$ occurs along the punctures of S and $Q_T = \varphi^n Q$, and similarly for $\partial_- W_{\varphi,n}$.

We may recover K_n from $J_n = S^3 \setminus L_n$, by performing $1/x_n$ surgery along Q_T and $1/y_n$ surgery along Q_B.

5. *Construct K_n^μ and $E_n^\mu = S^3 \setminus K_n^\mu$:* Let $F_n \subset E_n$ be a model copy of S, placed between the caps T and B. Let $\mu : S \to S$ be the hyper-elliptic involution corresponding to π-rotation about a horizontal line in S. Let K_n^μ be the knot obtained from K_n by mutation along F_n. In Figure 7, K_n^μ can be obtained by removing the tangle B, rotating it by π about a horizontal axis (thus interchanging the two horizontal twist regions), and gluing it back in.

6. *Construct X_n:* Let X_n be the cover of E_n corresponding to $\pi_1(F_n)$. As in Section 5.3, the hyper-elliptic involution μ extends over X_n, hence the cover of E_n^μ corresponding to $\pi_1(F_n)$ is isometric to X_n.

6.2. *Proving Theorem* 1.5

It remains to check that $E_n = S^3 \setminus K_n$ and $E_n^\mu = S^3 \setminus K_n^\mu$ have all the properties claimed in Theorem 1.5. Most of the steps follow the same outline as Section 5. The one main difference is in Lemma 6.6, where we must pay attention to geodesics that cut through the cusp of E_n.

LEMMA 6.2. *K_n and K_n^μ are hyperbolic, arborescent, alternating knots. Neither knot complement is arithmetic. Furthermore, neither knot has any lens space surgeries.*

Proof. K_n and K_n^μ are arborescent by construction, because we have connected two arborescent tangles by a 4-string braid. The knots are alternating by construction, because we have connected two alternating tangles via a braid that alternates in a consistent direction. By the work of Bonahon and Siebenmann [**8**], any non-hyperbolic arborescent knot is Montesinos (meaning, it is a cyclic sum of rational tangles, which K_n and K_n^μ are not). See also Futer–Guéritaud [**24**] and Wu [**59**]. Thus E_n and E_n^μ are hyperbolic.

A theorem of Wu [**59**] also implies that K_n and K_n^μ have no exceptional surgeries. By a theorem of Reid [**48**], the only knot in S^3 with an arithmetic complement is the figure-8 knot (which is Montesinos). Thus E_n and E_n^μ are non-arithmetic. □

LEMMA 6.3. $\mathrm{vol}(E_n) = \mathrm{vol}(E_n^\mu)$*, and this volume grows coarsely linearly with n.*

Proof. Since K_n and K_n^μ are mutant knots, their complements E_n and E_n^μ have equal volume by Ruberman's theorem [**53**].

By the same Dehn filling argument as in Proposition 5.4, $\mathrm{vol}(E_n) = \mathrm{vol}(W_{\varphi,n}) + O(1)$, and $\mathrm{vol}(W_{\varphi,n})$ grows linearly with n by Proposition 3.2. We remark that in the present context, where $S = S_{0,4}$, one could obtain explicit upper and lower bounds on $\mathrm{vol}(E_n)$ from the work of Guéritaud and Futer [**29**, Appendix B]. □

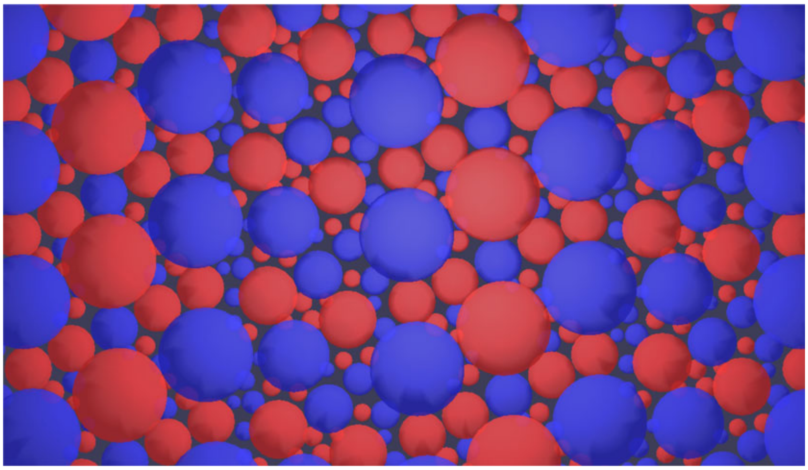

FIGURE 8 (colour online). *A horoball diagram for the pared version of the top tangle* $(T, L_T \cup Q_T)$. *The view is from the cusp of* K_n. *The blue (darker) horoballs correspond to arcs from* K_n *to* Q_T. *The totally geodesic pairs of pants* P *and* P' *lift to nearly vertical lines in this figure. The pattern of horoballs cannot be invariant under a rotation of order* 3 *or* 4, *implying that* $M_n = E_n \smallsetminus Q_T$ *cannot cover an orbifold with a rigid cusp.*

LEMMA 6.4. *The pared tangles* $(T, L_T \cup Q_T)$ *and* $(B, L_B \cup Q_B)$ *are simple, distinct, and asymmetric. It follows that for* $n \gg 0$, *the knot complements* $E_n = S^3 \smallsetminus K_n$ *and* $E_n^\mu = S^3 \smallsetminus K_n^\mu$ *are distinct and asymmetric.*

Proof. The conclusion about the tangles is checked using SnapPy, as in Lemma 5.3. Note that doubling $T \smallsetminus Q_T$ along its geodesic boundary produces a link $(L_T \cup \overline{L_T} \cup Q_T) \subset S^3$, where $\overline{L_T}$ is the mirror image of L_T. This link can be analyzed using SnapPy, verifying that it has no symmetries apart from reflection in the totally geodesic surface $\partial T \smallsetminus Q_T$. Regina checks that the only essential pants in $(T, L_T \cup Q_T)$ occur along the boundary. The same process works for the bottom tangle B.

Now, an argument using pairs of pants (just as in Lemma 5.7) implies that J_n and J_n^μ are non-isometric. Similarly, any self-isometry of J_n would have to preserve its decomposition into B, T, and $W_{\varphi,n}$, implying that the map is the identity. Because we have chosen very long Dehn filling slopes along Q_T and Q_B, any homeomorphism $E_n \to E_n^\mu$ or $E_n \to E_n$ would restrict to a homeomorphism (hence an isometry) on J_n, a contradiction. □

We remark that the distinctness and asymmetry of K_n and K_n^μ can also be checked using knot-theoretic tools. See the classification of arborescent knots by Bonahon and Siebenmann [**8**, Theorems 12.12 and 16.4], or the classification of alternating knots by Menasco and Thistlethwaite [**40**]. This approach does not need the hypothesis that $n \gg 0$.

LEMMA 6.5. *For* $n \gg 0$, *and an appropriate choice of* $x_n \gg y_n \gg 0$, *the following holds. Each of* E_n *and* E_n^μ *is the unique minimal orbifold in its commensurability class, and the only knot complement in its commensurability class.*

Proof. By Lemma 6.2, E_n is non-arithmetic, hence Theorem 1.3 says that it covers a minimal orbifold $\mathcal{Q}_n = \mathbb{H}^3/C^+(E_n)$. Suppose that the cover $\psi : E_n \to \mathcal{Q}_n$ is non-trivial. Then ψ must be an irregular cover, because E_n has no symmetries by Lemma 6.4. In other words, ψ is given by a so-called *hidden symmetry*. By a theorem of Neumann and Reid [**42**, Proposition 9.1], the cusp neighborhood of $\mathcal{Q}_n$ is *rigid*, meaning that its cross-section is a rigid 2-orbifold (recall Definition 2.5).

As in the proof of Theorem 1.7, we choose the filling coefficients $x_n \gg y_n \gg 0$ sufficiently long and sufficiently distinct to ensure that each of the core curves created by Dehn filling Q_T and Q_B is the complete preimage of its image in $\mathcal{Q}_n$. Thus the irregular covering map $\psi : E_n \to \mathcal{Q}_n$ restricts to a covering map $\psi : M_n \to \mathcal{O}_n$, where $M_n = E_n \setminus Q_T$. Each of the two cusps of M_n covers a distinct cusp of $\mathcal{O}_n$. Furthermore, by the above paragraph, the cusp of M_n corresponding to K_n must cover a rigid cusp of $\mathcal{O}_n$.

Observe that M_n contains an isometric copy of the pared tangle $(T, L_T \cup Q_T)$, with boundary along two pairs of pants, P and P'. We cannot apply Theorem 1.7 directly, because P and P' are not pairwise tangent (see Figure 8). Nevertheless, that figure shows that the horoball diagram for the top tangle $(T, L_T \cup Q_T)$ cannot be invariant under a rotation by angle $2\pi/3$ or $\pi/2$. On the other hand, every rigid 2-orbifold contains a cone point of order 3 or 4 (see Definition 2.5 and Lemma 2.6). Thus the quotient orbifold $\mathcal{O}_n$ cannot have a rigid cusp, which is a contradiction.

This contradiction implies that the cover $E_n \to \mathcal{Q}_n$ is trivial, hence E_n has no symmetries or hidden symmetries, and is minimal in its commensurability class. In addition, by Lemma 6.2, E_n has no lens space surgeries. Now, a theorem of Reid and Walsh [**51**, Proposition 5.1] implies that E_n is the only knot complement in its commensurability class. The same argument applies to E_n^μ. □

Lemma 6.6. *For $n \gg 0$, any closed geodesic $\gamma \subset E_n$ whose length is at most $2 \log n$ can be homotoped to be disjoint from the mutation surface F_n. Consequently, there is a bijection between the complex length spectra of E_n and E_n^μ up to length $2 \log n$.*

Proof. This follows by the same argument as in Proposition 5.4 and Lemma 5.9. However, we need to pay more attention to the cusp of E_n.

By Proposition 3.2, and our choice of φ as in Remark 3.3, the manifold $W_{\varphi,n}$ satisfies

$$d_0(\partial_- \mathrm{core}^0\, W_{\varphi,n},\, \partial_+ \mathrm{core}^0\, W_{\varphi,n}) > n/2.$$

Recall that for a general hyperbolic manifold M, and the Margulis constant ϵ, $\mathrm{core}^0(M)$ is the complement of the ϵ-thin neighborhood of the cusps of M. Furthermore, $d_0(\cdot,\cdot)$ is the shortest length of a path from the lower boundary of $\mathrm{core}^0\, W_{\varphi,n}$ to the upper boundary of $\mathrm{core}^0\, W_{\varphi,n}$, among paths that remain inside $\mathrm{core}^0\, W_{\varphi,n}$.

As in Proposition 5.4, we may construct pleated surfaces $\Sigma_\pm \subset E_n$ whose geometry closely approximates $\partial_\pm \mathrm{core}\, W_{\varphi,n}$. By the argument in that proposition (using the Brock–Bromberg bilipschitz theorem [**13**]), the portion of these surfaces in $\mathrm{core}^0 E_n$ satisfies

$$d_0(\Sigma_-,\, \Sigma_+) > n/2. \tag{6.1}$$

As above, $d_0(\cdot,\cdot)$ only considers paths that remain in $\mathrm{core}^0 E_n$.

Let $\gamma \subset E_n$ be a closed geodesic of length at most $2 \log n$. We claim that γ can be homotoped to be disjoint from at least one of Σ_- and Σ_+. Note that equation (6.1) is not immediately applicable, because γ may fail to be entirely contained in $\mathrm{core}^0 E_n$. However, γ must have non-trivial intersection with $\mathrm{core}^0 E_n$, because the ϵ-thin horocusp $C_n = E_n \setminus \mathrm{core}^0 E_n$ contains no closed geodesics.

If $\gamma \subset \mathrm{core}^0 E_n$, we already have $\ell(\gamma)/2 \leqslant \log n < n/2$. Otherwise, if γ is not contained in $\mathrm{core}^0 E_n$, we may break it up into geodesic sub-arcs $\gamma_1, \ldots, \gamma_{2k}$, such that every odd-numbered γ_i is contained in $\mathrm{core}^0 E_n$ and every even-numbered γ_i is contained in C_n. Every even-numbered geodesic arc γ_i is homotopic to a horocyclic segment $\gamma_i' \subset \partial C_n$ with the same endpoints. By [21, Lemma A.2],

$$\frac{\ell(\gamma_i')}{2} = \sinh \frac{\ell(\gamma_i)}{2}.$$

Now, construct a closed loop $\gamma' = \gamma_1 \cdot \gamma_2' \cdot \ldots \cdot \gamma_{2k-1} \cdot \gamma_{2k}'$, where $\cdot$ denotes concatenation. By construction, γ' is homotopic to γ, contained in $\mathrm{core}^0 E_n$, and its length satisfies

$$\begin{aligned}\frac{\ell(\gamma')}{2} &= \sum_{\text{odd } i} \frac{\ell(\gamma_i)}{2} + \sum_{\text{even } i} \frac{\ell(\gamma_i')}{2} \\ &< \sum_{\text{odd } i} \sinh \frac{\ell(\gamma_i)}{2} + \sum_{\text{even } i} \sinh \frac{\ell(\gamma_i)}{2} \\ &< \sinh \left(\sum_{i=1}^{2k} \frac{\ell(\gamma_i)}{2} \right) \\ &= \sinh \frac{\ell(\gamma)}{2} \\ &< \frac{1}{2} \exp \big(\ell(\gamma)/2\big) \\ &\leqslant n/2.\end{aligned}$$

Since $d_0(\Sigma_-, \Sigma_+) > n/2$, and the surfaces $\Sigma_\pm$ are (homologically) separating, the closed loop γ' must be disjoint from either Σ_- or Σ_+. As a consequence, the geodesic γ (which is homotopic to γ') can be homotoped to be disjoint from F_n.

Since all geodesics of length at most $2 \log n$ can be homotoped to be disjoint from F_n, the conclusion about length spectra follows by a theorem of Millichap [41, Proposition 4.4]. □

Lemma 6.7. *For $n \gg 0$, each of E_n and E_n^μ contains at least $n^2/(2 \log n)$ closed geodesics up to length $2 \log n$.*

Proof. This is proved by exactly the same argument as in Lemma 5.10. Applying Proposition 4.5 with the length cutoff $L = 2 \log n$ gives the desired conclusion for the manifold X_n that covers both E_n and E_n^μ. Since the top and bottom tangles T and B are simple by Lemma 6.4, the same conclusion holds for E_n and E_n^μ. □

Lemmas 6.3–6.7 complete the proof of Theorem 1.5. □

Acknowledgements. We thank Alan Reid and Sam Taylor for a number of enlightening conversations. We are also grateful to Ian Agol, Priyam Patel, Juan Souto, and Genevieve Walsh for their helpful comments and suggestions.

References

1. C. C. Adams, 'Thrice-punctured spheres in hyperbolic 3-manifolds', *Trans. Amer. Math. Soc.* 287 (1985) 645–656.
2. C. C. Adams, 'Augmented alternating link complements are hyperbolic', Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984), London Mathematical Society Lecture Note Series 112 (Cambridge University Press, Cambridge, 1986) 115–130.
3. C. C. Adams, 'Noncompact hyperbolic 3-orbifolds of small volume', *Topology '90* (Columbus, OH, 1990), Ohio State University Mathematical Research Institute Publications (de Gruyter, Berlin, 1992) 1–15.
4. R. Benedetti and C. Petronio, *Lectures on hyperbolic geometry* (Universitext, Springer, Berlin, 1992).
5. I. Biringer, 'Metrizing the Chabauty topology', Preprint, 2016, arXiv:1610.07396.

6. I. Biringer and J. Souto, 'Ranks of mapping tori via the curve complex', *J. reine angew. Math.* (2016), https://doi.org/10.1515/crelle-2016-0031.

7. F. Bonahon, 'Bouts des variétés hyperboliques de dimension 3', *Ann. of Math.* (2) 124 (1986) 71–158.

8. F. Bonahon and L. Siebenmann, 'New geometric splittings of classical knots, and the classification and symmetries of arborescent knots', http://www-bcf.usc.edu/∼fbonahon/Research/Preprints/BonSieb.pdf.

9. A. Borel, 'Commensurability classes and volumes of hyperbolic 3-manifolds', *Ann. Scuola Norm. Sup. Pisa Cl. Sci.* (4) 8 (1981) 1–33.

10. B. H. Bowditch, 'The ending lamination theorem', http://homepages.warwick.ac.uk/∼masgak/papers/elt.pdf.

11. R. Bowen, 'Hausdorff dimension of quasicircles', *Publ. Math. Inst. Hautes Etudes Sci.* 50 (1979) 11–25.

12. J. F. Brock, 'The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores', *J. Amer. Math. Soc.* 16 (2003) 495–535.

13. J. F. Brock and K. W. Bromberg, 'On the density of geometrically finite Kleinian groups', *Acta Math.* 192 (2004) 33–93.

14. J. F. Brock and K. W. Bromberg, 'Geometric inflexibility and 3-manifolds that fiber over the circle', *J. Topol.* 4 (2011) 1–38.

15. J. F. Brock and N. M. Dunfield, 'Injectivity radii of hyperbolic integer homology 3-spheres', *Geom. Topol.* 19 (2015) 497–523.

16. B. A. Burton, R. Budney, W. Pettersson, et al., *Regina: Software for 3-manifold topology and normal surface theory*, http://regina.sourceforge.net/, 1999–2014.

17. P. Buser, 'The collar theorem and examples', *Manuscripta Math.* 25 (1978) 349–357.

18. R. D. Canary, D. B. A. Epstein and P. L. Green, 'Notes on notes of Thurston', *Fundamentals of hyperbolic geometry: selected expositions*, London Mathematical Society Lecture Note Series 328 (Cambridge University Press, Cambridge, 2006), With a new foreword by Canary, 1–115.

19. E. Chesebro and J. DeBlois, 'Algebraic invariants, mutation, and commensurability of link complements', *Pacific J. Math.* 267 (2014) 341–398.

20. T. Chinburg, E. Hamilton, D. D. Long and A. W. Reid, 'Geodesics and commensurability classes of arithmetic hyperbolic 3-manifolds', *Duke Math. J.* 145 (2008) 25–44.

21. D. Cooper, D. Futer and J. S. Purcell, 'Dehn filling and the geometry of unknotting tunnels', *Geom. Topol.* 17 (2013) 1815–1876.

22. D. Cooper, C. D. Hodgson and S. P. Kerckhoff, *Three-dimensional orbifolds and cone-manifolds*, MSJ Memoirs 5 (Mathematical Society of Japan, Tokyo, 2000), With a postface by Sadayoshi Kojima.

23. M. Culler, N. M. Dunfield and J. R. Weeks, 'SnapPy, a computer program for studying the geometry and topology of 3-manifolds', http://snappy.computop.org/verify.html.

24. D. Futer and F. Guéritaud, 'Angled decompositions of arborescent link complements', *Proc. Lond. Math. Soc.* (3) 98 (2009) 325–364.

25. R. Gangolli and G. Warner, 'Zeta functions of Selberg's type for some noncompact quotients of symmetric spaces of rank one', *Nagoya Math. J.* 78 (1980) 1–44.

26. F. W. Gehring and G. J. Martin, 'Minimal co-volume hyperbolic lattices. I. The spherical points of a Kleinian group', *Ann. of Math.* (2) 170 (2009) 123–161.

27. O. Goodman, D. Heard and C. Hodgson, 'Commensurators of cusped hyperbolic manifolds', *Exp. Math.* 17 (2008) 283–306.

28. C. McA. Gordon, Small surfaces and Dehn filling, *Proceedings of the Kirbyfest* (Berkeley, CA, 1998), Geom. Topol. Monogr. 2 (Geom. Topol. Publ., Coventry, 1999) 177–199 (electronic).

29. F. Guéritaud, 'On canonical triangulations of once-punctured torus bundles and two-bridge link complements', *Geom. Topol.* 10 (2006) 1239–1284, With an appendix by David Futer.

30. C. D. Hodgson and S. P. Kerckhoff, 'The shape of hyperbolic Dehn surgery space', *Geom. Topol.* 12 (2008) 1033–1090.

31. H. Huber, 'Zur analytischen Theorie hyperbolischen Raumformen und Bewegungsgruppen', *Math. Ann.* 138 (1959) 1–26.

32. W. Jaco, *Lectures on three-manifold topology*, CBMS Regional Conference Series in Mathematics 43 (American Mathematical Society, Providence, R.I., 1980).

33. B. Linowitz, D. B. McReynolds, P. Pollack and L. Thompson, 'Counting and effective rigidity in algebra and geometry', Preprint, 2014, arXiv:1407.2294.

34. A. Lubotzky, B. Samuels and U. Vishne, 'Division algebras and noncommensurable isospectral manifolds', *Duke Math. J.* 135 (2006) 361–379.

35. C. Maclachlan and A. W. Reid, *The arithmetic of hyperbolic 3-manifolds*, Graduate Texts in Mathematics 219 (Springer, New York, 2003).

36. A. D. Magid, 'Deformation spaces of Kleinian surface groups are not locally connected', *Geom. Topol.* 16 (2012) 1247–1320.

37. G. A. Margulis, 'Certain applications of ergodic theory to the investigation of manifolds of negative curvature', *Funktsional. Anal. i Priložhen.* 3 (1969) 89–90.

38. G. A. Margulis, *Discrete subgroups of semisimple Lie groups*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)], vol. 17 (Springer, Berlin, 1991).

39. T. H. Marshall and G. J. Martin, 'Minimal co-volume hyperbolic lattices, II: simple torsion in a Kleinian group', *Ann. of Math.* (2) 176 (2012) 261–301.

40. W. Menasco and M. Thistlethwaite, 'The classification of alternating links', *Ann. of Math.* (2) 138 (1993) 113–171.
41. C. Millichap, 'Mutations and short geodesics in hyperbolic 3–manifolds', *Comm. Anal. Geom.*, Preprint, 2014, arXiv:1406.6033.
42. W. D. Neumann and A. W. Reid, 'Arithmetic of hyperbolic manifolds', *Topology '90* (Columbus, OH, 1990), Ohio State University Mathematical Research Institute Publications 1 (de Gruyter, Berlin, 1992) 273–310.
43. W. D. Neumann and D. Zagier, 'Volumes of hyperbolic three-manifolds', *Topology* 24 (1985) 307–332.
44. K. Ohshika, 'Kleinian groups which are limits of geometrically finite groups', *Mem. Amer. Math. Soc.* 177 (2005) xii+116.
45. M. Pollicott and R. Sharp, 'Exponential error terms for growth functions on negatively curved surfaces', *Amer. J. Math.* 120 (1998) 1019–1042.
46. G. Prasad and A. S. Rapinchuk, 'Weakly commensurable arithmetic groups and isospectral locally symmetric spaces', *Publ. Math. Inst. Hautes Études Sci.* 109 (2009) 113–184.
47. G. Prasad and A. S. Rapinchuk, 'Weakly commensurable groups, with applications to differential geometry', *Handbook of group actions*, vol. I, Advanced Lectures in Mathematics (ALM), vol. 31 (International Press, Somerville, MA, 2015) 495–524.
48. A. W. Reid, 'Isospectrality and commensurability of arithmetic hyperbolic 2- and 3-manifolds', *Duke Math. J.* 65 (1992) 215–228.
49. A. W. Reid, *The geometry and topology of arithmetic hyperbolic 3-manifolds*, Proceedings of the Symposium on Topology, Complex Analysis, and Arithmetic of Hyperbolic Spaces, Kyoto 2006, RIMS Kokyuroku Series, vol. 1571 (2007) 31–58.
50. A. W. Reid, 'Traces, lengths, axes and commensurability', *Ann. Fac. Sci. Toulouse Math.* (6) 23 (2014) 1103–1118.
51. A. W. Reid and G. S. Walsh, 'Commensurability classes of 2-bridge knot complements', *Algebr. Geom. Topol.* 8 (2008) 1031–1057.
52. T. Roblin, 'Ergodicité et équidistribution en courbure négative', *Mém. Soc. Math. Fr. (N.S.)* 95 (2003) 1–96.
53. D. Ruberman, 'Mutation and volumes of knots in S^3', *Invent. Math.* 90 (1987) 189–216.
54. T. Sunada, 'Riemannian coverings and isospectral manifolds', *Ann. of Math.* (2) 121 (1985) 169–186.
55. W. P. Thurston, *Hyperbolic structures on 3-manifolds, ii: Surface groups and 3-manifolds which fiber over the circle*, Preprint, 1998, arXiv:math/9801045.
56. W. P. Thurston, *Geometry and topology of 3-manifolds, lecture notes* (Princeton University, 1978).
57. W. P. Thurston, 'Three-dimensional manifolds, Kleinian groups and hyperbolic geometry', *Bull. Amer. Math. Soc. (N.S.)* 6 (1982) 357–381.
58. M.-F. Vignéras, 'Variétés riemanniennes isospectrales et non isométriques', *Ann. of Math.* (2) 112 (1980) 21–33.
59. Y.-Q. Wu, 'Dehn surgery on arborescent knots', *J. Differential Geom.* 43 (1996) 171–197.

David Futer
Department of Mathematics
Temple University
Philadelphia, PA 19122
USA

dfuter@temple.edu

Christian Millichap
Department of Mathematics
Linfield College
900 SE Baker Street
McMinnville, OR 97128
USA

Cmillich@linfield.edu