

COMMENSURATIONS OF $\text{Out}(F_n)$

by BENSON FARB and MICHAEL HANDEL*

ABSTRACT

Let $\text{Out}(F_n)$ denote the outer automorphism group of the free group F_n with $n > 3$. We prove that for any finite index subgroup $\Gamma < \text{Out}(F_n)$, the group $\text{Aut}(\Gamma)$ is isomorphic to the normalizer of Γ in $\text{Out}(F_n)$. We prove that Γ is *co-Hopfian*: every injective homomorphism $\Gamma \rightarrow \Gamma$ is surjective. Finally, we prove that the abstract commensurator $\text{Comm}(\text{Out}(F_n))$ is isomorphic to $\text{Out}(F_n)$.

CONTENTS

1. Introduction	1
1.1 Statement of results	2
1.2 Outline of the proof of Theorem 1.1	5
2. The topology of free group automorphisms	8
2.1 Automorphisms and graphs	9
2.2 Free factors	10
2.3 UL subgroups and Kolchin representatives	13
2.4 Axes and multiplicity	14
2.5 Fixed subgroups	15
2.6 Dehn twists	18
3. The endgame	20
3.1 Reduction to the action on elementary automorphisms	20
3.2 Proofs of the corollaries to Theorem 1.1	22
4. A commensurability invariant	22
4.1 The invariant $r(\phi, \mathcal{A})$	23
4.2 The subgroup $\text{O}(F_2, F_{n-2})$	23
4.3 Calculating $\text{WZ}(\mathcal{C}(\phi))$	25
5. The action on special abelian subgroups	27
5.1 Elementary abelian subgroups	28
5.2 Abelian subgroups of IA_n	28
6. Respecting a free factor while almost fixing an abelian subgroup	32
6.1 Comparing normalizations	32
6.2 Preserving $\text{O}(F_2)$ and $\text{O}(F_{n-2})$	34
7. Almost fixing certain subgroups attached to a free factor	37
7.1 Normalizing with respect to $\text{O}(F_2)$	37
7.2 Normalizing with respect to $\text{O}(F_{n-2})$	42
8. Moving between bases	44

1. Introduction

Let F_n denote the free group of rank n and let $\text{Out}(F_n) = \text{Aut}(F_n)/\text{Inn}(F_n)$ denote its group of outer automorphisms. The group $\text{Out}(F_n)$ has been a central example in combinatorial and geometric group theory ever since it was studied by Nielsen (1917), Magnus (1934) and J. H. C. Whitehead (1936). It is, along with the

* Both authors are supported in part by the NSF.

mapping class group Mod_g , a fundamental example to consider when trying to extend group theory ideas to a nonlinear context¹, and rigidity ideas beyond lattices in Lie groups. One reason that $\text{Out}(F_n)$ plays this role is that, while the basic tools and invariants from the theory of linear groups are no longer available, there is a well-known analogy between $\text{Out}(F_n)$ and lattices which has proven to be surprisingly useful (see, e.g., [Vo]). However, $\text{Out}(F_n)$ analogues of theorems about lattices or linear groups can be much harder to prove than their linear versions. A dramatic illustration of this is the Tits Alternative; see [BFH1, BFH2, BFH3].

In this paper we will prove an analogue of strong (Mostow) rigidity for $\text{Out}(F_n)$. As a start to explaining this, consider an irreducible lattice Γ in a semisimple Lie group $G \neq \text{SL}(2, \mathbf{R})$. One consequence of the strong rigidity of these Γ (proved by Mostow, Prasad and Margulis – see [Ma, Zi]) is that $\text{Out}(\Gamma)$ is finite. Incidentally, in the exceptional cases when $\Gamma < \text{SL}(2, \mathbf{R})$, we know that Γ is either a free group or a closed surface group, so that $\text{Out}(\Gamma)$ is either $\text{Out}(F_n)$ or Mod_g (the latter by a theorem of Dehn-Nielsen-Behr).

Some analogous results are known for automorphism groups of free groups. In 1975 Dyer-Formanek [DF] proved for $n \geq 3$ that $\text{Out}(\text{Aut}(F_n)) = 1$; Khramtsov [Kh] and Bridson-Vogtmann [BV] later proved that $\text{Out}(\text{Out}(F_n)) = 1$. While the proofs of these results are quite different from each other, each uses torsion in an essential way. As with most rigidity theorems, one really wants to prove the corresponding results for all finite index subgroups $\Gamma < \text{Out}(F_n)$. Such Γ are almost always torsion free. Further, one cannot use specific relations in $\text{Out}(F_n)$ because most of these disappear in Γ ; indeed it is still not known whether or not such Γ have finite abelianization, as does $\Gamma = \text{Out}(F_n)$, $n > 2$. Thus the computation of $\text{Out}(\Gamma)$ requires a new approach.

1.1. *Statement of results.* — The main result of this paper is the following theorem, which can be thought of as strong (Mostow) rigidity in this context.

1.1. *Theorem.* — *Let $n \geq 4$, let $\Gamma < \text{Out}(F_n)$ be any finite index subgroup and let $\Phi : \Gamma \rightarrow \text{Out}(F_n)$ be any injective homomorphism. Then there exists $g \in \text{Out}(F_n)$ such that $\Phi(\gamma) = g\gamma g^{-1}$ for all $\gamma \in \Gamma$.*

Theorem 1.1 implies in particular that $\Phi(\Gamma)$ must have finite index in $\text{Out}(F_n)$. We do not know a direct proof of this seemingly much easier fact. In Section 3.2 we use Theorem 1.1 to deduce the following.

1.2. *Corollary.* — *Let $n \geq 4$, let $\Gamma < \text{Out}(F_n)$ be any finite index subgroup, and let $N(\Gamma)$ denote the normalizer of Γ in $\text{Out}(F_n)$. Then the natural map*

$$N(\Gamma) \longrightarrow \text{Aut}(\Gamma)$$

¹ Unlike Mod_g , the group $\text{Out}(F_n)$, $n \geq 4$ is known (see [FP]) to be *nonlinear*, i.e. it admits no faithful representation into any matrix group over any field.

given by $f \mapsto \text{Conj}_f$ is an isomorphism. Here Conj_f is defined by $\text{Conj}_f(\gamma) := f\gamma f^{-1}$ for all $\gamma \in \Gamma$.

Taking $\Gamma = \text{Out}(\mathbb{F}_n)$ in Corollary 1.2 recovers the result $\text{Aut}(\text{Out}(\mathbb{F}_n)) = \text{Out}(\mathbb{F}_n)$. Note that $\text{Out}(\mathbb{F}_n)$ has infinitely many mutually nonconjugate finite index subgroups; indeed $\text{Out}(\mathbb{F}_n)$ is residually finite.

We now discuss two further corollaries of Theorem 1.1: a proof of the co-Hopf property for all finite index subgroups of $\text{Out}(\mathbb{F}_n)$, and a computation of the abstract commensurator of $\text{Out}(\mathbb{F}_n)$.

The co-Hopf property. — A group Λ is *co-Hopfian* if every injective endomorphism of Λ is an isomorphism. Unlike the Hopf property, which is true for example for all linear groups, the co-Hopf property holds much less often (consider, for example, any Λ which is free abelian or is a nontrivial free product), and is typically harder to prove. The co-Hopf property was proven for lattices in semisimple Lie groups by Prasad [Pr], and for mapping class groups by Ivanov [Iv2]. Theorem 1.1 immediately implies the following.

1.3. Corollary. — *For $n \geq 4$, every finite index subgroup $\Gamma < \text{Out}(\mathbb{F}_n)$ is co-Hopfian.*

Commensurators. — The (abstract) commensurator group $\text{Comm}(\Lambda)$ of a group Λ is defined to be the set of equivalence classes of isomorphisms $\phi : H \rightarrow N$ between finite index subgroups H, N of Λ , where the equivalence relation is the one generated by the relation that $\phi_1 : H_1 \rightarrow N_1$ is equivalent to $\phi_2 : H_2 \rightarrow N_2$ if $\phi_1 = \phi_2$ on some finite index subgroup of Λ . The set $\text{Comm}(\Lambda)$ is a group under composition. We think of $\text{Comm}(\Lambda)$ as the group of “hidden automorphisms” of Λ .

$\text{Comm}(\Lambda)$ is in general much larger than $\text{Aut}(\Lambda)$. For example $\text{Aut}(\mathbf{Z}^n) = \text{GL}(n, \mathbf{Z})$ whereas $\text{Comm}(\mathbf{Z}^n) = \text{GL}(n, \mathbf{Q})$. Margulis proved that an irreducible lattice Λ in a semisimple Lie group G is arithmetic if and only if it has infinite index in its commensurator in G . Mostow-Prasad-Margulis strong rigidity for the collection of irreducible lattices Λ in such a $G \neq \text{SL}(2, \mathbf{R})$ can be thought of as proving exactly that the abstract commensurator $\text{Comm}(\Lambda)$ is isomorphic to the commensurator of Λ in G , which in turn is computed concretely by Margulis and Borel-Harish-Chandra; see, e.g., [Ma, Zi]. The group $\text{Comm}(\text{Mod}_g)$ was computed for surface mapping class groups Mod_g by Ivanov [Iv2].

While for arbitrary groups Γ the group $\text{Comm}(\Gamma)$ can be much bigger than $\text{Aut}(\Gamma)$, we will see in Section 3.2 that Theorem 1.1 implies the following.

1.4. Corollary. — *For $n \geq 4$ the natural injection*

$$\text{Out}(\mathbb{F}_n) \rightarrow \text{Comm}(\text{Out}(\mathbb{F}_n))$$

is an isomorphism.

Corollary 1.4 answers Question 8 of K. Vogtmann’s list (see [Vo]) of open problems about $\text{Out}(F_n)$.

An application. — Recall that the *commensurator* $\text{Comm}_G(\Gamma)$ of a group Γ in a group G is defined as

$$\begin{aligned} \text{Comm}_G(\Gamma) \\ := \{g \in G : g\Gamma g^{-1} \cap \Gamma \text{ has finite index in both } \Gamma \text{ and } g\Gamma g^{-1}\}. \end{aligned}$$

Let Γ and G be discrete groups. A theorem of Mackey (see [BuH]) states that $\text{Comm}_G(\Gamma) = \Gamma$ if and only if the left regular representation of G on $\ell^2(G/\Gamma)$ is irreducible. He also proved that when this happens, the unitary induction map Ind_Γ^G on finite-dimensional representations is injective. Note that there is an exact sequence

$$1 \rightarrow \text{VC}_G(\Gamma) \rightarrow \text{Comm}_G(\Gamma) \rightarrow \text{Comm}(\Gamma) \rightarrow 1$$

where $\text{VC}_G(\Gamma)$ is the *virtual centralizer* of Γ in G , i.e. the group of elements $g \in G$ for which there is some finite index subgroup $H < \Gamma$ so that g commutes with H . Now consider a group Γ with $\text{Comm}(\Gamma) = \Gamma$, for example $\Gamma = \text{Out}(F_n)$ (by Corollary 1.4). We then see that for any discrete group G into which Γ embeds with $\text{VC}_G(\Gamma) = 1$ (note that this condition is easy to check), Mackey’s theorem applies. In this way the unitary representation theory of $\text{Out}(F_n)$ is “atomic”: it injects into the unitary representation theory of any group containing it in a “nontrivial” way.

The cases $n = 2$ and $n = 3$. — The conclusion of each result stated above is false when $n = 2$. Indeed, Nielsen proved that $\text{Out}(F_2) \approx \text{GL}(2, \mathbf{Z})$, which has non-abelian free subgroups of finite index. Thus $\text{Comm}(\text{Out}(F_2)) \approx \text{Comm}(F_2)$. Since F_2 contains each F_m , $m \geq 2$ as a finite index subgroup, $\text{Comm}(F_2) = \text{Comm}(F_m)$. It is easy to see that each F_m , $m \geq 2$ has self-injections of infinite index, and that $\text{Comm}(F_m)$ is enormous, in particular it contains F_m as an infinite index subgroup. Note also that any finite index subgroup of F_m is isomorphic to F_n for some $n \geq m$, and so has automorphism group $\text{Aut}(F_n)$, while the normalizer $N(F_n)$ in F_m is just F_n . We do not know what happens when $n = 3$, and propose each of the results above as an open question in this case. We note that Khramtsov and Bridson-Vogtmann’s proofs that $\text{Out}(\text{Out}(F_n)) = 1$ hold for $n = 3$.

Comparison with mapping class groups. — While some aspects of the general outline of our approach to Theorem 1.1 follow that of Ivanov for (extended) mapping class groups $\text{Mod}_g^\pm \approx \text{Out}(\pi_1 \Sigma_g)$ (see [Iv2]), there are fundamental differences between the two problems. The natural analogue of a Dehn twist in this context is played by the so-called *elementary automorphisms* of $\text{Out}(F_n)$ (see below). The key to understanding an injective endomorphism Φ of a finite index subgroup of Mod_g (resp. $\text{Out}(F_n)$) is to determine the image of each Dehn twist (resp. each elementary automorphism) under Φ . In the case of Mod_g , the following facts are crucial for such an understanding:

1. A Dehn twist is completely determined by specifying the conjugacy class in $\pi_1 \Sigma_g$ of a simple closed curve.
2. The set of all such curves (hence twists), along with the data recording whether or not they are disjoint (hence commute), is encoded in a simplicial complex, the *complex of curves* \mathcal{C}_g , whose automorphism group was determined (using topology) by Ivanov to be Mod_g .
3. Centralizers in Mod_g are essentially completely understood. This knowledge can be used to compute invariants characterizing certain elements of Mod_g , which in turn can be used to prove that any Φ as above induces an automorphism of \mathcal{C}_g .

Some of the serious obstacles to understanding the $\text{Out}(F_n)$ case now become apparent. First, an elementary automorphism is not simply determined by a single conjugacy class in F_n . Second, the powerful tool of Ivanov's theorem on automorphisms of \mathcal{C}_g is not available for $\text{Out}(F_n)$. Indeed we do not know of a simplicial complex that encodes commutations between elementary automorphisms. Finally, the theory of abelian subgroups and centralizers in $\text{Out}(F_n)$ is more complex and less well-developed than the corresponding theory for Mod_g (see [BFH2, FH]), making computations of the corresponding invariants more difficult. Thus a different approach is needed.

1.2. Outline of the proof of Theorem 1.1. — For $\psi \in \text{Out}(F_n)$, we denote by i_ψ the inner automorphism of $\text{Out}(F_n)$ defined by $i_\psi(\phi) = \psi\phi\psi^{-1}$ for all $\phi \in \text{Out}(F_n)$. Theorem 1.1 is a reconstruction problem: we are given an arbitrary injective homomorphism $\Phi : \Gamma \rightarrow \text{Out}(F_n)$, and we must construct some ψ for which $\Phi = i_\psi$. The automorphisms i_ψ have a number of special properties, and they preserve various special collections of elements and subgroups of $\text{Out}(F_n)$. The general strategy is to prove that Φ must do the same, so much so that we can eventually pin down Φ to be some i_ψ . More precisely, for any $\psi \in \text{Out}(F_n)$, we say that the injection $i_\psi \circ \Phi$ is a *normalization* of Φ . Our goal will be to perform repeated normalizations on Φ until the resulting map fixes every $\phi \in \Gamma$, thus proving the theorem.

In order to execute the above strategy one needs to give purely algebraic characterizations of (conjugacy classes of) various types of elements and subgroups; of course the characterizing properties must also be commensurability invariants. Another aspect is to encode the combinatorics of the collections of these subgroups and their intersection patterns in order to deduce finer structure.

Terminology. — As we are dealing with a finite index subgroup $\Gamma < \text{Out}(F_n)$, we will need to work with “almost” or “weak” versions of standard concepts. For example, we say that Φ *almost fixes* ϕ if there exist $s, t > 0$ such $\Phi(\phi^s) = \phi^t$, and that Φ *almost fixes a subgroup* \mathcal{A} if there exist $s, t > 0$ such $\Phi(\phi^s) = \phi^t$ for all $\phi \in \mathcal{A}$. Thus $\phi \in \Gamma$ is

almost fixed by some normalization of Φ if and only if $\Phi(\phi)$ is *weakly conjugate* to ϕ , meaning that $\Phi(\phi)^s$ is conjugate to ϕ^t for some $s, t > 0$.

Dynamics. — A typical way to understand elements $\phi \in \text{Out}(F_n)$ is via their dynamical properties, such as the rate of growth of the length of a word in F_n under repeated iterations of ϕ . Unfortunately these properties are not *a priori* commensurator invariants, and so they cannot be used to relate $\Phi(\phi)$ to ϕ . However, we will make repeated use of the set $\text{Fix}(\phi)$ of fixed subgroups associated to ϕ (see Section 2.5) to understand the centralizer of ϕ .

Our proof of Theorem 1.1 proceeds in steps.

Step 1 (Reduction to the action on elementaries). — Given a basis x_1, \dots, x_n for F_n , define automorphisms \hat{E}_{jk} and ${}_{kj}\hat{E}$ by

$$\begin{aligned} \hat{E}_{jk} &: x_j \mapsto x_j x_k \\ {}_{kj}\hat{E} &: x_j \mapsto \bar{x}_k x_j, \end{aligned}$$

where any basis element whose image is not explicitly mentioned is fixed. The outer automorphisms that they determine will be denoted E_{jk} and ${}_{kj}E$. A nontrivial outer automorphism μ is *elementary* if there is some choice of basis for F_n for which μ is an iterate of some E_{jk} or ${}_{kj}E$.

We begin by proving (Lemma 3.2), using an argument of Ivanov (§8.5 of [Iv]), that Theorem 1.1 can be reduced to finding a normalization of Φ that almost fixes each elementary outer automorphism in $\text{Out}(F_n)$.

Step 2 (Action on special abelian subgroups). — The injective homomorphism $\Phi : \Gamma \rightarrow \text{Out}(F_n)$ acts on the collection of (commensurability classes of) abelian subgroups of $\text{Out}(F_n)$. We will consider various special families of abelian subgroups of $\text{Out}(F_n)$ and, using some of the results from [FH], we will prove that these special classes of subgroups must be invariant under this action.

To give an example, we say that an element $\phi \in \text{Out}(F_n)$ is *unipotent* if its image in $\text{GL}(n, \mathbf{Z})$ is a unipotent matrix, and that ϕ has *linear growth* if the word length of any conjugacy class in F_n grows linearly under iteration by ϕ . We say that a subgroup of $\text{Out}(F_n)$ is *UL* if each of its elements is unipotent and linear. Define

$$\mathcal{A}_k := \langle E_{jk}, {}_{jk}E : j \neq k \rangle$$

where $\langle G \rangle$ denotes the group generated by G . Note that \mathcal{A}_k is free abelian of rank $2n - 3$.

A first step towards finding a normalization almost fixing each elementary outer automorphism is given in Corollary 5.2, which states that for any choice of basis for F_n there is a normalization of Φ that almost fixes \mathcal{A}_k . In particular, for each elemen-

tary $\psi \in \text{Out}(F_n)$ there exists a normalization of Φ which almost fixes ψ . The proof of Corollary 5.2 uses a commensurability invariant introduced in Section 4, together with results from [FH]. These results include the classification of abelian subgroups of maximal rank in $\text{Out}(F_n)$, as well as information on the rank of the weak center of the centralizer of an element of $\text{Out}(F_n)$.

Choosing an element $\psi \in \text{Out}(F_n)$ so that the normalization $i_\psi \circ \Phi$ almost fixes a given elementary ϕ can be viewed as choosing a basis with respect to which $\Phi(\psi)$ has a standard presentation. This brings to mind Kolchin's Theorem on linear groups, which states that if each element of a subgroup $H < \text{GL}(n, \mathbf{Z})$ has a basis with respect to which it is upper triangular with ones on the diagonal, then there is a single basis with respect to which every element of H has this form. The *UL Kolchin Theorem* of [BFH2] gives a version of Kolchin's Theorem for finitely generated UL subgroups of $\text{Out}(F_n)$, even those that are not abelian. This will be crucial in Step 3.

Warning on pinning down Φ via normalizations. — The ultimate normalization $i_\psi \circ \Phi$ that fixes every elementary in $\text{Out}(F_n)$ is unique, but the normalizations that occur as the proof progresses are not. It is easy to see (Lemma 6.2) that if $i_{\psi_1} \circ \Phi$ and $i_{\psi_2} \circ \Phi$ almost fix the elementary outer automorphism ϕ , then ψ_1 and ψ_2 differ by an element of the *weak centralizer* of ϕ ; i.e. by an element that commutes with some iterate of ϕ . Each time the list of elements weakly fixed by our given normalization grows, we lose some degree of freedom in choosing the normalization. Our challenge then is not only to find normalizations that weakly fix a growing list of elements, but to choose the list very carefully so that we do not use up all of our freedom prematurely.

Step 3 (Action on free factors). — We would like to find a normalization of Φ that almost fixes both E_{12} and E_{21} . Since this subgroup contains elements with exponential growth, the UL Kolchin theorem of [BFH2] does not directly apply. Instead, we consider elements T_w , defined as follows.

Given a basis $\{x_1, \dots, x_n\}$ for F_n , let F_2 and F_{n-2} be the free factors $\langle x_1, x_2 \rangle$ and $\langle x_3, \dots, x_n \rangle$. We denote by $O(F_2)$ the image of the composition

$$\text{Aut}(F_2) \hookrightarrow \text{Aut}(F_2) \times \text{Aut}(F_{n-2}) \hookrightarrow \text{Aut}(F_n) \rightarrow \text{Out}(F_n)$$

where the lefthand map is $\hat{\phi} \mapsto \hat{\phi} \times \text{Id}$, the middle map is inclusion, and the righthand map is the natural projection. We define $O(F_{n-2})$ similarly. The main next step in our proof of Theorem 1.1 is to prove that there is a normalization of Φ which “respects the decomposition $F_n = F_2 * F_{n-2}$ ” in the sense that it preserves both $O(F_2)$ and $O(F_{n-2})$ (and, in fact, further structure). This is done in Proposition 6.1.

To explain some of the key ideas in the proof of this proposition, we begin by letting $i_w \in \text{Aut}(F_2)$ with $w \in [F_2, F_2]$ denote “conjugation by w ”, and by letting $T_w \in O(F_2)$ be the element represented by $i_w \times \text{Id}$. Let IA_n be the subgroup of $\text{Out}(F_n)$ consisting of those elements that act trivially on $H_1(F_n, \mathbf{Z})$. There is a natural

abelian UL subgroup of IA_n that contains T_w (see Section 5.2). We use this fact, together with results from [FH] to prove, roughly speaking, that the set of all such T_w 's is Φ -invariant; see Lemma 5.3 for a precise statement. The UL Kolchin theorem applies to any subgroup generated by finitely many of the T_w because all such subgroups are UL.

Step 4 (Fixing a basis). — We say that a normalization Φ' of Φ *almost fixes a basis* B of F_n if it almost fixes each $\langle_j E, E_{ij} \rangle$ defined with respect to that basis. The next main step is to prove that, given any basis of F_n , there is a normalization of Φ which almost fixes that basis (see Lemma 8.2). This is perhaps the most delicate part of the proof of Theorem 1.1, since we use up all of the freedom in choosing the normalization of Φ before completing the proof. See Section 7.

Step 5 (Moving between bases). — If we could almost fix every basis at once, we would complete the proof of Theorem 1.1. This final piece of “rigidity” comes from an encoding of the space of bases for F_2 via the classical *Farey graph* \mathcal{F} , and from the fact that automorphisms of \mathcal{F} are determined by their action on 3 vertices.

2. The topology of free group automorphisms

In this section we recall some of the topological methods used to understand elements and subgroups of $\text{Out}(F_n)$, and we prove some results which will be used later in the paper.

Notational conventions. — We begin by giving some notation which will be used throughout the paper. We assume throughout that $n \geq 4$.

If a basis $\{x_1, \dots, x_n\}$ for F_n is understood then we will specify elements of $\text{Aut}(F_n)$ by defining their action on those x_i that are not fixed. Thus any unspecified generators are fixed.

If $\phi \in \text{Out}(F_n)$ then $\hat{\phi} \in \text{Aut}(F_n)$ will denote an automorphism representing it. Conversely if $\hat{\phi} \in \text{Aut}(F_n)$ then ϕ will denote the corresponding outer automorphism.

We will use the notation x^\pm to denote an element that might be either x or $\bar{x} = x^{-1}$. We will interpret x^{-k} to be \bar{x}^k .

We denote the conjugacy classes of $x \in F_n$ by $[x]$ and the unoriented conjugacy class by $[x]_u$. Thus $[x]_u = [y]_u$ if and only if $[x] = [y]$ or $[x] = [\bar{y}]$. Similarly the conjugacy class of a subgroup A is denoted $[A]$. An element $\phi \in \text{Out}(F_n)$ acts on the set of all conjugacy classes in F_n . We sometimes say that x or A is ϕ -invariant when, strictly speaking, we really mean that $[x]$ or $[A]$ is ϕ -invariant.

For $\psi \in \text{Out}(F_n)$, we denote by i_ψ the inner automorphism of $\text{Out}(F_n)$ defined by $i_\psi(\phi) = \psi\phi\psi^{-1}$. For $c \in F_n$, we denote by $i_c : F_n \rightarrow F_n$ the inner automorphism of F_n defined by $i_c(x) = cxc^{-1}$.

2.1. Automorphisms and graphs

Marked graphs and outer automorphisms. — Identify F_n , once and for all, with $\pi_1(\mathbf{R}_n, *)$ where \mathbf{R}_n is the *rose* (i.e. graph) with one vertex $*$ and with n edges. A *marked graph* G is a graph with $\pi_1(G) \approx F_n$, with each vertex having valence at least two, equipped with a homotopy equivalence $m : \mathbf{R}_n \rightarrow G$ called a *marking*. Letting $d = m(*) \in G$, the marking determines an identification of F_n with $\pi_1(G, d)$.

A homotopy equivalence $f : G \rightarrow G$ of G determines an outer automorphism of $\pi_1(G, d)$ and hence an element $\phi \in \text{Out}(F_n)$. We say that $f : G \rightarrow G$ *represents* ϕ . A path σ from d to $f(d)$ determines an automorphism of $\pi_1(G, d)$ and hence a representative $\hat{\phi} \in \text{Aut}(F_n)$ of ϕ that depends only on f and the homotopy class of σ . As the homotopy class of σ varies, $\hat{\phi}$ ranges over all representatives of ϕ . If f fixes d and no path is specified, then we use the trivial path.

We always assume that the restriction of f to any edge of G is an immersion.

Paths, circuits and edge paths. — Let Γ be the universal cover of a marked graph G and let $pr : \Gamma \rightarrow G$ be the covering projection. We always assume that a base point $\tilde{d} \in \Gamma$ projecting to $d = m(*) \in G$ has been chosen, thereby identifying the group of covering translations of Γ with $\pi_1(G, b)$, and so defining an action of F_n on Γ . The set of ends $\mathcal{E}(\Gamma)$ of Γ is naturally identified with the boundary ∂F_n of F_n and we make implicit use of this identification throughout the paper.

A proper map $\tilde{\sigma} : J \rightarrow \Gamma$ with domain a (possibly infinite) interval J will be called a *path in Γ* if it is an embedding or if J is finite and the image is a single point; in the latter case we say that $\tilde{\sigma}$ is a *trivial path*. If J is finite, then every map $\tilde{\sigma} : J \rightarrow \Gamma$ is homotopic rel endpoints to a unique (possibly trivial) path $[\tilde{\sigma}]$; we say that $[\tilde{\sigma}]$ is *obtained from $\tilde{\sigma}$ by tightening*. If $\tilde{f} : \Gamma \rightarrow \Gamma$ is a lift of a homotopy equivalence $f : G \rightarrow G$, we denote $[\tilde{f}([\tilde{\sigma}])]$ by $\tilde{f}_\#([\tilde{\sigma}])$.

We will not distinguish between paths in Γ that differ only by an orientation preserving change of parametrization. Thus we are interested in the oriented image of $\tilde{\sigma}$ and not $\tilde{\sigma}$ itself. If the domain of $\tilde{\sigma}$ is finite, then the image of $\tilde{\sigma}$ has a natural decomposition as a concatenation $\tilde{E}_1\tilde{E}_2\dots\tilde{E}_{k-1}\tilde{E}_k$ where \tilde{E}_i , $1 < i < k$, is an edge of Γ , \tilde{E}_1 is the terminal segment of an edge and \tilde{E}_k is the initial segment of an edge. If the endpoints of the image of $\tilde{\sigma}$ are vertices, then \tilde{E}_1 and \tilde{E}_k are full edges. The sequence $\tilde{E}_1\tilde{E}_2\dots\tilde{E}_k$ is called *the edge path associated to $\tilde{\sigma}$* . This notation extends naturally to the case that the interval of domain is half-infinite or bi-infinite. In the former case, an edge path has the form $\tilde{E}_1\tilde{E}_2\dots$ or $\dots\tilde{E}_{-2}\tilde{E}_{-1}$ and in the latter case has the form $\dots\tilde{E}_{-1}\tilde{E}_0\tilde{E}_1\tilde{E}_2\dots$.

A *path in G* is the composition of the projection map pr with a path in Γ . Thus a map $\sigma : J \rightarrow G$ with domain a (possibly infinite) interval will be called a path if it is an immersion or if J is finite and the image is a single point; paths of the latter type are said to be *trivial*. If J is finite, then every map $\sigma : J \rightarrow G$ is homotopic rel

endpoints to a unique (possibly trivial) path $[\sigma]$; we say that $[\sigma]$ is obtained from σ by *tightening*. For any lift $\tilde{\sigma} : J \rightarrow \Gamma$ of σ , $[\sigma] = pr[\tilde{\sigma}]$. We denote $[f(\sigma)]$ by $f_{\#}(\sigma)$. We do not distinguish between paths in G that differ by an orientation preserving change of parametrization. The *edge path associated to σ* is the projected image of the edge path associated to a lift $\tilde{\sigma}$. Thus the edge path associated to a path with finite domain has the form $E_1 E_2 \dots E_{k-1} E_k$ where E_i , $1 < i < k$, is an edge of G , E_1 is the terminal segment of an edge and E_k is the initial segment of an edge. We will identify paths with their associated edge paths whenever it is convenient.

We reserve the word *circuit* for an immersion $\sigma : S^1 \rightarrow G$. Any homotopically nontrivial map $\sigma : S^1 \rightarrow G$ is homotopic to a unique circuit $[\sigma]$. As was the case with paths, we do not distinguish between circuits that differ only by an orientation preserving change in parametrization and we identify a circuit σ with a *cyclically ordered edge path* $E_1 E_2 \dots E_k$. If $f : G \rightarrow G$ is a homotopy equivalence then we denote $[f(\sigma)]$ by $f_{\#}(\sigma)$. There is bijection between circuits in G and conjugacy classes in F_n ; if f represents $\phi \in \text{Out}(F_n)$ then the action of $f_{\#}$ on circuit corresponds to the action of ϕ on conjugacy classes in F_n .

A path or circuit *crosses* or *contains* an edge if that edge occurs in the associated edge path. For any path σ in G define $\bar{\sigma}$ to be ‘ σ with its orientation reversed’. For notational simplicity, we sometimes refer to the inverse of $\tilde{\sigma}$ by $\tilde{\sigma}^{-1}$.

A decomposition of a path or circuit into subpaths is a *splitting* for $f : G \rightarrow G$ and is denoted $\sigma = \dots \sigma_1 \cdot \sigma_2 \dots$ if $f_{\#}^k(\sigma) = \dots f_{\#}^k(\sigma_1) f_{\#}^k(\sigma_2) \dots$ for all $k \geq 0$. In other words, a decomposition of σ into subpaths σ_i is a splitting if one can tighten the image of σ under any iterate of $f_{\#}$ by tightening the images of the σ_i 's.

If $f_{\#}^k(\sigma) = \sigma$ then σ is a *periodic Nielsen path*; if $k = 1$ then σ is a *Nielsen path*. A (periodic) Nielsen path is *indivisible* if it does not decompose as a concatenation of nontrivial (periodic) Nielsen subpaths. A path is *primitive* if it is not multiple of a simpler path.

An unoriented bi-infinite properly embedded path in Γ is called a *line* in Γ . The ends of such a line converge to distinct points in ∂F_n (under the identification of ∂F_n with the set of ends of Γ .) Conversely, any distinct pair of points in ∂F_n are the endpoints of a unique line in Γ . This defines a bijection between lines in Γ and points in $((\partial F_n \times \partial F_n) \setminus \Delta) / Z_2$, where Δ is the diagonal and where Z_2 acts on $\partial F_n \times \partial F_n$ by interchanging the factors. There is an induced action of $\text{Aut}(F_n)$ on the space of lines in Γ . The projection of a line in Γ into G is a *line* in G . An element of $\text{Out}(F_n)$ acts on the space of lines in G .

2.2. Free factors. — If H is a subgroup of F_n and $H = A_1 * \dots * A_m$ is free decomposition then each A_i is a *free factor* of H and A_1, \dots, A_m are *cofactors* of H . We make use of the following special case of the Kurosh subgroup theorem where HcK is the (H, K) *double coset* determined by subgroups H, K and an element c .

2.1. Theorem. — Suppose that F is a free factor of F_n , that H is a subgroup of F_n and that $C = \{c_1, \dots, c_r\}$ where the c_i 's represent distinct (H, F) double cosets. Then $H_{F,C} := (H \cap i_{c_1}(F)) * \dots * (H \cap i_{c_r}(F))$ is a free factor of H . Moreover, if F^1, \dots, F^s are cofactors of F_n and C^j represent distinct (H, F^j) double cosets then $H_{F,C^1}, \dots, H_{F,C^s}$ are cofactors of H .

We record some easy corollaries.

2.2. Corollary. — If H is a subgroup of F_n and F is a free factor of F_n then any conjugate of F that is contained in H is a free factor of H .

Proof. — This is an immediate consequence of Theorem 2.1. \square

2.3. Corollary. — For any $c \in F_n$ and any free factor F of F_n , the following are equivalent.

1. $i_c(F) \cap F$ is nontrivial.
2. $i_c(F) = F$.
3. $c \in F$.

Proof. — It is obvious that (3) implies (2) implies (1). To see that (1) implies (3), note that the (F, F) double coset that contains the identity element is F , and so by Theorem 2.1 it is the only nontrivial (F, F) double coset. \square

2.4. Corollary. — Suppose that J and J' are subsets of $\{1, \dots, n\}$ and that $J \cap J' \neq \emptyset$.

1. If F is a free factor of F_n that is carried by both $\langle x_j : j \in J \rangle$ and $\langle x_j : j \in J' \rangle$ then F is also carried by $\langle x_j : j \in J \cap J' \rangle$.
2. If $\phi \in \text{Out}(F_n)$ and if both $[\langle x_j : j \in J \rangle]$ and $[\langle x_j : j \in J' \rangle]$ are ϕ -invariant then $[\langle x_j : j \in J \cap J' \rangle]$ is ϕ -invariant.

Proof. — Theorem 2.1 applied with $H = \langle x_j : j \in J \rangle$ implies that for all $c \in F_n$, $\langle x_j : j \in J \rangle \cap i_c \langle x_j : j \in J' \rangle$ either is trivial or is $\langle x_j : j \in J \cap J' \rangle$.

To prove (1), we may assume that $F \subset \langle x_j : j \in J \rangle$. By assumption, there exists $c \in F_n$ such that $F \subset i_c \langle x_j : j \in J' \rangle$. Thus

$$F \subset \langle x_j : j \in J \rangle \cap i_c \langle x_j : j \in J' \rangle = \langle x_j : j \in J \cap J' \rangle.$$

To prove (2), choose $\hat{\phi}$ and $a \in F_n$ so that $\hat{\phi}(\langle x_j : j \in J \rangle) = \langle x_j : j \in J \rangle$ and $\hat{\phi}(\langle x_j : j \in J' \rangle) = i_a(\langle x_j : j \in J' \rangle)$. Then

$$\hat{\phi}(\langle x_j : j \in J \cap J' \rangle) = (\langle x_j : j \in J \rangle \cap i_a(\langle x_j : j \in J' \rangle)) = \langle x_j : j \in J \cap J' \rangle. \quad \square$$

2.5. Corollary. — Suppose that $\hat{\phi} \in \text{Aut}(F_n)$, $w \in F_n$ and $\hat{\phi}(w) = w^\pm$. Then every ϕ -invariant free factor F that contains w is $\hat{\phi}$ -invariant.

Proof. — Since F is ϕ -invariant, $\hat{\phi}(F) = i_c(F)$ for some $c \in F_n$. Corollary 2.3 and the fact that $i_c(F) \cap F$ contains w implies that $i_c(F) = F$. \square

If G is a marked graph and G_r is a noncontractible connected subgraph then $[\pi_1(G_r)]$ is well defined and each representative of this conjugacy class is a free factor of F_n . There is a natural bijection between conjugacy classes $[a]$ in F_n and circuits $\sigma \subset G$. If F represents $[\pi_1(G_r)]$ then F contains a representative of $[a]$ if and only if the circuit $\sigma \subset G$ corresponding to $[a]$ is contained in G_r . In this case we say that F and G_r carry $[a]$; sometime we say that F and G_r carry a when we really mean that they carry $[a]$. A line γ in G corresponds to a bi-infinite word w in the generators of F_n . If $\gamma \subset G_r$ then we say that G_r carries γ and that F carries w .

2.6. Definition. — Suppose that A is a collection of conjugacy classes and bi-infinite words in F_n . If there is a free factor F such that:

- (i) F carries each element of A , and
- (ii) for any nontrivial decomposition $F = F_1 * F_2$ into free factors there exists $a \in A$ that is not carried by either F_1 or F_2 .

then we say that F is a minimal carrier of A and write $F = F(A)$.

2.7. Lemma. — If A is a collection of conjugacy classes and bi-infinite words in F_n and if $F(A)$ is a minimal carrier of A then the following are satisfied.

1. Every free factor that carries each element of A contains a subgroup that is conjugate to $F(A)$.
2. $[F(A)]$ does not depend on the choice of minimal carrier $F(A)$.
3. If $\psi \in \text{Out}(F_n)$ and if A is ψ -invariant, then $[F(A)]$ is ψ -invariant.

Proof. — (1) is proved in Section 2.6 of [BFH1]; see in particular, Lemma 2.6.4 and Corollary 2.6.5. (2) follows from (1) and Corollary 2.2. (3) follows from (2) and the fact that if $\hat{\psi}$ represents ψ then $\hat{\psi}(F(A))$ is a minimal carrier of $\psi(A)$. \square

We have the following pair of almost immediate corollaries.

2.8. Corollary. — Suppose that $\{x_1, \dots, x_n\}$ is a basis of F_n , that F is a free factor and that F carries w where w is the conjugacy class of either the commutator $[x_1, x_2]$ or a nonperiodic bi-infinite word in $\langle x_1, x_2 \rangle$. Then F contains a subgroup that is conjugate to $\langle x_1, x_2 \rangle$.

Proof. — Let $A = \{w\}$. Obviously w is carried by $\langle x_1, x_2 \rangle$ but not by any free factor of rank one. Thus $F(A) = \langle x_1, x_2 \rangle$ and the corollary follows from Lemma 2.7. \square

2.9. Corollary. — Suppose that $\phi \in \text{Out}(F_n)$ and that F is a free factor. If $\phi([a])$ is carried by F for each basis element $a \in F$, then $[F]$ is ϕ -invariant.

Proof. — Let A be the set of conjugacy classes of basis elements of F . Obviously F carries each element of A . For any decomposition $F = F_1 * F_2$, choose basis elements $b_i \in F_i$. Then $b_1 b_2$ is a basis element whose conjugacy class is not carried by either F_1 or F_2 . Thus F is a minimal carrier of A . For each $a \in A$, the conjugacy class $\phi(a)$ is represented by an element $b \in F$, which by Corollary 2.2 is a basis element of F . Thus A is ϕ -invariant and Lemma 2.7 implies that $[F]$ is ϕ -invariant. \square

Finally, we recall Lemma 3.2.1 of [BFH1].

2.10. Lemma. — Suppose that $\{x_1, \dots, x_n\}$ is a basis of F_n and that $1 \leq k \leq n - 1$. If $\hat{\phi} \in \text{Aut}(F_n)$ leaves both $\langle x_1, \dots, x_k \rangle$ and $\langle x_1, \dots, x_{k+1} \rangle$ invariant then $\hat{\phi}(x_{k+1}) = ux_{k+1}^\pm v$ for some elements $u, v \in \langle x_1, \dots, x_k \rangle$.

2.3. UL subgroups and Kolchin representatives. — A *filtered graph* is a marked graph along with a filtration

$$\emptyset = G_0 \subset G_1 \subset \dots \subset G_K = G$$

by subgraphs where each G_i is obtained from G_{i-1} by adding a single oriented edge e_i . A homotopy equivalence $f : G \rightarrow G$ of ϕ *respects the filtration* if for each non-fixed edge e_i , the path $f(e_i)$ has a splitting $f(e_i) = e_i \cdot u_i^{m_i}$ for some $m_i \in \mathbf{Z}$ and for some primitive closed path $u_i \subset G_{i-1}$ that is geodesic both as a path and as a loop. In particular, if e_i is non-fixed then its terminal vertex has valence at least two in G_{i-1} . It follows that the directions determined by the first two edges attached to a vertex $v \in G$ are fixed. If each u_i is a Nielsen path for f then we say that $f : G \rightarrow G$ is *UL*.

An element $\phi \in \text{Out}(F_n)$ has *linear growth* if it has infinite order and if the cyclic word length of $\phi^k([a])$ with respect to some, and hence any, fixed basis grows at most linearly in k for each $a \in F_n$. An element $\phi \in \text{Out}(F_n)$ is *unipotent* if its induced action on $H_1(F_n, \mathbf{Z})$ is unipotent. We say that ϕ is *UL* if it is unipotent and linear and that a *subgroup of $\text{Out}(F_n)$ is UL* if each of its elements is. It is an immediate consequence of the definitions that the outer automorphism determined by a UL homotopy equivalence is UL. Theorem 5.1.8 of [BFH3] implies that any UL ϕ is represented by a UL homotopy equivalence $f : G \rightarrow G$.

Let G be a filtered graph, let \mathcal{V} be the set of vertices of G and let $\text{FHE}(G, \mathcal{V})$ be the group (Lemma 6.2 of [BFH2]) of homotopy classes, relative to \mathcal{V} , of filtration-respecting homotopy equivalences of G . There is a natural homomorphism

$$\text{FHE}(G, \mathcal{V}) \rightarrow \text{Out}(F_n).$$

If a subgroup Q of $\text{Out}(F_n)$ lifts to a subgroup Q_G of $\text{FHE}(G, \mathcal{V})$, then we say that Q_G is a *Kolchin representative* of Q .

Recall (see, for example, Lemma 2.6 of [BFH3]) that if F is a free factor of F_n and $[F]$ is ϕ -invariant, then the restriction of ϕ to $[F]$ determines a well-defined outer automorphism $\phi|[F]$.

2.11. Proposition. — *Suppose that Q is a finitely generated UL subgroup of $\text{Out}(F_n)$ and that F is a (possibly trivial) ϕ -invariant free factor of F_n . Then Q has a Kolchin representative Q_G satisfying the following properties:*

- *There is a stratum G_m such that $[F] = [\pi_1(G_m)]$.*
- *If $\phi|[F]$ is trivial then G_m is Q_G -fixed; i.e. pointwise fixed by every element of Q_G .*

If Q is abelian then we may also assume the following.

- *The lift $f : G \rightarrow G$ to Q_G of $\phi \in Q$ is a UL representative of ϕ .*
- *If an edge e_i is not Q_G -fixed, then there is a nontrivial primitive closed path $u_i \subset G_{i-1}$ with basepoint equal to the terminal endpoint of e_i such that for all $f \in Q_G$, $f(e_i) = e_i u_i^{m_i(f)}$ for some $m_i(f) \in \mathbf{Z}$.*
- *If $[u_i]_u = [u_j]_u$ then $u_i = u_j$; in particular, the terminal endpoints of e_i and e_j are equal.*

Proof. — Theorem 1.1 of [BFH2] produces a Kolchin representative Q_G satisfying the first item. The second item is implicit in the construction of Q_G given on p. 57 of [BFH2]. The remaining items follows from Corollary 3.11 of [BFH3]. \square

Many arguments proceed by induction up the filtration of a UL representative $f : G \rightarrow G$ of ϕ . For any path $\sigma \subset G$ the *height* of σ is the smallest value of m for which $\sigma \subset G_m$.

2.4. Axes and multiplicity. — Suppose that $f : G \rightarrow G$ is a UL representative of ϕ and assume the usual notation that $f(e_i) = e_i u_i^{m_i}$ for each edge e_i . If u_i is nontrivial then we say that $\alpha = [u_i]_u$ is an *axis* for ϕ . If $\{e_j : j \in J\}$ is the set of edges with $[u_j]_u = \alpha$, then the *multiplicity of α with respect to ϕ* is the number of distinct nonzero values in $\{m_j : j \in J\}$.

Recall that the *centralizer* $C(H)$ of a subset $H \subset \text{Out}(F_n)$ is defined to be the subgroup of elements in $\text{Out}(F_n)$ that commute with every element of H .

2.12. Lemma. — *Suppose that $\phi \in \text{Out}(F_n)$ is UL.*

1. *The set of axes for ϕ and their multiplicities depend only on ϕ and not on the choice of UL representative.*
2. *If $[c]_u$ is an axis of ϕ with multiplicity m then $\psi([c]_u)$ is an axis of $\psi\phi\psi^{-1}$ with multiplicity m . In particular, each $\psi \in C(\phi)$ induces a multiplicity preserving permutation of the set of axes of ϕ .*
3. *If F is a ϕ -invariant free factor then $\phi|F$ is UL and each axis of $\phi|F$ is an axis of ϕ .*

Proof. — (1) is contained in Corollary 4.8 of [BFH3] and (2) is contained in Lemma 4.2 of [BFH3]. (3) follows from Proposition 2.11 and (1).

Remark. — Lemma 2.12 tells us that, in order to compute the axis of a UL element $\phi \in \text{Out}(F_n)$, it is enough to choose any UL representative for ϕ and compute its axis. We will do this numerous times (without further mention) throughout the paper.

We conclude this subsection with two examples.

2.13. Lemma. — *Suppose that $\{x_1, \dots, x_n\}$ is a basis for F_n , that $F_k = \langle x_1, \dots, x_k \rangle$, that $F_{n-k} = \langle x_{k+1}, \dots, x_n \rangle$ for some $1 \leq k \leq n-1$ and that $w \in F_k$ is primitive. If $\hat{\phi} = i_w^m \times \text{Id} \in \text{Aut}(F_k) \times \text{Aut}(F_{n-k}) \subset \text{Aut}(F_n)$ for some $m \neq 0$, then $[w]_u$ is the unique axis for ϕ and it has multiplicity one.*

Proof. — Let G be the graph with vertices v and v' and with edges X, e_1, \dots, e_n , where both ends of e_1, \dots, e_k and the terminal end of X are attached to v and all other ends of edges are attached to v' . The marking on G identifies e_i with x_i for $i > k$ and $Xe_i\bar{X}$ with x_i for $i \leq k$. The homotopy equivalence $f : G \rightarrow G$ defined by $X \mapsto Xw^m$ is a UL representative of ϕ and the lemma now follows from the definitions. \square

2.14. Lemma. — *Suppose that $\{x_1, \dots, x_n\}$ is a basis for F_n and that $w \in \langle x_1, x_2 \rangle$. For $3 \leq i \leq n$ define automorphisms $\hat{L}_{i,w}$ by $x_i \mapsto \bar{w}x_i$ and $\hat{R}_{i,w}$ by $x_i \mapsto x_iw$. Then*

1. *All elements of $\{\hat{L}_{i,w}\} \cup \{\hat{R}_{j,w}\} \cup \{\hat{L}_{i,w}\hat{R}_{j,w}, \hat{L}_{i,w}\hat{L}_{j,w}, \hat{R}_{i,w}\hat{L}_{j,w}, \hat{R}_{i,w}\hat{R}_{j,w} : i \neq j\}$ are conjugate.*
2. *If ϕ is any one of the elements of (1) then $[w]_u$ is the unique axis for ϕ and it has multiplicity one.*

Proof. — The automorphism defined by $x_i \mapsto \bar{x}_i$ conjugates $\hat{L}_{i,w}$ to $\hat{R}_{i,w}$ and vice-versa. The automorphism defined by $x_i \mapsto x_j$ and $x_j \mapsto x_i$ conjugates $\hat{L}_{i,w}$ to $\hat{L}_{j,w}$ and vice-versa. If $i \neq j$ then the automorphism defined by $x_j \mapsto x_j\bar{x}_i$ conjugates $\hat{R}_{i,w}\hat{L}_{j,w}$ to $\hat{R}_{i,w}$. Combining these moves completes the proof of (1).

If G is the rose with n edges e_1, \dots, e_n and if the marking identifies x_i with e_i , then $\hat{R}_{i,w}$ is realized by $f : G \rightarrow G$ where $f(e_i) = e_iw$ and where all other edges of G are fixed. This proves (2) for $\phi = \hat{R}_{i,w}$. Since the conjugating maps used in (1) preserve w , (2) follows. \square

2.5. Fixed subgroups. — Assume that $f : G \rightarrow G$ is a topological representative for $\phi \in \text{Out}(F_n)$.

If $x, y \in \text{Fix}(f)$ are the endpoints of a Nielsen path then they are *Nielsen equivalent* and belong to the same *Nielsen class* of fixed points. Equivalently x and y belong to the same Nielsen class if some, and hence every, lift $\tilde{f} : \Gamma \rightarrow \Gamma$ that fixes a lift \tilde{x} of x also

fixes a lift \tilde{y} of y . Each $x \in \text{Fix}(f)$ has contractible neighborhoods $V \subset U$ such that $f(V) \subset U$. It follows that all elements of $\text{Fix}(f) \cap V$ belong to the same Nielsen class and hence that there are only finitely many Nielsen classes.

If \tilde{f} is a lift of f and $\text{Fix}(\tilde{f}) \neq \emptyset$ then the projection of $\text{Fix}(\tilde{f})$ into G is an entire Nielsen class of $\text{Fix}(f)$. We say that \tilde{f} is a *lift for μ* and that μ is *the Nielsen class determined by \tilde{f}* . Another lift of f is also a lift for μ if and only if it equals $T\tilde{f}T^{-1}$ for some covering translation T .

If $b \in \text{Fix}(f)$, then there is an induced homomorphism $f_{\#} : \pi_1(G, b) \rightarrow \pi_1(G, b)$; we denote the fixed subgroup of this homomorphism by $\text{Fix}_b(f)$. Under the marking identification, $\text{Fix}_b(f)$ determines a conjugacy class $[\text{Fix}_b(f)]$ of subgroups in F_n . If b_1 and b_2 belong to the same Nielsen class in $\text{Fix}(f)$ then the Nielsen path that connects them provides an identification of $\text{Fix}_{b_1}(f)$ with $\text{Fix}_{b_2}(f)$. Thus $[\text{Fix}_b(f)]$ depends only on the Nielsen class of b .

Denote the fixed subgroup of an automorphism $\hat{\phi}$ by $\text{Fix}(\hat{\phi})$ and define

$$\text{Fix}(\phi) = \{[\text{Fix}(\hat{\phi})] : \hat{\phi} \text{ represents } \phi \text{ and } \text{rank}(\text{Fix}(\hat{\phi})) \geq 2\}.$$

2.15. Lemma. — *Suppose that $f : G \rightarrow G$ is a topological representative of $\phi \in \text{Out}(F_n)$. Then*

1. $\text{Fix}(\phi) = \{[\text{Fix}_{b_i}(f)] : b_i \in B\}$ where B contains one element for each Nielsen class of f whose associated (conjugacy class of) fixed subgroup has rank at least two.
2. $\text{Fix}(\phi)$ is finite.
2. Each $\psi \in C(\phi)$ permutes the elements of $\text{Fix}(\phi)$.

Proof. — The second item follows from the first and the third item follows from the observation that $\text{Fix}(\hat{\psi}\hat{\phi}\hat{\psi}^{-1}) = \hat{\psi}\text{Fix}(\hat{\phi})$. Corollary 2.2 of [BH] implies that each element of $\text{Fix}(\phi)$ is realized as $[\text{Fix}_b(f)]$ for some $b \in \text{Fix}(f)$. If $[\text{Fix}_{b_1}(f)] = [\text{Fix}_{b_2}(f)]$ then there is a path ρ connecting b_1 to b_2 such that $\rho\tau\bar{\rho}$ is a Nielsen path based at b_1 for each Nielsen path τ based at b_2 . The element $a \in \pi_1(G, b_2)$ determined by $\bar{\rho}f(\rho)$ is in the center of $\text{Fix}_{b_2}(f)$ and so is trivial. We conclude that ρ is a Nielsen path and hence that b_1 and b_2 belong to the same Nielsen class of $\text{Fix}(f)$. This completes the proof of the first item and so the lemma. \square

Remark. — If both $\hat{\phi}$ and $\hat{\phi}'$ represent $\phi \in \text{Out}(F_n)$, and if $\text{Fix}(\hat{\phi})$ and $\text{Fix}(\hat{\phi}')$ represent the same element of $\text{Fix}(\phi)$, then there exists $a \in F_n$ such that $\hat{\phi} = i_a\hat{\phi}'i_a^{-1}$. To see this, choose $a \in F_n$ so that $\text{Fix}(\hat{\phi}) = i_a\text{Fix}(\hat{\phi}')$. Then $\hat{\phi}$ and $i_a\hat{\phi}'i_a^{-1}$ agree on a subgroup of rank at least two and so are equal.

We next turn to the computation of $[\text{Fix}_b(f)]$.

Suppose that $f : G \rightarrow G$ is a UL representative of ϕ and that b is a vertex of G that is fixed by f . Denote the component of $\text{Fix}(f)$ that contains b by G_b and define

Σ_b to be the set of paths in G that can be written as a concatenation of subpaths, each of which is either contained in G_b or is of the form $e_i u_i^r \bar{e}_i$ for some $r \neq 0$ where e_i is a non-fixed edge with initial endpoint in G_b , u_i is a primitive closed path and $f(e_i) = e_i u_i^{m_i}$.

2.16. Lemma. — *Suppose that $f : G \rightarrow G$ is a UL representative of ϕ and that b is a vertex that is fixed by f . Assume further that if e_i and e_j are non-fixed edges with $[u_i]_u = [u_j]_u$ then $m_i \neq m_j$. Then $s \in \pi_1(G, b)$ is contained in $\text{Fix}_b(f)$ if and only if s is represented by a closed path in Σ_b based at b .*

Proof. — We have to show that a path σ with both endpoints at b is a Nielsen path if and only if $\sigma \in \Sigma_b$. The if direction is clear from the definitions.

By hypothesis, the number of non-fixed edges in G equals the sum of the multiplicities of the axes of ϕ and is therefore as small as possible. Assuming that e_i is a non-fixed edge, we apply this in two ways. The first is that there does not exist a path $\gamma \subset G_{i-1}$ such that $e_i \gamma$ is a Nielsen path. If there were such a path, then we could produce a new, more efficient UL representative $f' : G' \rightarrow G'$ of ϕ by the ‘sliding’ operation described in complete detail in Section 5.4 of [BFH1]. In this new representative the edge e_i is replaced by an edge e'_i that is marked so as to correspond to $e_i \gamma$. In particular e'_i is a fixed edge for $f' : G' \rightarrow G'$ and the total number of non-fixed edges would be decreased.

The second consequence, which we now prove, is that if $\gamma \subset G_{i-1}$ and if $e_i \gamma \bar{e}_i$ is a Nielsen path, then $\gamma = u_i^r$ for some $r \neq 0$. Choose a lift \tilde{e}_i to the universal cover Γ , let \tilde{x} be the initial endpoint of \tilde{e}_i , let \tilde{p} be the terminal endpoint of \tilde{e}_i , let $p \in G_{i-1}$ be the projected image of \tilde{p} and let $\tilde{f} : \Gamma \rightarrow \Gamma$ be the lift of $f : G \rightarrow G$ that fixes \tilde{x} . Let C be the component of G_{i-1} that contains p , let Γ_{i-1} be the component of the universal cover of C that contains \tilde{p} and let $h : \Gamma_{i-1} \rightarrow \Gamma_{i-1}$ be the restriction of \tilde{f} . There is a lift $\tilde{\gamma}$ of γ that begins at \tilde{p} . The covering translation $T : \Gamma \rightarrow \Gamma$ that sends \tilde{p} to the terminal endpoint of $\tilde{\gamma}$ sends \tilde{x} to the terminal endpoint \tilde{y} of the lift of $e_i \gamma \bar{e}_i$ that begins with $\tilde{e}_i \tilde{\gamma}$. Since $e_i \gamma \bar{e}_i$ is a Nielsen path for f and $\tilde{x} \in \text{Fix}(\tilde{f})$ it follows that $\tilde{y} \in \text{Fix}(\tilde{f})$ and hence that T commutes with \tilde{f} . Since T preserves Γ_{i-1} it restricts to a covering translation $T' : \Gamma_{i-1} \rightarrow \Gamma_{i-1}$ that commutes with h . It suffices to show that the subgroup of all such T' has rank one. If this fails, then $\text{Fix}(h) \neq \emptyset$ by Lemma 2.1 of [BH]. If $\gamma' \subset G_{i-1}$ is a path connecting \tilde{p} to an element of $\text{Fix}(h)$ then $e_i \gamma'$ is a Nielsen path for f . As we have already shown that this is impossible, we have verified our second consequence.

We can now prove the only if direction. It suffices to show that if σ is a Nielsen path with one endpoint in G_b then $\sigma \in \Sigma_b$. We will induct on the height of σ . Since $G_1 \subset \text{Fix}(f)$ the height 1 case is clear, and we may assume by induction that σ has height m and that the statement is true for paths with height less than m . By Lemma 4.1.4 of [BFH1], σ has a decomposition into Nielsen subpaths σ_i where each

σ_i or its inverse has the form $\gamma, e_m\gamma$ or $e_m\gamma\bar{e}_m$ for some path $\gamma \subset G_{m-1}$. As we have seen $\sigma_i = e_m\gamma$ can not occur and if $\sigma_i = e_m\gamma\bar{e}_m$ occurs then $\sigma_i \in \Sigma_b$. The case that $\sigma_i = \gamma$ follows from the inductive hypothesis and we have now completed the induction step. \square

We record the following example as a lemma for future reference.

2.17. *Lemma.* — Suppose that $\{x_1, \dots, x_n\}$ is a basis for F_n and that $\hat{\phi}$ is defined by $x_n \mapsto x_1^{-k} x_n x_1^k$ for some $k \neq 0$. Then

1. $\text{Fix}(\phi) = \{[\langle x_1, \dots, x_{n-1} \rangle], [\langle x_1, x_n \rangle]\}$.
2. $[x_1]_u, [\langle x_1, \dots, x_{n-1} \rangle]$ and $[\langle x_1, x_n \rangle]$ are ψ -invariant for all $\psi \in C(\phi)$.
3. Suppose that F is a free factor, that $[F]$ is ϕ -invariant and that $\phi|[F]$ is not the identity. Then F contains a representative of $[\langle x_1, x_n \rangle]$ and F has rank at least three.

Proof. — Let G be the graph with vertices v and w and with edges X, e_1, \dots, e_n , where both ends of e_n and the initial end of X are attached to w and all other ends of edges are attached to v . The marking on G identifies e_i to x_i for $i < n$ and $\bar{X}e_nX$ to x_n . The homotopy equivalence $f : G \rightarrow G$ defined by $X \mapsto Xe_1^k$ is a UL representative of ϕ . Lemmas 2.15 and 2.16 imply that $\text{Fix}(\phi) = \{[\langle x_1, \dots, x_{n-1} \rangle], [\langle x_1, x_n \rangle]\}$ is ψ -invariant for all $\psi \in C(\phi)$. Since the two elements of $\text{Fix}(\phi)$ have different ranks they are each ψ -invariant. Lemma 2.12 implies that $[x_1]_u$ is ψ -invariant. This completes the proof of (1) and (2).

A loop σ in G has a cyclic splitting into subpaths $\sigma = \sigma_1 \dots \sigma_r$ defined in three steps as follows. For $l \neq 0$, denote $Xe_1^l\bar{X}$ by τ^l . Any occurrence of τ^l as a subpath of σ defines a σ_i ; each of these subpaths is a Nielsen path based at w . In the complementary subpaths, each maximal length subpath of the form Xe_1^j or $e_1^j\bar{X}$ for some integer j is a σ_i . The third step is to define each remaining edge to be a σ_i ; each of these subpaths is a Nielsen path based at v . Thus $f_{\#}^m(\sigma)$ is obtained from σ by replacing each Xe_1^j with Xe_1^{j+mk} and each $e_1^j\bar{X}$ with $e_1^{j-mk}\bar{X}$. If σ is a loop whose free homotopy class is not fixed by f then there is at least one σ_i of the form Xe_1^j and at least one of the form $e_1^j\bar{X}$. These can be chosen to be separated in σ by a Nielsen path μ based at w . Thus $f_{\#}^m(\sigma)$ contains $e_1^{-kn+p}\bar{X}\mu Xe_1^{kn+q}$ as a subpath for all $m \geq 0$ and some $p, q \in \mathbf{Z}$.

Carrying this back to ϕ and F_n via the marking and taking limits, we conclude that if $\phi|[F]$ is not the identity then F carries a bi-infinite nonperiodic word in $\langle x_1, x_n \rangle$. Corollary 2.8 implies that F contains a representative of $[\langle x_1, x_n \rangle]$. Since $\phi|[\langle x_1, x_n \rangle]$ is trivial, F must properly contain $[\langle x_1, x_n \rangle]$ and so must have rank at least three. \square

2.6. *Dehn twists.* — The group $\text{Out}(F_2)$ plays a special role in understanding $\text{Out}(F_n)$. One of the reasons for this is that, as shown by Nielsen, it can be understood via surface topology.

The once-punctured torus S is homotopy equivalent to the rose R_2 , so we may assume that S is marked. Recall that the (extended) *mapping class group* $\text{Mod}^\pm(S)$ of S is the group of homotopy classes of homeomorphisms of S . It is well known that the natural homomorphism $\text{Mod}^\pm(S) \rightarrow \text{Out}(F_2)$ given by the action of $\text{Mod}^\pm(S)$ on $\pi_1(S) \approx F_2$ is an isomorphism. It is also well known that there is a bijective correspondence between the set \mathcal{S} of isotopy classes of essential, nonperipheral (i.e. not isotopic to the puncture) simple closed curves on S and the set \mathcal{C} of unoriented conjugacy classes of basis elements of F_2 . Recall that a *Dehn twist* about a simple closed curve α in S is defined as the element of $\text{Mod}^\pm(S)$ represented by cutting S along α , twisting one of the resulting boundary circles by a complete rotation, and regluing.

2.18. Lemma. — *The following are equivalent.*

- $\phi_1 \in \text{Out}(F_2)$ is UL.
- There is a basis $\{z_1, z_2\}$ of F_2 and $b > 0$ so that $z_2 \mapsto z_2 z_1^b$ defines a representative of ϕ_1 .
- ϕ_1 corresponds to a Dehn twist of the once-punctured torus about the simple closed curve represented by $[z_1]$.

Proof. — This is immediate from the definitions and the fact that every UL outer automorphism is represented by a UL homotopy equivalence. \square

2.19. Corollary. — *Suppose that $\{x_1, x_2\}$ is a basis for F_2 and that $\hat{E}_{21} \in \text{Aut}(F_2)$ is defined by $x_2 \mapsto x_2 x_1$. Let $\rho = [x_1, x_2]$. Then for any $k \neq 0$:*

1. $[x_1^\pm]$ are the only E_{21}^k -invariant conjugacy classes represented by basis elements of F_2 .
2. If $\mu \in \text{Out}(F_2)$ has infinite order and if there is a conjugacy class $[a] \neq [\rho^l]$ that is fixed both by μ and by E_{21}^k , then $\mu^2 \in \langle E_{21} \rangle$.
3. Elements of $\text{Fix}(\hat{E}_{21}^k)$ that are conjugate in F_2 are conjugate in $\text{Fix}(\hat{E}_{21}^k)$.

Proof. — By Lemma 2.18, the mapping class element θ determined by E_{21} is represented by a Dehn twist $f : S \rightarrow S$ about a simple closed curve β that corresponds to $[x_1]$. The complement S' of an open annulus neighborhood of β is topologically a 3-times punctured sphere. The free homotopy class of a closed curve is fixed by θ^k if and only if it is represented by a closed curve in S' . Part (1) now follows from the fact that a basis element is represented by a simple closed curve in S and the fact that the only simple closed curves in S' are peripheral.

There is an orientation-preserving homeomorphism $h : S \rightarrow S$ whose mapping class ν corresponds to μ^2 . The Thurston classification theorem implies that ν preserves the free homotopy class of some simple closed curve β' and that $[a]$ is represented by a closed curve that is disjoint from β' and by a closed curve that is disjoint from β . It follows that $\beta = \beta'$ and that $\nu \in \langle \theta \rangle$. This proves (2).

Part (3) follows from the fact that closed curves of S' that are freely homotopic in S are also freely homotopic in S' . \square

We will also make use of the following.

2.20. Lemma. — *Suppose that $\hat{\phi}_1 \in \text{Aut}(F_2)$ is nontrivial and that $\text{Fix}(\hat{\phi}_1)$ has rank bigger than one. Then there exists $s > 0$ and there exists some basis $\{z_1, z_2\}$ of F_2 in which $\hat{\phi}_1$ is defined by $z_2 \mapsto z_2 z_1^s$.*

Proof. — We view ϕ_1 as an element of the mapping class group of the once punctured torus S . It is well known that $\text{Fix}(\hat{\phi})$ corresponds to the fundamental group of a proper essential subsurface S_0 and that there exists a homomorphism $h : S \rightarrow S$ representing ϕ such that $h|_{S_0}$ is the identity. Thus S_0 has rank two and is the complement of an open annulus neighborhood of a simple closed curve α . Up to isotopy, h must be a Dehn twist of nonzero order s about α . Lemma 2.18, Lemma 2.15 and Remark 2.5 complete the proof. \square

3. The endgame

For the remainder of this paper, Γ will denote an arbitrary finite index subgroup of $\text{Out}(F_n)$ and $\Phi : \Gamma \rightarrow \text{Out}(F_n)$ will be an arbitrary injective homomorphism.

In this section we prove that Theorem 1.1 can be reduced to understanding the image under Φ of the so-called elementary outer automorphisms. We then prove that Theorem 1.1 implies the corollaries stated in the introduction. Having dispatched with these necessities, we can then proceed with the heart of the argument of Theorem 1.1, which occupies the remainder of the paper.

3.1. Reduction to the action on elementary automorphisms. — Given a basis x_1, \dots, x_n for F_n , define for $j \neq k$ automorphisms \hat{E}_{jk} and ${}_{kj}\hat{E}$ by

$$\begin{aligned} \hat{E}_{jk} &: x_j \mapsto x_j x_k \\ {}_{kj}\hat{E} &: x_j \mapsto \bar{x}_k x_j. \end{aligned}$$

The elements of $\text{Out}(F_n)$ determined by these automorphisms will be denoted by E_{jk} and ${}_{kj}E$, respectively. Lemma 2.14 implies that $[x_k]_u$ is the unique axis of any iterate of E_{jk} or of ${}_{kj}E$, and that the multiplicity is one in each case.

3.1. Definition (Elementary Automorphism). — *A nontrivial element $\phi \in \text{Out}(F_n)$ is called elementary if there exists a choice of basis for F_n so that in this basis the element ϕ is an iterate of either E_{jk} or of ${}_{kj}E$ for some $j \neq k$.*

Since the set of bases is $\text{Aut}(F_n)$ -invariant, the set of elementary elements of $\text{Out}(F_n)$ is invariant under the conjugation action of $\text{Out}(F_n)$ on itself.

For any $\psi \in \text{Out}(F_n)$ we say that the injective homomorphism $i_\psi \circ \Phi : \Gamma \rightarrow \text{Out}(F_n)$ is a *normalization* of Φ . We say that $\phi \in \text{Out}(F_n)$ is *almost fixed* by Φ if there exists $s, t > 0$ such that $\Phi(\phi^s) = \phi^t$. If there exists $s, t > 0$ such that $\Phi(\phi^s) = \phi^t$ for every ϕ in a subgroup then we say that the *subgroup is almost fixed*.

Our strategy in proving Theorem 1.1 is to show that Φ has a normalization that almost fixes each elementary element of $\text{Out}(F_n)$. The following lemma, based on an argument of Ivanov in the context of mapping class groups (see Section 8.5 of [Iv]), shows that this is sufficient.

3.2. Lemma (Action on elementaries suffices). — *Let $\Gamma < \text{Out}(F_n)$, $n \geq 3$ be any finite index subgroup, and let $\Phi : \Gamma \rightarrow \text{Out}(F_n)$ be any injective homomorphism. If Φ has a normalization that almost fixes every elementary element of $\text{Out}(F_n)$, then there exists $g \in \text{Out}(F_n)$ such that $\Phi(\gamma) = g\gamma g^{-1}$ for all $\gamma \in \Gamma$.*

Proof. — It clearly suffices to show that if Φ almost fixes each elementary element of $\text{Out}(F_n)$, then Φ restricted to Γ is the identity. Given any $\phi \in \Gamma$, let $\eta = \phi^{-1}\Phi(\phi)$. Given any basis element x_1 , extend x_1 to a basis $\{x_1, \dots, x_n\}$. The assumption that Φ almost fixes every elementary outer automorphism gives that, for some $s, t, u, v > 0$,

$$\Phi(E_{21}^u) = E_{21}^v$$

and

$$(1) \quad \Phi(\phi)\Phi(E_{21}^s)\Phi(\phi)^{-1} = \Phi((\phi E_{21}\phi^{-1})^s) = (\phi E_{21}\phi^{-1})^t = \phi E_{21}^t \phi^{-1}.$$

Equation (1) implies that

$$(\eta\Phi(E_{21}^s)\eta^{-1})^u = E_{21}^{tu}$$

and so

$$\eta E_{21}^{sv} \eta^{-1} = \eta \Phi(E_{21}^u)^s \eta^{-1} = \eta \Phi(E_{21}^s)^u \eta^{-1} = (\eta \Phi(E_{21}^s) \eta^{-1})^u = E_{21}^{tu}.$$

By Lemma 2.14(b), we have that $[x_1]_u$ is the unique axis for E_{21}^{sv} and for E_{21}^{tu} . Lemma 2.12 then implies that η takes $[x_1]_u$ to a power of itself. As η clearly preserves the property of being primitive, it follows that η fixes $[x_1]_u$. As x_1 was arbitrary, the following lemma then completes the proof. \square

3.3. Lemma. — *If $\phi \in \text{Out}(F_n)$ fixes $[x]_u$ for each basis element x , then ϕ is the identity.*

Proof. — Corollary 2.9 implies that every free factor of F_n is ϕ -invariant. Choose a basis $\{x_1, \dots, x_n\}$ for F_n . By Corollary 2.5 there is an automorphism $\hat{\phi}$ representing ϕ such that $\langle x_1, x_2 \rangle$ is $\hat{\phi}$ -invariant and such that $\hat{\phi}(x_1) = x_1^\pm$. Lemma 2.10 implies that $\hat{\phi}(x_2) = x_1^j x_2^\pm x_1^k$ for some $j, k \in \mathbf{Z}$. By hypothesis $j + k = 0$, so after replacing $\hat{\phi}$ with $i_{x_1}^k \hat{\phi}$, we may assume that $\hat{\phi}(x_1) = x_1^\pm$ and $\hat{\phi}(x_2) = x_2^\pm$.

We now claim that $\hat{\phi}(x_i) = x_i^\pm$ for all $1 \leq i \leq n$. Assume by induction that the claim is true for $i = m - 1$ with $m \geq 3$. By hypothesis $\hat{\phi}(x_m) = wx_m^\pm \bar{w}$ for some $w \in F_n$. Either x_1^\pm or x_2^\pm , say x_1^\pm , is not the first letter of w . Then $\hat{\phi}(x_1 x_m) = x_1^\pm wx_m^\pm \bar{w}$ is cyclically reduced and, unless w is trivial, does not cyclically reduce to $(x_1 x_m)^\pm 1$, as it should by assumption since $x_1 x_m$ is a basis element. Thus w must be trivial, completing the proof of the claim.

For any distinct i, j, k we have that

$$\hat{\phi}(x_i x_j x_k) = x_i^\pm x_j^\pm x_k^\pm.$$

On the other hand, since $x_i x_j x_k$ is a basis element, we also have that $\hat{\phi}(x_i x_j x_k)$ is conjugate either to $x_i x_j x_k$ or to $\bar{x}_k \bar{x}_j \bar{x}_i$. As the latter clearly cannot occur, it follows that $\hat{\phi}(x_i) = x_i$ for each i . \square

3.2. Proofs of the corollaries to Theorem 1.1. — We now give short arguments to show how to derive the other claimed results in the introduction from Theorem 1.1.

Proof of Corollary 1.2. — The given map is clearly a homomorphism. Its kernel is precisely the centralizer $C(\Gamma)$ of Γ in $\text{Out}(F_n)$. Since Γ contains an iterate of each element of $\text{Out}(F_n)$ Lemma 3.2 implies that the map is injective. Surjectivity is immediate from Theorem 1.1.

Proof of Corollary 1.4. — The proof here is essentially the same as that of Corollary 1.2 just given. One need only remark that, by definition, an element $f \in \text{Out}(F_n)$ is trivial in $\text{Comm}(\text{Out}(F_n))$ precisely when Conj_f is the identity when restricted to some finite index subgroup $\Gamma \leq \text{Out}(F_n)$. This happens precisely when f centralizes Γ , which by Lemma 3.2 happens only when f is the identity.

4. A commensurability invariant

In this section we introduce and compute a commensurability invariant which will be crucial for understanding Φ . An analogous invariant for the mapping class group was studied by Ivanov-McCarthy in [IM]. We assume that $\{x_1, \dots, x_n\}$ is a basis for F_n , and we denote $\langle x_1, x_2 \rangle$ by F_2 and $\langle x_3, \dots, x_n \rangle$ by F_{n-2} .

4.1. *The invariant $r(\phi, \mathcal{A})$.* — Recall that the *centralizer* $C(H)$ of a subset $H \subseteq \Gamma$ is the subgroup of Γ consisting of elements commuting with every element of H . The *center* $Z(\Gamma)$ is the group of elements commuting with every element of Γ . We will need coarse versions of these basic group-theoretic notions.

4.1. *Definition (Weak center and centralizers).* — We define the weak centralizer of a subset $H \subseteq \text{Out}(F_n)$ to be the subgroup $\text{WC}(H) < \text{Out}(F_n)$ consisting of those $g \in \text{Out}(F_n)$ with the property that for each $h \in H$ there exists $s \neq 0$ so that g commutes with h^s . We define the weak center of H , denoted by $\text{WZ}(H)$, to be

$$\text{WZ}(H) := \text{WC}(H) \cap H.$$

By the *rank* of an abelian subgroup we will mean the rank of its free abelian direct factor. It is easy to see that any automorphism $\Phi^* : \text{Out}(F_n) \rightarrow \text{Out}(F_n)$ preserves centers of centralizers; that is, for each $\phi \in \text{Out}(F_n)$ we have $\Phi^*(Z(C(\phi))) = Z(C(\Phi^*(\phi)))$. In particular, $\text{rank}(Z(C(\phi))) = \text{rank}(Z(C(\Phi^*(\phi))))$. This is not obvious if Φ^* is replaced by an arbitrary injective homomorphism $\Phi : \Gamma \rightarrow \text{Out}(F_n)$ of a finite index subgroup Γ of $\text{Out}(F_n)$. In place of $\text{rank}(Z(C(\phi)))$ we use the following invariant.

For any abelian subgroup $\mathcal{A} \leq \text{Out}(F_n)$ and any $\phi \in \mathcal{A}$ define

$$r(\phi, \mathcal{A}) := \text{rank}(\mathcal{A} \cap \text{WZ}(C(\phi))).$$

Note that if \mathcal{A} is infinite and if ϕ has infinite order then $r(\phi, \mathcal{A}) \geq 1$. We are particularly interested in the case that $r(\phi, \mathcal{A}) = 1$. The following lemma states that Φ preserves pairs with this property.

4.2. *Lemma.* — Let $\Gamma \subseteq \text{Out}(F_n)$ be any finite index subgroup, and let $\mathcal{A} \subset \text{Out}(F_n)$ be any abelian subgroup. If $\phi \in \mathcal{A} \cap \Gamma$ then $r(\Phi(\phi), \Phi(\mathcal{A} \cap \Gamma)) \leq r(\phi, \mathcal{A})$. In particular, if $r(\phi, \mathcal{A}) = 1$ then $r(\Phi(\phi), \Phi(\mathcal{A} \cap \Gamma)) = 1$.

Proof. — Since $\mathcal{A} \cap \Gamma$ has finite index in \mathcal{A} it is clear that $r(\phi, \mathcal{A} \cap \Gamma) = r(\phi, \mathcal{A})$. Thus without loss of generality we can assume that $\mathcal{A} \subset \Gamma$. If $\psi \in \mathcal{A}$ and $\psi \notin \text{WZ}(C(\phi))$ then there exists $\mu \in C(\phi)$ such that ψ does not commute with any iterate of μ . We can clearly assume that $\mu \in \Gamma$. Thus $\Phi(\psi)$ does not commute with any iterate of $\Phi(\mu) \in C(\Phi(\phi))$, which implies that $\Phi(\psi) \notin \text{WZ}(C(\Phi(\phi)))$. This proves that the Φ -image of $\text{WZ}(C(\phi)) \cap \mathcal{A}$ contains $\text{WZ}(C(\Phi(\phi))) \cap \Phi(\mathcal{A})$ and the lemma follows. \square

4.2. *The subgroup $O(F_2, F_{n-2})$.* — Define a subgroup $O(F_2, F_{n-2})$ of $\text{Out}(F_n)$ by

$$O(F_2, F_{n-2}) := \{\phi \in \text{Out}(F_n) : \text{both } [F_2] \text{ and } [F_{n-2}] \text{ are } \phi\text{-invariant}\}.$$

The natural inclusion of $\text{Aut}(F_2)$ into $\text{Aut}(F_n)$ given by $\hat{\theta}_1 \mapsto \hat{\theta}_1 \times \text{Id} \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2}) \subset \text{Aut}(F_n)$ defines an embedding

$$\text{Aut}(F_2) \hookrightarrow \text{O}(F_2, F_{n-2})$$

whose image we denote by $\text{O}(F_2)$. Define $\text{O}(F_{n-2})$ similarly using the natural inclusion of $\text{Aut}(F_{n-2})$ into $\text{Aut}(F_n)$. Each element of $\text{O}(F_2)$ commutes with each element of $\text{O}(F_{n-2})$.

4.3. Lemma. — *Let notation be as above. Then:*

1. $\text{O}(F_2, F_{n-2}) \cong \text{O}(F_2) \times \text{O}(F_{n-2}) \cong \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$.
2. If $\phi \in \text{O}(F_2)$ then $\text{WZ}(\text{C}(\phi)) \subset \text{O}(F_2)$.

Proof. — The natural homomorphism $\text{Aut}(F_2) \times \text{Aut}(F_{n-2}) \rightarrow \text{Aut}(F_n)$ induces an injection

$$\text{Aut}(F_2) \times \text{Aut}(F_{n-2}) \hookrightarrow \text{O}(F_2, F_{n-2}).$$

To prove the first item it suffices to show this injection is onto.

Each $\eta \in \text{O}(F_2, F_{n-2})$ is (non-uniquely) represented by an automorphism $\hat{\eta}$ that leaves F_2 invariant. Define $\hat{\mu} = \hat{\eta}|_{F_2} \times \text{Id} \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$. There is no loss in replacing η by $\mu^{-1}\eta$ so we may assume that $\eta|_{F_2}$ determines the trivial element of $\text{Out}(F_2)$. Thus $\hat{\eta}|_{F_2} = i_a$ for some $a \in F_2$. By the symmetric argument we may assume that $\eta|_{F_{n-2}}$ is trivial and hence that $\hat{\eta}|_{F_{n-2}} = i_w$ for some $w \in F_n$. (We cannot assume that $w \in F_{n-2}$ because we do not yet know that F_{n-2} is $\hat{\eta}$ -invariant.) If there is a nontrivial initial segment \hat{a} of w that belongs to F_2 then replace $\hat{\eta}$ by $i_{\hat{a}}^{-1}\hat{\eta}$. Thus $w = b_1 a_1 b_2 \dots$ is an alternating concatenation where $b_i \subset F_{n-2}$ and $a_i \subset F_2$.

By the same argument, there is a representative $\hat{\mu}$ of $\mu = \eta^{-1}$ such that $\hat{\mu}|_{F_2} = i_{a'}$ for some $a' \in F_2$ and $\hat{\mu}|_{F_{n-2}} = i_v$ for some $v \in F_n$ that begins in F_{n-2} . Since $\hat{\mu}\hat{\eta}|_{F_2}$ is conjugation by a (possibly trivial) element of F_2 , the same must be true for $\hat{\mu}\hat{\eta}|_{F_{n-2}} = i_{\hat{\mu}(w)v}$ which implies that $\hat{\mu}(w)v \in F_2$. Letting $\#$ stand for the reducing operation, we have

$$(\hat{\mu}(w))_{\#} = (i_v(b_1))_{\#}(i_{a'}(a_1))_{\#}(i_v(b_2))_{\#}\dots$$

where each $(i_{a'}(a_1))_{\#} \in F_2$ is nontrivial and each $(i_v(b_i))_{\#}$ is nontrivial and begins and ends in F_{n-2} . If w ends with an a_l then $(\hat{\mu}(w)v)_{\#} = \hat{\mu}(w)_{\#}v$ in contradiction to the fact that $\hat{\mu}(w)v \in F_2$. Thus w ends with b_l and

$$(\hat{\mu}(w)v)_{\#} = (i_v(b_1))_{\#}(i_{a'}(a_1))_{\#}\dots(i_{a'}(a_l))_{\#}(i_v(b_l)v)_{\#}.$$

It follows that $l = 1$ and that $w = v^{-1} = b_1 \in F_{n-2}$. Thus $\hat{\eta} = i_a \times i_{b_1} \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$ which completes the proof of (1).

Suppose now that $\phi \in \text{O}(F_2)$ and that $\psi \in \text{WZ}(\text{C}(\phi))$. Choose $\hat{\mu}_2 \in \text{Aut}(F_{n-2})$ so that $\text{Fix}(\hat{\mu}_2^k)$ and $\text{Fix}(\mu_2^k)$ are trivial for all $k > 0$. For example, μ can be represented by a pseudo-Anosov homeomorphism $h : S \rightarrow S$ of a surface with boundary and $\hat{\mu}$ can be the automorphism of $\pi_1(S, b)$ determined by h at a fixed point b in the interior of S . Let $\hat{\mu} = \text{Id} \times \hat{\mu}_2$. Then $\text{Fix}(\hat{\mu}^k) = F_2$ and $\text{Fix}(\mu^k) = \{[F_2]\}$ for all $k > 0$. Since μ is an element of $\text{O}(F_{n-2})$, it commutes with ϕ and ψ commutes with some μ^k . Lemma 2.15 implies that F_2 is ψ -invariant.

Choose $w \in F_{n-2}$ and define $\hat{\eta} = \text{Id} \times i_w \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$. Then ϕ commutes with η and so ψ commutes with some η^k . Lemma 2.13 and Lemma 2.12 imply that $[w]_u$ is ψ -invariant. Since w is arbitrary, Corollary 2.9 implies that F_{n-2} is ψ -invariant. By (1), ψ has a representation of the form $\hat{\psi} = \hat{\psi}_1 \times \hat{\psi}_2 \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$. Since ψ commutes with η^k and $\hat{\eta}|_{F_2} = \text{Id}$, $\hat{\psi}_2$ commutes with i_w^k . It follows that $\hat{\psi}_2$ fixes w for all w and so is the identity. Thus $\psi \in \text{O}(F_2)$ as desired. \square

Notation. — Each $\phi \in \text{O}(F_2, F_{n-2})$ is represented by a unique $\hat{\phi}$ that preserves both F_2 and F_{n-2} . The restrictions $\hat{\phi}|_{F_2}$ and $\hat{\phi}|_{F_{n-2}}$ are denoted $\hat{\phi}_1$ and $\hat{\phi}_2$.

Remark. — If $\phi, \psi \in \text{O}(F_2, F_{n-2})$ then ϕ commutes with ψ if and only if $\hat{\phi}_1$ commutes with $\hat{\psi}_1$ and $\hat{\phi}_2$ commutes with $\hat{\psi}_2$.

4.3. *Calculating $\text{WZ}(\text{C}(\phi))$.* — Our first calculation is related to Lemma 2.14. We change the notation from that lemma to make it more consistent with future applications. Suppose that $w \in F_2$. For $3 \leq l \leq n$ we define automorphisms

$$\begin{aligned} \hat{\mu}_{2l-5,w} : x_1 &\mapsto \bar{w}x_l \\ \hat{\mu}_{2l-4,w} : x_l &\mapsto x_l w. \end{aligned}$$

We say that $\hat{\mu}_{2l-5,w}$ and $\hat{\mu}_{2l-4,w}$ are *paired*. In the notation of Lemma 2.14, $\mu_{i,w}$ for odd values of i corresponds to an $L_{j,w}$, and $\mu_{i,w}$ for even values of i corresponds to an $R_{j,w}$.

4.4. *Lemma.* — Suppose that s and t are nonzero and that $w \in F_2$ is primitive.

1. If $\phi = \mu_{i,w}$, or if $\phi = \mu_{i,w}\mu_{j,w}$ where $\mu_{i,w}$ and $\mu_{j,w}$ are unpaired, then $\text{WZ}(\text{C}(\phi^s)) = \langle \phi \rangle$ for any $s \neq 0$.
2. If $\mu_{i,w}$ and $\mu_{j,w}$ are paired, or if $s \neq t$, then $\text{WZ}(\text{C}(\mu_{i,w}^s \mu_{j,w}^t)) \supset \langle \mu_{i,w}, \mu_{j,w} \rangle$.

Proof. — All of the ϕ considered in (1) are conjugate by Lemma 2.14. We may therefore assume, for (1), that $\hat{\phi}$ is defined by $x_n \mapsto x_n w$.

For any $y \in \langle x_1, \dots, x_{n-1} \rangle$ define $\hat{\theta}_y$ by $x_n \mapsto yx_n$. Then ϕ^s commutes with every θ_y . If $\psi \in \text{WZ}(\text{C}(\phi^s))$ then ψ commutes with θ_y^p for some $p > 0$. Lemma 2.12 implies that every $[y]_u$ is ψ -invariant. Corollary 2.9 then implies that $\langle x_1, \dots, x_{n-1} \rangle$ is ψ -invariant

and Lemma 3.3 implies that $\psi|[\langle x_1, \dots, x_{n-1} \rangle]$ is the identity. Lemma 2.10 implies that ψ is represented by $\hat{\psi}$ defined by $x_n \mapsto ux_nv$ where $u, v \in \langle x_1, \dots, x_{n-1} \rangle$.

Since ψ commutes with both θ_y^p and ϕ^s , and since $\hat{\psi}, \hat{\theta}_y^p$ and $\hat{\phi}^s$ agree on subgroup of rank bigger than one, $\hat{\psi}$ commutes with $\hat{\theta}_y^p$ and $\hat{\phi}^s$. Direct computation now shows that u is trivial and $v \in \langle w \rangle$. Thus $\psi \in \langle \phi \rangle$ as desired. This completes the proof of (1).

Theorem 6.8 of [FH] imply (2) in the case that $s \neq t$. It remains to consider the case that $s = t$ and that $\mu_{i,w}$ and $\mu_{j,w}$ are paired. There is no loss in assuming that $\hat{\mu}_{i,w}$ is defined by $x_n \mapsto \bar{w}x_n$ and $\hat{\mu}_{j,w}$ is defined by $x_n \mapsto x_nw$. Thus $\hat{\eta}$ is defined by $x_n \mapsto \bar{w}^s x_n w^s$. An argument exactly like that given in the proof of Lemma 2.17(1) shows that $\text{Fix}(\eta)$ has two elements, one represented by $\langle x_1, x_2, \dots, x_{n-1} \rangle$ and the other by $\langle w, x_n \rangle$. If $\theta \in \text{C}(\eta)$, then $[\langle x_1, x_2, \dots, x_{n-1} \rangle], [\langle w, x_n \rangle]$ and $[w]_u$ are θ -invariant. After replacing θ by θ^2 if necessary, there is an automorphism $\hat{\theta}$ representing θ such that $\hat{\theta}(w) = w$. Corollary 2.5 and Lemma 2.10 imply that $\langle x_1, x_2, \dots, x_{n-1} \rangle$ is $\hat{\theta}$ -invariant and that $\hat{\theta}(x_n) = ux_nv$ for some $u, v \in \langle x_1, x_2, \dots, x_{n-1} \rangle$. Since θ commutes with η and the restrictions of $\hat{\theta}$ and $\hat{\eta}$ to $\langle x_1, x_2, \dots, x_{n-1} \rangle$ commute, $\hat{\theta}$ and $\hat{\eta}$ commute. Since

$$\hat{\theta}\hat{\eta}(x_n) = \hat{\theta}(\bar{w}^s x_n w^s) = \bar{w}^s ux_nv w^s$$

and

$$\hat{\eta}\hat{\theta}(x_n) = \hat{\eta}(ux_nv) = u\bar{w}^s x_n w^s v$$

it follows that $u, v \in \langle w \rangle$ which implies that $\hat{\theta}$ commutes with $\hat{\mu}_{i,w}$ and $\hat{\mu}_{j,w}$. \square

4.5. Definition (Twists). — For $w \in \mathbb{F}_2$, define $T_w \in \text{O}(\mathbb{F}_2)$ by $\hat{T}_w = i_w \times \text{Id}$.

4.6. Lemma. — $\text{Fix}(T_w) = \{[\mathbb{F}_2], [\langle \mathbb{F}_{n-2}, w \rangle]\}$.

Proof. — Let G be the graph with vertices v and v' , with edges e_1, e_2 attached to v , edges e_3, \dots, e_n attached to v' and an edge X with initial endpoint at v' and terminal endpoint at v . The homotopy equivalence $f : G \rightarrow G$ by $f(X) = Xw$ is a UL representative of T_w , and the lemma now follows from Lemma 2.15 and Lemma 2.16. \square

We say that $\rho \in \mathbb{F}_2$ is *peripheral* if it is the commutator of two basis elements. We think of T_ρ as a Dehn twist about a peripheral curve on a once-punctured torus representing \mathbb{F}_2 in the decomposition $\mathbb{F}_2 * \mathbb{F}_{n-2}$.

4.7. Lemma. — If $\rho \in \mathbb{F}_2$ is peripheral then $\langle T_\rho \rangle$ has finite index in $\text{WZ}(\text{C}(T_\rho))$.

Proof. — The group $\text{WZ}(\text{C}(T_\rho))$ has a torsion free subgroup of finite index so it suffices to show that each infinite order $\psi \in \text{WZ}(\text{C}(T_\rho))$ is an iterate of T_ρ . By

Lemma 4.3, ψ is represented by $\hat{\psi}_1 \times \text{Id} \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$. Every $\phi_1 \in \text{Out}(F_2)$ has a representative $\hat{\phi}_1 \in \text{Aut}(F_2)$ that fixes ρ ; this is because any two peripheral elements of F_2 are conjugate in F_2 . The outer automorphism represented by $\hat{\phi}_1 \times \text{Id}$ is an element of $C(T_\rho)$. Thus $\hat{\phi}_1^k$ commutes with $\hat{\psi}_1$ for some $k > 0$. This proves that ψ_1 commutes with an iterate of every element of $\text{Out}(F_2)$ and, having infinite order, is therefore trivial. In other words $\hat{\psi}_1 = i_w$ for some $w \in F_2$. Since $\hat{\psi}_1$ commutes with i_ρ^k , we have $w \in \langle \rho \rangle$ as desired. \square

4.8. Lemma. — *If $w \in F_2$ is a nontrivial nonperipheral element of $\text{Fix}(\hat{E}_{21})$ then $E_{21} \in \text{WZ}(C(T_w))$.*

Proof. — We must show that some iterate of each $\theta \in C(T_w)$ commutes with E_{21} . Lemma 2.12 and Lemma 4.6 imply that $[w]_a$, $[F_2]$ and $H := [(F_{n-2}, w)]$ are θ -invariant. After replacing θ with θ^2 if necessary there exists $\hat{\theta}$ representing θ that fixes w . Lemma 2.5 implies that F_2 is $\hat{\theta}$ -invariant. Corollary 2.19 implies that $\theta|_{F_2}$ is an iterate of $E_{21}|_{F_2}$ and hence that $\hat{\theta}|_{F_2} = i_w^p \hat{E}_{21}^m|_{F_2}$ for some $m, p \neq 0$. In particular, $\hat{\theta}|_{F_2}$ commutes with $\hat{E}_{21}|_{F_2}$.

There exists $c \in F_n$ such that $w \in \hat{\theta}(H) = i_c(H)$. Theorem 2.1 implies that $\bar{c} = ha$ for some $h \in H$ and $a \in F_2$. Thus $w \in i_a(H) \cap F_2$ which implies that $i_a w \in H \cap F_2 = \langle w \rangle$. It follows that $a \in \langle w \rangle$ and hence that $\hat{\theta}(H) = H$. Since $\hat{E}_{21}|_H$ is the identity, it commutes with $\hat{\theta}|_H$. As we have already seen that $\hat{E}_{21}^m|_{F_2}$ commutes with $\hat{\theta}|_{F_2}$, we conclude that $\hat{\theta}$ commutes with \hat{E}_{21}^m . \square

5. The action on special abelian subgroups

To obtain constraints on the injective homomorphism $\Phi : \Gamma \rightarrow \text{Out}(F_n)$ we will consider two special families of abelian subgroups of $\text{Out}(F_n)$, one of rank $2n - 3$ and one of rank $2n - 4$. We will use [FH] to isolate properties which characterize such subgroups and at the same time are preserved by Φ .

To fix notation, we let $\{x_1, \dots, x_n\}$ be a basis for F_n , denote the group $\langle x_1, x_2 \rangle$ by F_2 , denote the group $\langle x_3, \dots, x_n \rangle$ by F_{n-2} , and denote $[x_1, x_2]$ by ρ . The following definition is relevant to both special families of abelian subgroups we will study.

If $\mathcal{A} < \text{Out}(F_n)$ is an abelian subgroup, we say that a set of elements $\{\phi_1, \dots, \phi_{2n-4}\} \subset \mathcal{A}$ satisfies the *pairing property for \mathcal{A}* if the following two conditions hold for all $m \neq 0$:

1. $r(\phi_j^m, \mathcal{A}) = 1$ for all j .
2. $r(\phi_k^m \phi_l^m, \mathcal{A}) = 1$ if the unordered pair $(k, l) \notin \{(1, 2), (3, 4), \dots, (2n - 5, 2n - 4)\}$.

5.1. Elementary abelian subgroups. — For $s > 0$ define

$$\mathcal{A}_E^s = \langle \{E_{j1}^s, {}_{1k}E^s : 2 \leq j \leq n, \quad 3 \leq k \leq n\} \rangle.$$

We say that a subgroup $\mathcal{A} < \text{Out}(F_n)$ has type E (for “elementary”) if there exists $s > 0$ and some basis for F_n in which \mathcal{A} equals \mathcal{A}_E^s . Equivalently, if one prefers to work with a fixed basis, then \mathcal{A} has type E if it equals $i_\psi \mathcal{A}_E^s$ for some s and some $\psi \in \text{Out}(F_n)$. We sometimes write \mathcal{A}_E for \mathcal{A}_E^1 . Note that the nontrivial elements of a type E subgroup \mathcal{A} have the same (unique) axis. We refer to this axis as the *characteristic axis* of \mathcal{A} .

In the notation of Lemma 4.4, ${}_{1j}E = \mu_{2j-5, e_1}$ and $E_{j1} = \mu_{2j-4, e_1}$ for $3 \leq j \leq n$. We extend this notation slightly and denote E_{21} by μ_{2n-3, e_1} .

5.1. Lemma. — *Let $\{\phi_1, \dots, \phi_{2n-3}\}$ be a basis for a torsion-free abelian subgroup \mathcal{A} . Then there exists $\psi \in \text{Out}(F_n)$ and $s, t > 0$ such that $i_\psi(\phi_i^s) = \mu_{i, e_1}^t$ for all i , if and only if each of the following conditions holds:*

1. $\{\phi_1, \dots, \phi_{2n-4}\}$ satisfies the pairing property for \mathcal{A} .
2. $r(\phi_j^m \phi_{2n-3}^m, \mathcal{A}) = 1$ for $1 \leq j \leq 2n-4$ and for all $m \neq 0$.

Proof. — The “only if” direction follows from Lemma 4.4 and Lemma 2.14. The “if” direction follows directly from Lemma 9.3 of [FH]. \square

The following corollary includes, as a special case, that the Φ -image of an elementary outer automorphism is elementary.

5.2. Corollary. — *If \mathcal{A} has type E then there is a normalization Φ' of Φ that almost fixes \mathcal{A} . Equivalently, there exists $\psi \in \text{Out}(F_n)$ and $s, t > 0$ so that $\Phi(\eta^s) = i_\psi(\eta^t)$ for each $\eta \in \mathcal{A}$.*

Proof. — There is no loss in assuming that $\mathcal{A} = \mathcal{A}_E^s \subset \Gamma$. The corollary then follows from Lemma 4.2 and from Lemma 5.1 applied to $\{\phi_i = \Phi(\mu_{i, e_1}^s)\}$. \square

5.2. Abelian subgroups of IA_n . — An element of \mathcal{A}_E is represented by an automorphism that multiplies each $x_j, j > 1$, on the left and on the right by various powers of x_1 . In this section we consider the analogous subgroup where we replace x_1 by a non-basis element $w \in F_2$, and we restrict the action to those x_j 's with $j > 2$. We impose a homology condition on w to control the image under Φ .

Let IA_n denote the subgroup of $\text{Out}(F_n)$ consisting of those elements which act trivially on $H_1(F_n, \mathbf{Z})$. For any nontrivial w in the commutator subgroup $[F_2, F_2]$, and for any fixed $s > 0$, define

$$\mathcal{A}_w^s = \langle \mu_{i, w}^s : 3 \leq i \leq 2n-4 \rangle$$

where $\mu_{i,w}$ is defined as in Section 4.3. Note that $A_w^s \subset \text{IA}_n$ and that

$$T_w^s = \mu_{3,w}^s \mu_{4,w}^s \cdots \mu_{2n-1,w}^s \in \mathcal{A}_w^s.$$

We say that a subgroup $\mathcal{A} < \text{Out}(\mathbb{F}_n)$ has *type C* if it equals $i_\eta(A_w^s)$ for some $\eta \in \text{Out}(\mathbb{F}_n)$, for some $w \in [\mathbb{F}_2, \mathbb{F}_2]$, and for some $s > 0$. We say that an element of $\text{Out}(\mathbb{F}_n)$ is a *C-twist* if it equals $i_\eta T_w^s$ for some $\eta \in \text{Out}(\mathbb{F}_n)$, some $w \in [\mathbb{F}_2, \mathbb{F}_2]$ and some $s > 0$. We sometimes write \mathcal{A}_w for \mathcal{A}_w^1 .

The nontrivial elements of a type C subgroup $\mathcal{A} < \text{Out}(\mathbb{F}_n)$ have a common (unique) axis, which we will refer to as the *characteristic axis* of \mathcal{A} . If $\mathcal{A} = i_\eta(A_w^s)$ then the characteristic axis is $\eta([w]_a)$. In order to recognize type C subgroups, we begin by recalling the following.

5.3. Lemma ([FH], Lemma 9.4). — *Suppose that $\{\phi_1, \dots, \phi_{2n-4}\}$ is a basis for a torsion-free abelian subgroup $\mathcal{A} \subset \text{IA}_n$ and that $\{\phi_1, \dots, \phi_{2n-4}\}$ satisfies the pairing property for \mathcal{A} . Then there exists $\psi \in \text{Out}(\mathbb{F}_n)$, a primitive element $w \in [\mathbb{F}_2, \mathbb{F}_2]$ and integers $s, t > 0$ such that $i_\psi(\phi_i^s) = \mu_{i,w}^t$ for each i .*

For each $1 \leq i \leq 2n - 4$, the map $a \mapsto \mu_{i,a}$ defines an injective homomorphism $\mathbb{F}_2 \rightarrow \text{Out}(\mathbb{F}_n)$. Given an arbitrary finite index subgroup $\Gamma \subseteq \text{Out}(\mathbb{F}_n)$, define

$$\Gamma_2 := \{a \in \mathbb{F}_2 : \mu_{i,a} \in \Gamma \text{ for all } 1 \leq i \leq 2n - 4\}$$

which is a finite index subgroup of \mathbb{F}_2 . The first half of the next lemma produces type C subgroups in $\Phi(\Gamma)$ and C-twists whose Φ -images are C-twists. The second half relates the Φ -images of E_{21} and T_w .

5.4. Lemma. — *For all nontrivial $w \in [\Gamma_2, \Gamma_2]$ there exist $s, t > 0$, a normalization $\Phi' = i_\psi \Phi$ and a primitive $v \in [\mathbb{F}_2, \mathbb{F}_2]$ such that:*

1. $\Phi'(\mu_{i,w}^s) = \mu_{i,v}^t$ for all $1 \leq i \leq 2n - 4$.
2. $\Phi'(T_w^s) = T_v^t$.
3. The characteristic axis of $\Phi'(\mathcal{A}_E^s)$ is carried by $[\mathbb{F}_2]$.
4. The characteristic axis of $\Phi(\mathcal{A}_E^s)$ is carried by $F([c]_u)$, where $[c]_u$ is the unique axis of $\Phi(T_w^s)$ and where $F([c]_u)$ is the unique conjugacy class of free factor of rank two that carries $[c]_u$.

Proof. — $\{\mu_{1,w}, \dots, \mu_{2n-4,w}\}$ satisfies the pairing property by Lemma 4.4 and is contained in $[\Gamma, \Gamma]$ by construction. The latter implies that each $\Phi(\mu_{i,w})$ is an element of $[\text{Out}(\mathbb{F}_n), \text{Out}(\mathbb{F}_n)]$ and hence an element of IA_n and the former, in conjunction with Lemma 4.2, implies that $\{\Phi(\mu_{1,w}), \dots, \Phi(\mu_{2n-4,w})\}$ satisfies the pairing property. (1) is therefore a consequence of Lemma 5.3. (2) follows from (1) and the fact that T_w^s is

represented by $\hat{\mu}_{1,w}^s \hat{\mu}_{2,w}^s \cdots \hat{\mu}_{2n-4,w}^s$. Assuming (3) for the moment, the characteristic axis of $\Phi(\mathbb{A}_E^s)$ is carried by

$$\psi^{-1}([\mathbb{F}_2]) = \psi^{-1}(\mathbb{F}([v]_u)) = \mathbb{F}(\psi^{-1}([v]_u)) = \mathbb{F}([c]_u)$$

where the last equality follows from (2). Thus (3) implies (4) and it remains only to verify (3).

For $4 \leq j \leq n$, define $\hat{\theta}_j$ by $x_j \mapsto \bar{v}x_jv$. Thus $\theta_j^t = \Phi'(\mu_{2j-5,w}^s \mu_{2j-4,w}^s)$ and θ_j^t commutes with $\eta := \Phi'(E_{31}^s)$, where we assume without loss that $E_{31}^s \in \Gamma$. Corollary 5.2 implies that η is elementary. Lemma 2.15 and Lemma 2.16 (see also Lemma 2.17) imply that

$$\text{Fix}(\theta_j^t) = \{[\langle x_i : i \neq j \rangle], [\langle v, x_j \rangle]\}.$$

It follows that $[\langle x_i : i \neq j \rangle]$, $[\langle v, x_j \rangle]$ and $[\langle \mathbb{F}_2, x_3 \rangle] = [\cap_{j=4}^n \langle x_i : i \neq j \rangle]$ are η -invariant, where the last fact follows from Corollary 2.4.

The set A_j of conjugacy classes of elements in $\langle v, x_j \rangle$ is η -invariant. If $\langle \mathbb{F}_2, x_j \rangle$ is not the minimal carrier $\mathbb{F}(A_j)$ of A_j then there is a free factor \mathbb{F}' of rank one and a free factor \mathbb{F}'' of rank two such that $\langle \mathbb{F}_2, x_j \rangle$ is conjugate to $\mathbb{F}' * \mathbb{F}''$ and such that each conjugacy class in $\langle v, x_j \rangle$ is carried by either \mathbb{F}' or \mathbb{F}'' . Since v is not a basis element, $[v]$ is carried by \mathbb{F}'' . It follows that \mathbb{F}'' is conjugate to \mathbb{F}_2 , and we may assume without loss that $\mathbb{F}'' = \mathbb{F}_2$. But then \mathbb{F}' would have to carry $[vx_j^k]$ for all k which is impossible. We may therefore assume that $\langle \mathbb{F}_2, x_j \rangle$ equals $\mathbb{F}(A_j)$ and so is η -invariant by Lemma 2.7.

We next assume that $\eta|_{\langle \mathbb{F}_2, x_3 \rangle}$ is trivial and argue to a contradiction. Choose $\hat{\eta} \in \text{Aut}(\mathbb{F}_n)$ such that $\hat{\eta}|_{\langle \mathbb{F}_2, x_3 \rangle} = \text{Id}$. Lemma 2.10 implies that $\hat{\eta}(x_j) = \alpha x_j^\pm \beta$ for some $\alpha, \beta \in \mathbb{F}_2$. Since η and θ_j^t commute and $\hat{\eta}$ and $\hat{\theta}_j^t$ both restrict to the identity on \mathbb{F}_2 , $\hat{\eta}$ commutes with $\hat{\theta}_j^t$. It follows that $\alpha = v^p$ and $\beta = v^q$ for some p and q . Since v is homologically trivial and η is elementary, $[v]_u$ is not the axis of η ; Lemma 2.12(3) implies that $p = q = 0$. As this holds for all $j \geq 4$, η^2 is the identity, which is a contradiction. We have now shown that $\eta|_{\langle \mathbb{F}_2, x_3 \rangle}$ is nontrivial and hence that $\eta|_{\langle \mathbb{F}_2, x_3 \rangle}$ contains the unique axis a of η .

The symmetric argument, with the roles of x_3 and x_4 reversed, implies that a is carried by $\langle \mathbb{F}_2, x_4 \rangle$. Corollary 2.3 implies that a is carried by $\mathbb{F}_2 = \langle \mathbb{F}_2, x_3 \rangle \cap \langle \mathbb{F}_2, x_4 \rangle$. Since a is the characteristic axis of $\Phi'(\mathcal{A}_E^s)$, this completes the proof of (3). \square

If a \mathbb{C} twist T_1 is defined with respect to $\{x_1, \dots, x_n\}$ then it is represented by the automorphism defined by $x_1 \mapsto w_1 x_1 \bar{w}_1$ and $x_2 \mapsto w_1 x_2 \bar{w}_1$ for some $w_1 \in \langle x_1, x_2 \rangle$. If a \mathbb{C} twist T_2 is defined with respect to the basis $\{x_3, x_4, x_1, x_2, x_5, \dots, x_n\}$, then it is represented the automorphism defined by $x_3 \mapsto w_2 x_3 \bar{w}_2$ and $x_4 \mapsto w_2 x_4 \bar{w}_2$ for some $w_2 \in \langle x_3, x_4 \rangle$. Thus T_1 and T_2 generate a rank two abelian subgroup. The following lemma, which uses the Kolchin theorem (Proposition 2.11), can thought of as a converse to this observation.

5.5. Lemma. — *Let T_1 and T_2 be C -twists, and suppose that $\mathcal{A} = \langle T_1, T_2 \rangle$ is a rank 2 abelian subgroup. If $[w_1]$ and $[w_2]$ are the characteristic axes of T_1 and T_2 , then there exist rank 2 free factors, F^1 carrying w_1 and F^2 carrying w_2 , such that $F^1 * F^2$ is a free factor of F_n .*

Proof. — Let F^1 be a rank two free factor that carries $[w_1]$. Then $[F^1]$ is invariant under both T_1 and T_2 . Obviously $T_1|[F^1]$ is trivial. If F^1 carries $[w_2]$, then $T_2|[F^1]$ is trivial because the unique axis $[w_2]$ of T_2 is not carried by any proper free factor of F^1 and so cannot be an axis of $T_2|[F^1]$. If F^1 does not carry $[w_2]$, then $T_2|[F^1]$ is trivial because $T_2|[F^1]$ has no axes. Thus F^1 is \mathcal{A} invariant and $\mathcal{A}|[F^1]$ is trivial.

Since T_1 and T_2 are UL so is \mathcal{A} . By Proposition 2.11, there is a filtered graph

$$\emptyset = G_0 \subset G_1 \subset \cdots \subset G_K = G$$

and a Kolchin representative \mathcal{A}_G such that $[G_2] = [F^1]$ and such that $f|G_2 = \text{Id}$ for each $f \in \mathcal{A}_G$. Moreover, the lifts $f_1 : G \rightarrow G$ and $f_2 : G \rightarrow G$ of T_1 and T_2 are UL.

Let Y be the component of $\text{Fix}(f_1)$ that contains G_2 . Lemmas 4.6, 2.15 and 2.16 imply that Y has rank two and that no non-fixed edge of G has initial endpoint in Y . There are at least two fixed directions at every vertex in G and the terminal endpoint of a non-fixed edge is never attached to a valence one vertex, so Y does not have valence one vertices and must equal G_2 . Since the only axis of T_1 is carried by G_2 and since this axis has multiplicity one, every non-fixed edge e_j for f_1 has the same terminal endpoint in G_2 , and both u_j and $m_j(f_1)$ are independent of j .

The same analysis shows that the smallest subgraph X that carries $[w_2]$ has rank two and is a component of $\text{Fix}(f_2)$. If $X \cap G_2 \neq \emptyset$ then $X = G_2$. In that case, F^1 carries $[w_2]$ and the above argument shows that f_1 and f_2 have the same non-fixed edges $\{e_j\}$ and that $m_j(f_2)$ is independent of j . This contradicts the assumption that \mathcal{A}_G is abelian with rank two. Thus X is disjoint from G_2 . Choose a basepoint in G_2 and let F^2 be the subgroup of $\pi_1(G)$ determined by X . \square

For $s > 0$ and $1 \leq i \leq n$, define

$$\hat{\mathcal{A}}_i^s := \langle \{ \hat{E}_j^s, \hat{E}_i^s : j \neq i \} \rangle.$$

We sometimes write $\hat{\mathcal{A}}_i$ for $\hat{\mathcal{A}}_i^1$. Thus each \mathcal{A}_i is a type E subgroup and $\mathcal{A}_1 = \mathcal{A}_E$.

A more general statement of the following corollary is possible but we limit ourselves to what is needed later in the proof.

5.6. Corollary. — *For $i = 1, 2, 3$, let a'_i be the characteristic axis of $\Phi(\mathcal{A}_i)$. Then*

1. a'_i is represented by y_i , where y_1, y_2, y_3 are part of a basis for F_n .
2. If a rank two free factor F carries a'_1 and a'_2 then there are representatives y_1 of a'_1 and y_2 of a'_2 such that $F = \langle y_1, y_2 \rangle$.

Proof. — Theorem 2.1 and (1) imply (2) so it suffices to prove (1).

Choose $w_1, w_2 \in [\Gamma_2, \Gamma_2]$ and let $\hat{\mu}$ be the order two automorphism that switches x_1 with x_3 and switches x_2 with x_4 . Then $\langle T_{w_1}, i_\mu T_{w_2} \rangle$ is a rank two abelian subgroup. Lemma 5.4 implies that $T_1 := \Phi(T_{w_1}^s)$ and $T_2 := \Phi(i_\mu T_{w_2}^s)$ are C-twists for some $s > 0$. Moreover, if $[c_i]_\mu$ is the characteristic axis of T_i then a'_1 and a'_2 are carried by $[F(c_1)]$ and a'_3 is carried by $[F(c_2)]$. By Lemma 5.5, we may choose $F(c_1)$ and $F(c_2)$ so that $F(c_1) * F(c_2)$ is a free factor of F_n . Choose y_1 and y_2 in $F(c_1)$ representing a'_1 and a'_2 and choose $y_3 \in F(c_2)$ representing a'_3 . Then y_1 is a basis element of $F(c_1)$ and y_3 is a basis element of $F(c_2)$ which implies that y_1 and y_3 are cobasis elements.

By symmetry (not of the construction in the preceding paragraph but of the roles of a'_2 and a'_3 in this corollary), there is a representative y'_2 of a'_2 (i.e. a conjugate of y_2) such that y_1 and y'_2 are cobasis elements. Theorem 2.1 implies that $F(c) = \langle y_1, y_2 \rangle$ and (1) follows. \square

6. Respecting a free factor while almost fixing an abelian subgroup

We continue with the notation of the previous section. In addition, for $s > 0$ define

$$\hat{H}^s = \langle \hat{E}_{j_3}^s, {}_3\hat{E}^s : j = 4, \dots, n \rangle.$$

We sometimes write \hat{H} for \hat{H}^1 .

We say that Φ respects the decomposition $F_n = F_2 * F_{n-2}$ if it preserves $O(F_2)$, $O(F_{n-2})$ and $O(F_2, F_{n-2})$. In Lemma 5.4(2) we showed that Φ can always be normalized so that a single C twist defined with respect to the decomposition $F_n = F_2 * F_{n-2}$ is mapped to a C twist defined with respect to the same decomposition. Our main goal in this section is to prove the following proposition, which in turn will be an important step in the proof of Theorem 1.1.

6.1. Proposition (Respecting a decomposition). — *There is normalization of Φ that respects the decomposition $F_n = F_2 * F_{n-2}$ and that almost fixes H .*

We work throughout with a fixed basis $\{x_1, \dots, x_n\}$.

6.1. Comparing normalizations. — The following lemma is used throughout the normalization process. It relates the weak centralizer of an element to the set of normalizations of Φ that fix that element.

6.2. Lemma. — *If both Φ and $i_\psi \circ \Phi$ almost fix η , then $\psi \in \text{WC}(\eta)$.*

Proof. — There exist $s, t, u, v > 0$ such that $\Phi(\eta^s) = \eta^t$ and $i_\psi \circ \Phi(\eta^u) = \eta^v$. Thus $i_\psi \eta^{tu} = i_\psi \circ \Phi(\eta^s)^u = i_\psi \circ \Phi(\eta^u)^s = \eta^{sv}$. Since i_ψ is an automorphism of $\text{Out}(\mathbf{F}_n)$ and since $tu, sv > 0$, we have that $tu = sv$. Thus ψ commutes with η^{tu} . \square

Motivated by Lemma 6.2, we calculate some weak centralizers.

6.3. Lemma. — *The following statements hold.*

1. If $\psi \in \text{WC}(\mathbf{H})$ then ψ is represented by $\hat{\psi}$, where $\hat{\psi}|_{\mathbf{F}_{n-2}} \in \hat{\mathbf{H}}|_{\mathbf{F}_{n-2}}$ and $\langle x_1, x_2, x_3 \rangle$ is $\hat{\psi}$ -invariant.
2. If $\psi \in \text{WC}(\mathbf{H})$ and $[\mathbf{F}_2]$ is ψ -invariant, then ψ is represented by $\hat{\psi}_1 \times \hat{\psi}_2 \in \text{Aut}(\mathbf{F}_2) \times \text{Aut}(\mathbf{F}_{n-2})$ where $\hat{\psi}_2 \in \hat{\mathbf{H}}|_{\mathbf{F}_{n-2}}$.
3. $\text{WC}(\mathcal{A}_i) = \mathcal{A}_i$.

Proof. — Assume that $\psi \in \text{WC}(\mathbf{H})$ and choose $s > 0$ so that ψ commutes with \mathbf{H}^s . Lemma 2.17 implies that $[x_3]_u, \langle x_3, x_j \rangle$ and $\langle \{x_k : k \neq j\} \rangle$ are ψ -invariant for all $j \geq 4$. Choose $\hat{\psi}$ so that

$$\hat{\psi}(x_3) = x_3^\epsilon$$

with $\epsilon = \pm 1$. Corollary 2.5 and Lemma 2.10 imply that for each $j \geq 4$, the groups $\langle x_3, x_j \rangle$ and $\langle \{x_k : k \neq j\} \rangle$ are $\hat{\psi}$ -invariant and that

$$\hat{\psi}(x_j) = x_3^p x_j^\delta x_3^q$$

for some $p, q \in \mathbf{Z}$ and $\delta = \pm 1$. The intersection $\langle x_1, x_2, x_3 \rangle = \bigcap_{j=4}^n \langle \{x_k : k \neq j\} \rangle$ is therefore $\hat{\psi}$ -invariant. For (1) it suffices to prove that $\epsilon = \delta = 1$.

For each $j \geq 4$, the automorphisms $\hat{\psi} \hat{\mathbf{E}}_{j3}^s$ and $\hat{\mathbf{E}}_{j3}^s \hat{\psi}$ represent the same outer automorphism and agree on $\langle x_1, x_2, x_3 \rangle$, and so must be equal. If $\delta = -1$ then

$$\hat{\psi} \hat{\mathbf{E}}_{j3}^s(x_j) = \hat{\psi}(x_j x_3^s) = x_3^p \bar{x}_j x_3^{q+\epsilon s}$$

and

$$\hat{\mathbf{E}}_{j3}^s \hat{\psi}(x_j) = \hat{\mathbf{E}}_{j3}^s(x_3^p \bar{x}_j x_3^q) = x_3^{p-s} \bar{x}_j x_3^q$$

which are unequal; thus $\delta = 1$. If $\epsilon = -1$ then

$$\hat{\psi} \hat{\mathbf{E}}_{j3}^s(x_j) = \hat{\psi}(x_j x_3^s) = x_3^p x_j x_3^{q-s}$$

and

$$\hat{\mathbf{E}}_{j3}^s \hat{\psi}(x_j) = \hat{\mathbf{E}}_{j3}^s(x_3^p x_j x_3^q) = x_3^p x_j x_3^{q+s}$$

which are unequal; thus $\epsilon = 1$. This proves (1).

Suppose now that $[F_2]$ is ψ -invariant. Then $\psi \in O(F_2, F_{n-2})$ and so is represented by $\hat{\psi}' = \hat{\psi}'_1 \times \hat{\psi}'_2 \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$. Since each element of \hat{H}^s restricts to the identity on F_2 , $\hat{\psi}'$ commutes with \hat{H}^s . If $\hat{\psi}$ is as in (1) then $\hat{\psi}'\hat{\psi}^{-1}$ is an inner automorphism i_c that commutes with \hat{H}^s and preserves F_{n-2} . It follows that $c \in F_{n-2} \cap \text{Fix}(\hat{H}^s) \subset F_{n-2} \cap \bigcap_{j=4}^n \langle x_k : k \neq j \rangle = \langle x_3 \rangle$ and hence that $\hat{\psi}'_2 = i_c \hat{\psi}|_{F_{n-2}} \in \hat{H}|_{F_{n-2}}$. This proves (2).

For (3) we may assume without loss that $i = 3$. The automorphism $\hat{\psi}$ commutes with ${}_3\hat{E}^s$ and \hat{E}_{j3}^s for $j = 1, 2$ because they commute on $\langle x_3, x_4 \rangle$ and their corresponding outer automorphisms commute. The same calculation as in the proof of (1) now applies to show that $\hat{\psi} \in \mathcal{A}_3$. \square

6.2. Preserving $O(F_2)$ and $O(F_{n-2})$. — We are now ready for the following.

Proof of Proposition 6.1. — We may assume by Corollary 5.2 that Φ almost fixes H . We divide the proof into steps to clarify the logic.

Step 1 (Defining W and Q). — Choose a finite generating set B for $\Gamma \cap O(F_2)$ and let Γ_2 be the finite index subgroup of F_2 defined in Section 5.2. Each $\mu \in B$ is represented by $\hat{\mu}_1 \times \text{Id}$ for some $\hat{\mu}_1 \in \text{Aut}(F_2)$. We will show that there is a finite subset W of $[\Gamma_2, \Gamma_2]$ with the following properties.

- (1) W is not contained in a cyclic subgroup of F_n .
- (2) For each $\mu \in B$ there exists $w \in W$ such that $\hat{\mu}_1(w) \in W$.

To construct W , note that for each $\mu \in B$, the group $\Gamma_2 \cap \hat{\mu}_1^{-1}(\Gamma_2)$ has finite index in F_2 and so contains noncommuting elements α and β . Setting $w = [\alpha, \beta]$ we have $w, \hat{\mu}_1(w) \in [\Gamma_2, \Gamma_2]$. If W contains one such pair for each μ then (2) is satisfied. If (1) is not satisfied then add any element of $[\Gamma_2, \Gamma_2]$ that is not contained in the maximal cyclic subgroup containing W . This is always possible since Γ_2 has finite index in F_2 .

Let $Q = \langle \Phi(T_w) : w \in W \rangle$ which as a set equals $\{\Phi(T_w) : w \in \langle W \rangle\}$. Since $\langle W \rangle \subset [\Gamma_2, \Gamma_2]$, Lemma 5.4 implies that each element of Q has an iterate that is a C-twist. Corollary 5.7.6 of [BFH1] implies that Q has a UL subgroup of finite index. After replacing each $w \in W$ with a suitable power we may assume that Q itself is UL and that

- (3) $\Phi(T_w)$ is a C twist for each $w \in W$.

Since H is almost fixed by Φ and commutes with each T_w , we have $Q \subset \text{WC}(H)$. Lemma 6.3(1) and the fact that no element of Q has $[x_3]_u$ as an axis, imply that $[F_{n-2}]$ is Q -invariant and that $Q|[F_{n-2}]$ is trivial.

Step 2 (A preliminary Kolchin representative \mathcal{Q}_G). — By Proposition 2.11, there exists a filtered graph $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_K = G$, a Kolchin representative \mathcal{Q}_G and a filtration element G_m such that $[\pi_1(G_m)] = [F_{n-2}]$ and such that $f|_{G_m}$ is the identity for all $f \in \mathcal{Q}_G$. After collapsing edges to points if necessary, we may assume that if $j > m$ and if the unique edge e_j of $G_j \setminus G_{j-1}$ is \mathcal{Q}_G -fixed and does not have both endpoints in G_{j-1} then it is a loop that is disjoint from G_{j-1} .

Choose $w \in W$ and let $T' = \Phi(T_w)$. We claim that the unique axis a' of T' is not carried by G_m . Since Φ almost fixes H , we know that $[\pi_1(G_m)] = [F_{n-2}]$ carries the characteristic axis of $\Phi(A_3)$. If $[\pi_1(G_m)]$ carries a' then, by Lemma 5.4, it also carries the characteristic axis of $\Phi(A_2)$ and $\Phi(A_1)$. Lemma 5.6 then implies that $[\pi_1(G_m)]$ has rank at least three. On the other hand, Lemma 4.6 implies that there is a unique T' -invariant free factor that carries a' and on which the restriction of T' represents the trivial outer automorphism; moreover, this free factor has rank two. This completes the proof of the claim.

An immediate consequence is that the unique edge e_{m+1} of $G_{m+1} \setminus G_m$ must be \mathcal{Q}_G -fixed. By Lemma 4.6, e_{m+1} does not have both endpoints in G_m and must therefore be a loop in the complement of G_m . Since a' is not represented by a basis element this same argument can be repeated to conclude that e_{m+2} is a \mathcal{Q}_G -fixed loop that is disjoint from G_m . Rank considerations prevent this argument from being repeated yet again so the basepoints of e_{m+1} and e_{m+2} must be equal. Let X be the subgraph with edges e_{m+1} and e_{m+2} . Then $G = G_m \cup X \cup e_{m+3}$ where G_m and X are disjoint and \mathcal{Q}_G -fixed, where X carries a' and where e_{m+3} is an edge with initial endpoint in G_m and terminal endpoint in X . The subgraph $e_{m+3} \cup X$ determines a free factor F'_2 that carries a' and satisfies $F_n = F'_2 * F_{n-2}$. Note that all of this is independent of the choice of $w \in W$ used to define T' .

Step 3 (Improving \mathcal{Q}_G and choosing the normalization). — Choose $t > 0$ so that $H^t \subset \Phi(H \cap \Gamma)$. Then H^t commutes with T' and $[F'_2]$ is H^t -invariant. Since $[F'_2]$ does not carry $[x_3]$, we have that $H^t|[F'_2]$ is trivial. Lemma 2.17(1) implies that $[F'_2]$ is carried by $[\langle x_i : i \neq k \rangle]$ for each $4 \leq k \leq n$ and so by Corollary 2.4 is carried by $[\langle x_1, x_2, x_3 \rangle]$. Equivalently, $F''_2 := i_\gamma F'_2 \subset \langle x_1, x_2, x_3 \rangle$ for some $\gamma \in F_n$. We claim that γ can be chosen in F_{n-2} .

Theorem 2.1 implies that $\langle x_1, x_2, x_3 \rangle = F''_2 * \langle x_3 \rangle$ and hence that $F_n = F''_2 * F_{n-2}$. Thus

$$(i_\gamma, \text{Id}) : F'_2 * F_{n-2} \rightarrow F''_2 * F_{n-2}$$

is an isomorphism, which we can realize by a homotopy equivalence $h : G \rightarrow G$ by letting u be the closed path based at the initial basepoint of e_{m+3} that determines γ , and by defining h by $h|(G_m \cup X) = \text{Id}$ and by letting $h(e_{m+3})$ be the path obtained from ue_{m+3} by tightening. Lemma 3.2.2 of [BFH1] implies that $h(e_{m+3}) = u_1 e_{m+3} u_2$ where u_1 is a (possibly trivial) closed loop in G_m and u_2 is a (possibly trivial) closed loop

in X . Thus u is obtained by tightening $u_1 e_{m+3} u_2 \bar{e}_{m+3}$. Let $\gamma_1 \in F_{n-2}$ be the element determined by u_1 and let $\gamma_2 \in F'_2$ be the element determined by $e_{m+3} u_2 \bar{e}_{m+3}$. Then $\gamma = \gamma_1 \gamma_2$ and $i_\gamma(F'_2) = i_{\gamma_1}(F'_2)$. Replacing γ with γ_1 completes the proof of the claim that γ can be chosen in F_{n-2} .

We now assume that $\gamma \in F_{n-2}$ and that $u \subset G_m$. Thus h commutes with each $f \in Q_G$ and we may change the marking on G by postcomposing the given marking with h and still have that Q_G is a Kolchin representative of Q . This results in F'_2 , which is defined to be the free factor determined by subgraph $e_{m+3} \cup X$, being replaced by F''_2 . In particular, we may assume that $F'_2 \subset \langle x_1, x_2, x_3 \rangle$ and hence that $\hat{H}|_{F'_2}$ is the identity. Choose $\hat{\psi} \in \text{Aut}(F_n)$ such that $\hat{\psi}(F'_2) = F_2$ and $\hat{\psi}|_{F_{n-2}} = \text{Id}$. Then ψ commutes with H because $\hat{\psi}$ commutes with \hat{H} . Replace Φ with $i_\psi \circ \Phi$ and note that Φ still almost fixes H . The effect on Q and Q_G is that Q is replaced by $i_\psi(Q)$ and the marking on G is changed by precomposing with $\hat{\psi}^{-1}$. Thus F'_2 is replaced with F_2 and $T' = T_v$ for some $v \in [F_2, F_2]$.

Step 4 (Checking the properties). — By choosing $w_1, w_2 \in W$ that do not commute, we have $\Phi(T_{w_i}) = T_{v_i}$ for noncommuting $v_1, v_2 \in [F_2, F_2] \subset F_2$. Thus v_1 and v_2 are not multiples of a common indivisible element and, with one possible exception, the only conjugacy classes carried by both $\langle F_{n-2}, v_1 \rangle$ and $\langle F_{n-2}, v_2 \rangle$ are those carried by F_{n-2} . The one exception is the conjugacy class of v_1 and v_2 if v_1 and v_2 happen to be conjugate. Note that this exceptional case is not the conjugacy class of a basis element. For every $\eta \in \Gamma \cap O(F_{n-2})$, the element $\eta' := \Phi(\eta)$ commutes with both T_{v_1} and T_{v_2} . Lemma 4.6 implies that F_2 and $\langle F_{n-2}, v_i \rangle$ are η' -invariant. In particular, if y is a basis element of F_{n-2} then $\eta'([y])$ is carried by both $\langle F_{n-2}, v_1 \rangle$ and $\langle F_{n-2}, v_2 \rangle$ and so also by F_{n-2} . Corollary 2.9 implies that F_{n-2} is η' -invariant and hence that $\eta' \in O(F_2, F_{n-2})$.

Lemma 4.3 implies that η' is represented by $\hat{\eta}'_1 \times \hat{\eta}'_2 \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$. Choose $s > 0$ so that $E_{21}^s, E_{12}^s \in \Gamma$. Then η' commutes with both $\Phi(E_{21}^s)$ and $\Phi(E_{12}^s)$ and so preserves their unique axes a'_1 and a'_2 . Lemma 5.4 and Lemma 5.6 imply that $a'_1 = [y'_1]_u$ and $a'_2 = [y'_2]_u$ where $\{y'_1, y'_2\}$ is a basis for F_2 . As an element of the mapping class group of the once punctured torus, η'_1 preserves the unoriented isotopy class of a pair of non-isotopic simple closed curves and so has finite order. We also know that $\hat{\eta}'_1$ commutes with both i_{v_1} and i_{v_2} because η' commutes with T_{v_1} and T_{v_2} . Thus $\text{Fix}(\hat{\eta}'_1)$ has rank at least two. Lemma 2.20 implies that $\hat{\eta}'_1$ is the identity. This completes the proof that $\Phi(O(F_{n-2}) \cap \Gamma) \subset O(F_{n-2})$.

Suppose now that $\mu \in B$ and that $w, \hat{\mu}_1(w) \in W$. Then $\Phi(T_w) = T_v$ and $\Phi(T_{\hat{\mu}_1(w)}) = T_{v'}$ for $v, v' \in F_2$. Denote $\Phi(\mu)$ by μ' . Then

$$T_{v'} = \Phi(T_{\hat{\mu}_1(w)}) = \Phi(i_\mu T_w) = i_{\mu'} T_v$$

from which it follows that $F_2 = F([v']) = F([v])$ is μ' -invariant.

By Lemma 6.3(2), μ' is represented by $\hat{\mu}'_1 \times \hat{\mu}'_2 \in \text{Aut}(F_2) \times \text{Aut}(F_{n-2})$ where $\hat{\mu}'_2 \in \hat{H}|_{F_{n-2}}$. Choose $\theta \in O(F_{n-2}) \cap \Gamma$ that does not commute with any nontrivial

element of \mathbf{H} and let $\theta' = \Phi(\theta) \in \mathbf{O}(\mathbb{F}_{n-2})$. Then θ' commutes with μ' but does not commute with any nontrivial element of \mathbf{H} (because Φ almost fixes \mathbf{H}). The former implies that $\hat{\theta}'_2$ commutes with $\hat{\mu}'_2$ and hence commutes with $\text{Id} \times \hat{\mu}'_2 \in \hat{\mathbf{H}}$. The latter then implies that $\hat{\mu}'_2$ is the identity. This proves that $\Phi(\mu) \in \mathbf{O}(\mathbb{F}_2)$ and since this holds for each $\mu \in \mathbf{B}$, $\Phi(\mathbf{O}(\mathbb{F}_2) \cap \Gamma) \subset \mathbf{O}(\mathbb{F}_2)$. \square

Notation. — Let $\mathbf{D}_{ij} = {}_j\mathbf{E} \circ \mathbf{E}_{ij}$.

6.4. Lemma. — *The following properties hold for all $1 \leq i \neq j \leq n$.*

1. *The restriction of \mathbf{D}_{ij} to any invariant free factor of rank two is trivial.*
2. *$[\langle x_i, x_j \rangle]$ is the unique rank two element of $\text{Fix}(\mathbf{D}_{ij})$.*
3. *\mathbf{D}_{ij} is almost fixed by some normalization of Φ .*

Proof. — (1) and (2) follow from Lemma 2.17 and (3) follows from Corollary 5.2. \square

We say that an outer automorphism η has *type D* if it is equal to \mathbf{D}_{ij}^s for some choice of basis $\{x_1, \dots, x_n\}$, some $s \neq 0$ and some $1 \leq i \neq j \leq n$. We write $\text{FS}(\eta)$ for the unique rank two element of $\text{Fix}(\eta)$. If $\{x_1, \dots, x_n\}$ is a basis with respect to which $\eta = \mathbf{D}_{ij}$, then $\text{FS}(\eta) = [\langle x_i, x_j \rangle]$. Note that $i_\psi \eta$ has type D for any $\psi \in \text{Out}(\mathbb{F}_n)$ and that $\text{FS}(i_\psi \eta) = \psi(\text{FS}(\eta))$. An immediate consequence of Lemma 6.4(3) is that if η has type D then there exists $s > 0$ so that $\Phi(\eta^s)$ has type D.

We make frequent use of the following easy consequence of Proposition 6.1.

6.5. Corollary. — $\text{FS}(\Phi(\mathbf{D}_{ij})) = \text{FS}(\Phi(\mathbf{D}_{ji}))$ for all $1 \leq i \neq j \leq n$.

Proof. — If $\Phi' = i_\psi \Phi$ then $\text{FS}(\Phi'(\mathbf{D}_{ij})) = \psi(\text{FS}(\Phi(\mathbf{D}_{ij})))$. We may therefore replace Φ with a normalization that respects the decomposition $\mathbb{F}_n = \langle x_i, x_j \rangle * \langle x_k : k \neq i, j \rangle$. In this case, both $\text{FS}(\Phi(\mathbf{D}_{ij}))$ and $\text{FS}(\Phi(\mathbf{D}_{ji}))$ equal $\langle x_i, x_j \rangle$. \square

7. Almost fixing certain subgroups attached to a free factor

In this section we build on what we showed in Section 6 by further normalizing Φ . More precisely, we prove in Section 7.1 that Φ can further be normalized by conjugating with an element of $\mathbf{O}(\mathbb{F}_2)$ so that the resulting map almost fixes each of \mathcal{A}_3 , $\langle {}_{12}\mathbf{E}, \mathbf{E}_{21} \rangle$ and $\langle {}_{21}\mathbf{E}, \mathbf{E}_{12} \rangle$. We then prove in Section 7.2 that Φ can be normalized even further so that the resulting map almost fixes each of $\langle \mathbf{E}_{ij}, {}_j\mathbf{E} \rangle$ for all $1 \leq i \neq j \leq n$.

7.1. Normalizing with respect to $\mathbf{O}(\mathbb{F}_2)$. — The next step in the ultimate normalization of Φ is to modify its induced action on $\mathbf{O}(\mathbb{F}_2)$. If $\psi \in \mathbf{O}(\mathbb{F}_2)$ then we say

that $i_\psi \circ \Phi$ is an $\mathbf{O}(\mathbf{F}_2)$ -normalization of Φ . If there exists $s > 0$ and $t \neq 0$ such that $\Phi(\eta^s) = \eta^t$ then we say that Φ *almost fixes* η *up to sign*.

The following proposition, whose proof appears at the end of the section, is the main result of this section.

7.1. Proposition. — *Assume that Φ respects the decomposition $\mathbf{F}_n = \mathbf{F}_2 * \mathbf{F}_{n-2}$ and almost fixes \mathbf{H} . Then there is an $\mathbf{O}(\mathbf{F}_2)$ -normalization of Φ that almost fixes $\langle {}_{21}\mathbf{E}, \mathbf{E}_{12} \rangle$, $\langle {}_{12}\mathbf{E}, \mathbf{E}_{21} \rangle$ and \mathbf{A}_3 and that almost fixes \mathbf{T}_ρ up to sign.*

The following lemma lists properties of $\hat{\phi}_1$ for $\phi \in \mathbf{O}(\mathbf{F}_2)$ of type D or E.

7.2. Lemma

1. If $\phi \in \mathbf{O}(\mathbf{F}_2)$ is elementary then $\hat{\phi}_1$ is defined by $z_2 \mapsto z_2 z_1^p$ for some $p > 0$ and some basis $\{z_1, z_2\}$ of \mathbf{F}_2 .
2. Each UL $\phi_1 \in \text{Out}(\mathbf{F}_2)$ has a unique representative $\hat{\phi}_1$ that fixes ρ and such that $\hat{\phi}_1 \times \text{Id}$ represents an elementary $\phi \in \mathbf{O}(\mathbf{F}_2)$.
3. $\phi \in \mathbf{O}(\mathbf{F}_2)$ has type D if and only if $\hat{\phi}_1 = i_a^s$ for some basis element $a \in \mathbf{F}_2$ and some $s > 0$.

Proof. — If $\phi \in \mathbf{O}(\mathbf{F}_2)$ is elementary then $\text{Fix}(\hat{\phi}_1 \times \text{Id})$ has rank n and so $\text{Fix}(\hat{\phi}_1)$ has rank two. (1) follows from Lemma 2.20.

If $\phi \in \mathbf{O}(\mathbf{F}_2)$ is elementary then by Lemma 2.18 there exists a representative $\hat{\phi}_1$ that is defined by $z_2 \mapsto z_2 z_1^b$ for some $b > 0$ and some basis $\{z_1, z_2\}$ of \mathbf{F}_2 . In particular, $\hat{\phi}_1$ fixes $\rho' := [z_1, z_2]$. There exists $c \in \mathbf{F}_2$ such that $\rho = i_c \rho'$. After replacing z_1 by $i_c(z_1)$, z_2 by $i_c(z_2)$ and $\hat{\phi}_1$ by $i_c \hat{\phi}_1 i_c^{-1}$, we may assume that $\hat{\phi}_1$ fixes ρ . It is clear that $\hat{\phi}_1 \times \text{Id}$ is elementary. To prove uniqueness, suppose that $\hat{\phi}'_1 \neq \hat{\phi}_1$ also fixes ρ and represents ϕ_1 . Then $\phi'_1 = i_\rho^k \hat{\phi}_1$ for some $k \neq 0$ and ϕ' is represented by $z_2 \mapsto z_2 z_1^b$ and $x_j \mapsto \bar{\rho}^k x_j \rho^k$ for $j \geq 3$. Thus $[\rho]_u$ is an axis for ϕ' and ϕ' is not elementary. This completes the proof of (2).

The if part of (3) is obvious. If $\phi \in \mathbf{O}(\mathbf{F}_2)$ has type D then $\text{Fix}(\hat{\phi}_1 \times \text{Id})$ is a rank $n - 1$ free factor. It follows that $\text{Fix}(\hat{\phi}_1) = \langle a \rangle$ for some basis element a and hence that $\hat{\phi}_1$ is defined by $b \mapsto a^t b^\delta a^s$ where $\delta = \pm 1$, $s, t \in \mathbf{Z}$ and $\mathbf{F}_2 = \langle a, b \rangle$. Since the unique axis of ϕ has multiplicity one and ϕ is not elementary, $s = -t$. Since ϕ is UL, $\delta = 1$. Thus $\hat{\phi}_1 = i_a^s$ and after replacing a by \bar{a} if necessary, $s > 0$. \square

Recall the notation $\mathbf{D}_{ij} = {}_{j\bar{i}}\mathbf{E} \circ \mathbf{E}_{ij}$.

7.3. Lemma. — *Assume that Φ respects the decomposition $\mathbf{F}_n = \mathbf{F}_2 * \mathbf{F}_{n-2}$ and almost fixes \mathbf{H} .*

1. There exists an $\mathbf{O}(\mathbf{F}_2)$ -normalization Φ' of Φ that almost fixes \mathbf{D}_{21} and \mathbf{D}_{12} up to sign.
2. If Φ' is as in (1) and if Φ' almost fixes \mathbf{E}_{21} and \mathbf{E}_{12} up to sign then Φ' almost fixes \mathbf{A}_3 .

Proof. — Choose $s > 0$ so that $\Phi(D_{ij}^s)$ has type D for all $1 \leq i \neq j \leq 3$. By Lemma 7.2(3) there is a basis element $a \in F_2$ and $r \neq 0$ such that $i_a^r \times \text{Id}$ represents $\Phi(D_{21}^s)$. Choose $\omega \in O(F_2)$ such that $\hat{\omega}_1(a) = x_1$ and replace Φ by $i_\omega \Phi$. Then $\Phi(D_{21}^s) = D_{21}^r$ or in other words, Φ almost fixes D_{21} up to sign.

Corollary 5.2 implies, after increasing s if necessary, that there exists $\psi \in \text{Out}(F_n)$ and $t > 0$ so that $\Phi(\theta^s) = i_\psi \theta^t$ for all $\theta \in \mathcal{A}_3$. Since Φ almost fixes H , Lemma 6.2 implies that $\psi \in \text{WC}(H)$. By Lemma 6.3 there is a representative $\hat{\psi}$ that leaves $\langle x_1, x_2, x_3 \rangle$ and F_{n-2} invariant, and whose restriction to F_{n-2} agrees with the restriction of an element of \hat{H} . In particular, $\hat{\psi}(x_3) = x_3$ and neither $[\hat{\psi}(\langle x_1, x_3 \rangle)]$ nor $[\hat{\psi}(\langle x_2, x_3 \rangle)]$ is equal to $[\langle x_1, x_2 \rangle]$.

The next section of the proof of (1) uses only the fact that Φ almost fixes D_{21} up to sign and will be referred to as the ‘ $\langle x_1, x_3 \rangle$ step’ when we prove (2).

Let $\sigma = D_{13}$, $\tau = D_{31}$ and $\mu = D_{21}$. Then $\Phi(\mu^s) = \mu^r$ and

$$\text{FS}(\sigma^s) = \text{FS}(\tau^s) = [\langle x_1, x_3 \rangle]$$

and

$$\text{FS}(\mu^s) = \text{FS}(\mu^r) = [\langle x_1, x_2 \rangle].$$

By Corollary 6.5 we have

$$\text{FS}(\Phi(\tau^s)) = \text{FS}(\Phi(\sigma^s)) = \text{FS}(i_\psi(\sigma^t)) = \psi(\text{FS}(\sigma^t)) = \psi([\langle x_1, x_3 \rangle]).$$

Since μ^r commutes with $\Phi(\tau^s)$, we have that $[\hat{\psi}(\langle x_1, x_3 \rangle)] = \text{FS}(\Phi(\tau^s))$ is μ^r -invariant. The restriction of μ^r to $[\hat{\psi}(\langle x_1, x_3 \rangle)]$ is trivial by Lemma 6.4(1). The only element of $\text{Fix}(\mu^r)$ other than $[\langle x_1, x_2 \rangle]$ is $[\langle x_i : i \neq 2 \rangle]$, and so $\hat{\psi}(\langle x_1, x_3 \rangle)$ is carried by $[\langle x_i : i \neq 2 \rangle]$. Thus

$$\hat{\psi}(\langle x_1, x_3 \rangle) \subset \langle x_1, x_2, x_3 \rangle \cap i_c \langle x_i : i \neq 2 \rangle$$

for some $c \in F_n$. We may assume, by Theorem 2.1 applied to $H = \langle x_1, x_2, x_3 \rangle$, that $c \in \langle x_1, x_2, x_3 \rangle$. Thus

$$\langle x_1, x_2, x_3 \rangle \cap i_c \langle x_i : i \neq 2 \rangle = i_c(\langle x_1, x_2, x_3 \rangle \cap \langle x_i : i \neq 2 \rangle) = i_c \langle x_1, x_3 \rangle$$

from which it follows that $\hat{\psi}(\langle x_1, x_3 \rangle) = i_c \langle x_1, x_3 \rangle$. Since $x_3 \in \hat{\psi}(\langle x_1, x_3 \rangle)$, Lemma 2.3 implies that $\hat{\psi}(\langle x_1, x_3 \rangle) = \langle x_1, x_3 \rangle$. This completes the ‘‘ $\langle x_1, x_3 \rangle$ step’’.

We now turn our attention to

$$[\hat{\psi}(\langle x_2, x_3 \rangle)] = [\hat{\psi}(\text{FS}(D_{23}))] = \text{FS}(\Phi(D_{23}^s)) = \text{FS}(\Phi(D_{32}^s)).$$

Lemma 6.4(1) and the fact that $\Phi(D_{32}^s)$ commutes with $\Phi(D_{12}^s)$ imply that $[\hat{\psi}(\langle x_2, x_3 \rangle)]$ is $\Phi(D_{12}^s)$ -invariant and that the restriction of $\Phi(D_{12}^s)$ to $[\hat{\psi}(\langle x_2, x_3 \rangle)]$ is trivial. By Lemma 7.2(3), there is a basis element $b \in F_2$ such that $\Phi(D_{12}^s)$ is represented by a positive iterate of $i_b \times \text{Id}$. The only element of $\text{Fix}(\Phi(D_{12}^s))$ other than $[\langle x_1, x_2 \rangle]$ is $[\langle \{b, x_i : i \geq 3\} \rangle]$, and so $\hat{\psi}(\langle x_2, x_3 \rangle)$ is carried by $[\langle \{b, x_i : i \geq 3\} \rangle]$. Thus

$$\hat{\psi}(\langle x_2, x_3 \rangle) \subset \langle x_1, x_2, x_3 \rangle \cap i_{c'} \langle \{b, x_i : i \geq 3\} \rangle$$

for some $c' \in F_n$. We may assume, by Theorem 2.1 applied to $H = \langle x_1, x_2, x_3 \rangle$, that $c' \in \langle x_1, x_2, x_3 \rangle$. Thus

$$\begin{aligned} \langle x_1, x_2, x_3 \rangle \cap i_{c'} \langle \{b, x_i : i \geq 3\} \rangle &= i_{c'}(\langle x_1, x_2, x_3 \rangle \cap \langle \{b, x_i : i \geq 3\} \rangle) \\ &= i_{c'} \langle b, x_3 \rangle \end{aligned}$$

from which it follows that $\hat{\psi}(\langle x_1, x_3 \rangle) = i_{c'} \langle b, x_3 \rangle$. Since $x_3 \in \hat{\psi}(\langle b, x_3 \rangle)$, Lemma 2.3 applied to $\langle b, x_3 \rangle$ implies that $\hat{\psi}(\langle x_2, x_3 \rangle) = \langle b, x_3 \rangle$. This implies that $\hat{\psi}(x_2)$ and x_3 are a basis for $\langle b, x_3 \rangle$ and so by Lemma 2.10 we have that $\hat{\psi}(x_2) = x_3^k b^\delta x_3^l$ for some $k, l \in \mathbf{Z}$ and $\delta = \pm 1$.

On the other hand, Lemma 2.10 applied to $\langle x_1, x_2, x_3 \rangle$ also implies that $\hat{\psi}(x_2) = ux_2^\epsilon v$ for some $u, v \in \langle x_1, x_3 \rangle$ and $\epsilon = \pm 1$. Since $b \in F_2$, it follows that $b = x_1^i x_2^{\delta \epsilon} x_3^j$ for some $i, j \in \mathbf{Z}$. Define $\hat{\eta}^{-1} \in O(F_2)$ by $x_2 \mapsto x_1^i x_2^{\delta \epsilon} x_3^j$ and let $\psi' = \eta\psi$. Then $\Phi' := i_\eta \Phi$ satisfies $\Phi'(\theta^s) = i_\psi \theta^t$ for all $\theta \in \mathcal{A}_3$. Since η commutes with H and with D_{21} , Φ' almost fixes H and almost fixes D_{21} up to sign. Moreover, Φ almost fixes D_{12} because $\hat{\eta}(b) = x_2$. This completes the proof of (1).

To prove (2), assume that Φ' is an $O(F_2)$ -normalization of Φ that almost fixes D_{21} and D_{12} up to sign. As above, $\Phi'(\theta^s) = i_\psi \theta^t$ for all $\theta \in \mathcal{A}_3$ and some $s, t > 0$ where ψ is represented by $\hat{\psi}$ such that $\hat{\psi}|_{F_{n-2}} \in \hat{H}$. The $\langle x_1, x_3 \rangle$ step used in the proof of (1) applies to both $\langle x_1, x_3 \rangle$ and $\langle x_2, x_3 \rangle$ and proves that both $\langle x_1, x_3 \rangle$ and $\langle x_2, x_3 \rangle$ are $\hat{\psi}$ -invariant.

There exist $d, e \in \mathbf{Z}$ and $\gamma = \pm 1$ such that $\hat{\psi}(x_2) = x_3^d x_2^\gamma x_3^e$. We claim that if Φ' almost fixes E_{21} up to sign then $\gamma = 1$. Indeed, if $\gamma = -1$ then direct computation shows that $\Phi'({}_{32}E^s) = i_\psi \circ {}_{32}E^t = E_{23}^t$. This contradicts the fact that ${}_{32}E^s$ commutes with E_{21} but E_{23}^t does not commute with E_{21}^m for any $m \neq 0$. The symmetric argument shows that if Φ' almost fixes E_{12} up to sign then $\hat{\psi}(x_1) = x_3^u x_1 x_3^v$. This completes the proof of (2). \square

The following corollary is a strengthening of Corollary 5.6.

7.4. Corollary. — *Assume that Φ respects the decomposition $F_n = F_2 * F_{n-2}$ and almost fixes H . Then there is a basis $\{a, b\}$ for F_2 , $s > 0$ and $t, u \neq 0$ such that $i_a^t \times \text{Id}$ represents $\Phi(D_{21}^s)$ and $i_b^u \times \text{Id}$ represents $\Phi(D_{12}^s)$.*

Proof. — By Lemma 7.3, there exists $\psi \in \text{O}(F_2)$ such that $\Phi' = i_\psi \circ \Phi$ almost fixes D_{21} and D_{12} up to sign. The conclusions of the corollary are satisfied with $a = \hat{\psi}_1^{-1}(x_1)$ and $b = \hat{\psi}_1^{-1}(x_2)$. \square

The next lemma produces a $\text{O}(F_2)$ -normalization of Φ with different useful properties than the one produced in Lemma 7.3. We will combine these in the ultimate proof of Proposition 7.1. Recall that $\rho = [x_1, x_2]$ and that $\hat{T}_\rho = i_\rho \times \text{Id}$.

7.5. Lemma. — *Suppose that Φ respects the decomposition $F_n = F_2 * F_{n-2}$ and almost fixes H . Then there is a $\text{O}(F_2)$ -normalization of Φ that almost fixes E_{21} and that almost fixes E_{12} and T_ρ up to sign.*

Proof. — Choose $s > 0$ so that $\Phi(T_\rho^s)$ is defined and so that $\Phi({}_{12}E^s)$, $\Phi(E_{21}^s)$, $\Phi({}_{21}E^s)$ and $\Phi(E_{12}^s)$ are defined and elementary. By Lemma 7.2(1), $\Phi(E_{21}^s)$ differs from E_{21}^t for some $t > 0$ only by a change of basis in F_2 . After replacing Φ with a $\text{O}(F_2)$ -normalization of Φ , we may assume that $\Phi(E_{21}^s) = E_{21}^t$ and in particular that Φ almost fixes E_{21} .

Denote E_{21}^t by η and $\Phi(E_{12}^s)$ by μ . Corollary 5.6(2) implies that the axis of η and the axis of μ are represented by elements that form a basis for F_2 . Since the former is $[x_1]_u$, the latter must be $[x_2 x_1^k]_u$ for some $k \neq 0$. Let $T' = \Phi(T_\rho^s) \in \text{O}(F_2)$. Then T' commutes with μ and η , which implies that $T'_1 \in \text{Out}(F_2)$ preserves their axes and so has finite order. After replacing s by an iterate if necessary, we may assume that T'_1 is trivial. Thus $\hat{T}'_1 = i_x$ where $x \in \text{Fix}(\hat{\eta}_1) \cap \text{Fix}(\hat{\mu}_1)$. The second and third items of Corollary 2.19 imply that $x = i_a(\rho^l)$ for some $a \in \text{Fix}(\hat{\eta}_1)$ and some $l \neq 0$. Denote $i_a \times \text{Id}$ by $\hat{\sigma}$. Then σ commutes with η and $(i_\sigma \circ \Phi)(T_\rho^s) = T_\rho^l$. Replacing Φ with $i_\sigma \circ \Phi$, we may assume that Φ almost fixes E_{21} and almost fixes T_ρ up to sign.

Let $\hat{\psi} = \hat{E}_{21}^k$ and let $\hat{v} = i_{\hat{\psi}} \hat{E}_{12}$. Then \hat{v}_1 fixes ρ and $x_2 x_1^k$. Since μ_1 and v_1 are UL and fix the conjugacy class of the same basis element, they are iterates of a common element of $\text{Out}(F_2)$. We also know that μ commutes with an iterate of T_ρ and hence that $\hat{\mu}_1$ fixes ρ . Lemma 7.2(2) implies that $\hat{\mu}_1$ and \hat{v}_1 are iterates of some common element. Thus $\mu^r = v^q$ for some $r > 0$ and $q \neq 0$. In other words $(i_{\hat{\psi}}^{-1} \circ \Phi)(E_{12}^s)^{rs} = E_{12}^q$. Since T_ρ and E_{21} commute with ψ , we can replace Φ with $i_{\hat{\psi}}^{-1} \circ \Phi$. Thus Φ almost fixes E_{21} , almost fixes T_ρ up to sign, and almost fixes E_{12} up to sign. \square

The following two lemmas are used to show that certain elements that are almost fixed up to sign are in fact almost fixed.

7.6. Lemma. — *If Φ almost fixes E_{21} and almost fixes D_{21} up to sign then Φ almost fixes D_{21} and almost fixes $\langle {}_{12}E, E_{21} \rangle$.*

Proof. — There exist $s, t > 0$ and $r \neq 0$ so that $\Phi(E_{21}^s) = E_{21}^t$, $\Phi(D_{21}^s) = D_{21}^r$ and so that $\sigma := \Phi({}_{12}E^s)$ is elementary. Since σ commutes with D_{21}^r , it follows that

$\hat{\sigma}_1$ commutes with i'_{x_1} and hence that $\hat{\sigma}_1$ fixes x_1 . Thus $\hat{\sigma}_1$ is defined by $x_2 \mapsto \bar{x}_1^i x_2 x_1^j$ where either i or j is zero. Since $\hat{\sigma}_1$ and \hat{E}_{21} generate a rank two abelian subgroup, $j = 0$. Thus $\Phi({}_{12}E^s) = {}_{12}E^i$. It follows from $D_{21} = {}_{12}EE_{21}$, that $i = r = t$. \square

7.7. Lemma. — *If Φ almost fixes E_{12} and $i_\alpha \circ \Phi$ almost fixes E_{12} up to sign where $\alpha \in A_3$, then $i_\alpha \circ \Phi$ almost fixes E_{12} .*

Proof. — It suffices to show that if $\alpha E_{12}^p \alpha^{-1} = E_{12}^q$ then $p = q$. Let $\hat{\alpha}$ be the lift of α into \hat{A}_3 . Then $\hat{\alpha} \hat{E}_{12}^p \hat{\alpha}^{-1}$ and \hat{E}_{12}^q agree on F_{n-2} and represent the same outer automorphism so are equal. If $\hat{\alpha}(x_1) = \bar{x}_3^a x_1 x_3^b$ and $\hat{\alpha}(x_2) = \bar{x}_3^c x_2 x_3^d$ then $x_1 x_2^q = \hat{E}_{12}^q(x_1) = \hat{\alpha} \hat{E}_{12}^p \hat{\alpha}^{-1}(x_1) = x_1 x_3^b (\bar{x}_3^c x_2 x_3^d)^p \bar{x}_3^b$. This proves that $p = q$ as desired. \square

Proof of Proposition 7.1. — We may assume by Lemma 7.5 that Φ almost fixes E_{21} and almost fixes E_{12} and T_ρ up to sign. Thus $\Phi(E_{21}^s) = E_{21}^t$, $\Phi(E_{12}^s) = E_{12}^m$ and $\Phi(T_\rho^s) = T_\rho^r$ for some $s, t > 0$ and some $r, m \neq 0$. Denote $\Phi(D_{21}^s)$ by μ and $\Phi(D_{12}^s)$ by ν .

Corollary 7.4 implies, after increasing s if necessary, that there is a basis $\{a, b\}$ of F_2 and $p, q \neq 0$ such that $\hat{\mu}_1 = i_a^p$ and $\hat{\nu}_1 = i_b^q$. Since μ commutes with E_{21}^t we have $a \in \text{Fix}(\hat{E}_{21})$. Corollary 2.19 implies that $a = i_u(x_1^\pm)$ for some $u \in \text{Fix}(\hat{E}_{21})$. The symmetric argument shows that $b = i_v(x_2^\pm)$ for some $v \in \text{Fix}(\hat{E}_{12})$. It follows that $\{x_1, i_{\bar{u}}(x_2)\}$ is a basis of F_2 and hence that $\bar{u}v = x_1^i x_2^j$ for some $i, j \in \mathbf{Z}$. In particular, $u x_1^i \in \text{Fix}(\hat{E}_{21})$ equals $v \bar{x}_2^j \in \text{Fix}(\hat{E}_{12})$. Since $\text{Fix}(\hat{E}_{21}) \cap \text{Fix}(\hat{E}_{12}) = \langle \rho \rangle$ we have $u = \rho^l \bar{x}_1^i$ and $v = \rho^l x_2^j$ for some $l \in \mathbf{Z}$. Thus $a = i_\rho^l(x_1^\pm)$ and $b = i_\rho^l(x_2^\pm)$. After replacing Φ with $\Phi' = i_{T_\rho}^{-l} \circ \Phi$, we may assume that $a = x_1^\pm$ and $b = x_2^\pm$. Thus Φ' almost fixes E_{21} and almost fixes E_{12}, T_ρ, D_{21} and D_{12} up to sign.

Lemma 7.3(2) implies that Φ' almost fixes A_3 . The roles of x_1 and x_2 are interchangeable in this argument, so there is an $O(F_2)$ -normalization Φ'' of Φ that almost fixes A_3 and E_{12} . Lemma 6.2 and Lemma 6.3 imply that $\Phi' = i_\alpha \circ \Phi''$ where $\alpha \in \mathcal{A}_3$. Lemma 7.7 then implies that Φ' almost fixes E_{12} and Lemma 7.6 completes the proof of the proposition. \square

7.2. Normalizing with respect to $O(F_{n-2})$. — The final normalizing step involves only $O(F_{n-2})$.

7.8. Proposition. — *There is a unique normalization of Φ that respects the decomposition $F_n = F_2 * F_{n-2}$, that almost fixes $A_3, \langle {}_{21}E, E_{12} \rangle$ and $\langle {}_{j2}E, E_{2j} \rangle$ for all $j \neq 2$ and that almost fixes T_ρ up to sign.*

Proof. — By Proposition 6.1 and Proposition 7.1 there is a normalization Φ_1 of Φ that respects the decomposition $F_n = F_2 * F_{n-2}$ and almost fixes $A_3, \langle {}_{21}E, E_{12} \rangle, \langle {}_{12}E, E_{12} \rangle$ and that almost fixes T_ρ up to sign. All of these properties are preserved if

Φ_1 is replaced by $i_\mu \circ \Phi$ where $\mu \in \mathbb{A}_3 \cap \mathcal{O}(\mathbb{F}_{n-2})$. We show below that for each $j \geq 4$ there exists $\mu_j \in \langle {}_{3j}\mathbb{E}, \mathbb{E}_{j3} \rangle$ such that $i_{\mu_j} \circ \Phi_1$ almost fixes $\langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle$. If $\mu = \mu_4 \circ \dots \circ \mu_n$ then $\Phi' = i_\mu \Phi_1$ satisfies the conclusions of the proposition. Uniqueness follows from Lemma 6.2 and from Lemma 7.9 below.

Fix $j \geq 4$. Proposition 6.1 and Proposition 7.1, applied with j replacing 1, imply that there exists ψ_j such that $\Phi_2 := i_{\psi_j} \circ \Phi_1$ almost fixes \mathcal{A}_3 , $\langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle$ and $\langle {}_{2j}\mathbb{E}, \mathbb{E}_{j2} \rangle$. Lemma 6.2 and Lemma 6.3 imply that $\psi_j \in \mathcal{A}_3$. Let $\eta_j = \psi_j^{-1}$. From the fact that Φ_2 almost fixes $\langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle$ we conclude that

$$(1) \quad \Phi_1(\tau^s) = i_{\eta_j} \tau^t \text{ for some } s, t > 0 \text{ and for all } \tau \in \langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle.$$

From the fact that Φ_2 almost fixes $\langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle$ we conclude that $[x_2]_u$ is the unique axis of $\Phi_2(\mathbb{E}_{j2}^p)$ and hence is the characteristic axis of $\Phi_2(\mathbb{A}_2^p)$ where p is chosen so that $\Phi_2(\mathbb{A}_2^p)$ has type E. Similarly, $[x_2]_u$ is the unique axis of $\Phi_1(\mathbb{E}_{12}^p)$ and so is the characteristic axis of $\Phi_1(\mathbb{A}_2^p)$. It follows that

$$(2) \quad [x_2]_u \text{ is } \eta_j\text{-invariant.}$$

Write η_j as a composition $\eta_j = \eta'_j \eta''_j$ where $\hat{\eta}'_j \in \hat{\mathbb{A}}_3$ is the identity on $\langle \{x_k : k \neq 2, 3, j\} \rangle$ and $\hat{\eta}''_j \in \hat{\mathbb{A}}_3$ is the identity on $\langle x_2, x_3, x_j \rangle$. Then η''_j commutes with each $\tau \in \langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle$ and preserves $[x_2]_u$. We may therefore replace η_j with η'_j and maintain (1) and (2). In other words, we may assume that $\hat{\eta}_j$ is defined by $x_2 \mapsto x_3^a x_2 x_3^b$ and $x_j \mapsto x_3^c x_j x_3^d$ for some a, b, c and d . (2) implies that $b = -a$.

Define

$$u = x_3^{c-a} x_j x_3^{d+a}.$$

Direct computation shows that $\hat{\eta}_j \hat{\mathbb{E}}_{j2}^t \hat{\eta}_j^{-1}$ is defined by

$$x_2 \mapsto x_3^{-a} (x_3^c x_j x_3^d)^{-t} x_3^a x_2 = \bar{u}^t x_2,$$

and that $\hat{\eta}_j \hat{\mathbb{E}}_{2j}^t \hat{\eta}_j^{-1}$ is defined by

$$x_2 \mapsto x_2 x_3^{-a} (x_3^c x_j x_3^d)^t x_3^a = x_2 u^t.$$

Define $\hat{v}_j \in \langle {}_{3j}\hat{\mathbb{E}}, \hat{\mathbb{E}}_{j3} \rangle$ by $x_j \mapsto x_3^{c-a} x_j x_3^{d+a}$ or equivalently $x_j \mapsto u$. Then $i_{\hat{v}_j}(\hat{\mathbb{E}}_{j2}^t)$ is defined by $x_2 \mapsto \bar{u}^t x_2$ and $i_{\hat{v}_j}(\hat{\mathbb{E}}_{2j}^t)$ is defined by $x_2 \mapsto x_2 u^t$. We conclude that $\Phi_1(\tau^s) = i_{\hat{v}_j}(\tau^t) = i_{\hat{v}_j} \tau^t$ for all $\tau \in \langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle$. Letting $\mu_j = \hat{v}_j^{-1}$, we have that $i_{\mu_j} \Phi_1$ almost fixes $\langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle$ as desired. \square

7.9. Lemma. — *If $\psi \in \text{WC}(\langle {}_{j2}\mathbb{E}, \mathbb{E}_{2j} \rangle)$ for each $j \neq 2$ then $\psi = \text{identity}$.*

Proof. — Lemma 2.17 implies that ψ fixes $[x_j]_u$ and leaves $\langle x_2, x_j \rangle$ invariant for all $j \neq 2$. Since $[x_2]_u$ is the only unoriented conjugacy class carried by both $\langle x_2, x_1 \rangle$ and $\langle x_2, x_3 \rangle$, we know that $[x_2]_u$ is ψ -invariant. By Corollary 2.5, there exists $\hat{\psi}$ such that

$\hat{\psi}(x_2) = x_2^\epsilon$ for $\epsilon = \pm 1$ and such that $\langle x_2, x_j \rangle$ is $\hat{\psi}$ -invariant for all $j \neq 2$. Lemma 2.10 implies that $\hat{\psi}(x_j) = x_2^{-a_j} x_j^{\delta_j} x_2^{a_j}$ for some $a_j \in \mathbf{Z}$ and $\delta_j = \pm 1$. Replacing $\hat{\psi}$ by $i_{x_2}^{a_1} \hat{\psi}$ we may assume that $a_1 = 0$.

Assuming now that $j > 2$, choose distinct $s, t > 0$ so that ${}_{j2}\mathbf{E}^t \circ \mathbf{E}_{2j}^s$ commutes with ψ . Since ${}_{j2}\hat{\mathbf{E}}^t \circ \hat{\mathbf{E}}_{2j}^s$ fixes x_1 and $\hat{\psi}(x_1) = x_1^\pm$, it follows that $[{}_{j2}\hat{\mathbf{E}}^t \circ \hat{\mathbf{E}}_{2j}^s, \hat{\psi}] = i_{x_1}^l$ for some $l \in \mathbf{Z}$. We now compute

$$\begin{aligned} \hat{\psi}({}_{j2}\hat{\mathbf{E}}^t \circ \hat{\mathbf{E}}_{2j}^s)(x_2) &= \hat{\psi}(x_j^{-t} x_2 x_j^s) = (x_2^{-a_j} x_j^{-\delta_j t} x_2^{a_j}) x_2^\epsilon (x_2^{-a_j} x_j^{\delta_j s} x_2^{a_j}) \\ &= x_2^{-a_j} x_j^{-\delta_j t} x_2^\epsilon x_j^{\delta_j s} x_2^{a_j} \end{aligned}$$

and

$$({}_{j2}\hat{\mathbf{E}}^t \circ \hat{\mathbf{E}}_{2j}^s) \hat{\psi}(x_2) = ({}_{j2}\hat{\mathbf{E}}^t \circ \hat{\mathbf{E}}_{2j}^s)(x_2^\epsilon) = x_j^{-t} x_2 x_j^s \text{ or } x_j^{-s} x_2^{-1} x_j^t$$

depending on whether $\epsilon = 1$ or $\epsilon = -1$.

It follows that $l = 0$, $a_j = 0$ and $\delta_j = \epsilon = 1$ which proves that $\hat{\psi}$ is the identity. \square

8. Moving between bases

We say that a normalization Φ' of Φ *almost fixes a basis* \mathbf{B} of F_n if it almost fixes each $\langle {}_{ji}\mathbf{E}, \mathbf{E}_{ij} \rangle$ defined with respect to that basis. The goal of this section is to prove the following.

8.1. Proposition. — *There is a normalization of Φ that almost fixes every basis of F_n .*

Combining Proposition 8.1 with Lemma 3.2 immediately gives the main theorem of this paper, namely Theorem 1.1. We divide the proof of Proposition 8.1 into a number of steps.

Step 1 (Normalizing on any basis). — We begin with the much weaker claim that any basis can be fixed by some normalization (depending on that basis).

8.2. Lemma. — *Each basis \mathbf{B} is almost fixed by a unique normalization Φ' of Φ . If $\mathbf{B} = \{x_1, \dots, x_n\}$ and if $\rho = [x_1, x_2]$ then Φ' almost fixes Γ_ρ up to sign.*

Proof. — The normalization Φ' given by Proposition 7.8 applied to \mathbf{B} almost fixes \mathcal{A}_3 , $\langle {}_{21}\mathbf{E}, \mathbf{E}_{12} \rangle$ and $\langle {}_{j2}\mathbf{E}, \mathbf{E}_{2j} \rangle$ for $j \neq 2$ and almost fixes Γ_ρ up to sign. By Proposition 7.8 with 1 replaced by $l \neq 2$ and 3 replaced by $k \neq 2$, there exists ψ_k so that $i_{\psi_k} \Phi'$ almost fixes \mathcal{A}_k , $\langle {}_{2l}\mathbf{E}, \mathbf{E}_{l2} \rangle$ and $\langle {}_{j2}\mathbf{E}, \mathbf{E}_{2j} \rangle$ for $j \neq 2$. Lemma 6.2 and Lemma 7.9 imply that ψ_k is the identity, and hence that Φ' almost fixes $\langle {}_{ki}\mathbf{E}, \mathbf{E}_{ik} \rangle$ for $k \neq 2$ and $i \neq k$ and almost fixes $\langle {}_{2l}\mathbf{E}, \mathbf{E}_{l2} \rangle$ for $l \neq 2$. This proves that Φ' almost fixes \mathbf{B} . Uniqueness follows from Lemma 6.2 and Lemma 7.9. \square

Step 2 (The Farey graph and the set of bases). — We will need to understand the set of all bases of F_2 . A useful tool to do this is the Farey graph, which we now recall.

Recall from Section 2.6 that the natural homomorphism from the extended mapping class group of the once-punctured torus S to $\text{Out}(F_2)$ is an isomorphism. Further, the set \mathcal{S} of isotopy classes of essential, nonperipheral simple closed curves on S are in bijective correspondence with the set \mathcal{C} of (unoriented) conjugacy classes of basis elements of F_2 . A marking on S also induces a bijective correspondence between \mathcal{S} and $\mathbf{Q} \cup \infty$, where (p, q) , which is identified with $\frac{p}{q} \in \mathbf{Q}$, represents the “slope” of the corresponding element in \mathcal{S} , that is the simple closed curve representing the element $(p, q) \in H_1(S, \mathbf{Z}) \approx \mathbf{Z} \times \mathbf{Z}$. We assume that $[x_1]_u$ corresponds to $(1, 0)$ and that $[x_2]_u$ corresponds to $(0, 1)$.

The *Farey graph*, denoted \mathcal{F} , is defined to be the graph with one vertex for each element of \mathcal{S} , and with an edge connecting (p, q) to (r, s) when $|ps - rq| = 1$. Note that this is equivalent to the corresponding curves on S having geometric intersection number one and, more importantly, it happens precisely when the associated unoriented conjugacy classes can be represented by *cobasis elements*, which means that together they generate F_2 .

There is a standard embedding of \mathcal{F} into the hyperbolic disc D^2 defined by embedding $\mathbf{Q} \cup \infty$ into S^1 in the obvious way and then connecting (p, q) to (r, s) for $|ps - rq| = 1$ with the unique hyperbolic geodesic between them. This gives the well-known *Farey tessellation* of D^2 , denoted $\widehat{\mathcal{F}}$ which is a (not locally finite) 2-dimensional simplicial complex K .

We would like to pin down general set maps $F_2 \rightarrow F_2$ using purely combinatorial information about their action on basis elements. The usefulness of the Farey graph is that it converts this problem into a geometric one.

8.3. Lemma (Farey lemma). — *Let $h : \mathcal{S} \rightarrow \mathcal{S}$ be any bijective map. Suppose that if c_1 and c_2 are represented by cobasis elements then so are $h(c_1)$ and $h(c_2)$. Suppose further that h fixes $(0, 1)$ and $(1, 0)$, and that $h(s, 1) = (t, 1)$ for some $s, t > 0$. Then h is the identity map.*

Proof. — We use the $\mathbf{Q} \cup \infty$ notation. Let σ denote the 2-simplex in $\widehat{\mathcal{F}}$ with vertices $(1, 0)$, $(0, 1)$ and $(1, 1)$. Every edge of \mathcal{F} is a face of precisely two 2-simplices in $\widehat{\mathcal{F}}$. From this, an easy induction on combinatorial distance to σ gives that an automorphism of \mathcal{F} is completely determined by its action on σ .

There is no loss in identifying h with its induced automorphism of \mathcal{F} . By hypothesis, h fixes $(0, 1)$ and $(1, 0)$ so it suffices to show that $h(1, 1)$ is $(1, 1)$ rather than $(-1, 1)$. The edge e of \mathcal{F} that connects $(0, 1)$ and $(1, 0)$ separates \mathcal{F} into two components, one containing all the positive slopes and the other containing all the negative slopes. It therefore suffices to show that h setwise fixes the components of the complement of e . This is immediate from our hypothesis that $h(s, 1) = (t, 1)$ for some $s, t > 0$. \square

Suppose that a basis has been chosen and that Φ' respects the decomposition $F_n = F_2 * F_{n-2}$. Then Φ' induces a self-map $\Phi'_\#$ of \mathcal{C} as follows. Given $c \in \mathcal{C}$, choose a primitive $\mu_1 \in \text{Out}(F_2)$ that fixes c and is UL. In other words, think of c as an unoriented simple closed curve on S and let μ_1 be the Dehn twist about this curve. By Lemma 7.2, there is a unique $\hat{\mu}_1 \in \text{Aut}(F_2)$ such that $\hat{\mu} = \hat{\mu}_1 \times \text{Id} \in \text{O}(F_2)$ represents μ_1 , fixes ρ and is elementary. Choose $s > 0$ so that $\mu' = \Phi'(\mu^s) \in \text{O}(F_2)$ is elementary. Lemma 7.2 and Lemma 2.18 imply that μ'_1 fixes some $c' \in \mathcal{C}$. Define $\Phi'_\#(c) = c'$. Thus μ'_1 is a Dehn twist about an unoriented simple closed curve representing c' . As s varies, the resulting μ'_1 belong to a cyclic subgroup of $\text{Out}(F_2)$, which shows that c' is independent of s and $\Phi'_\#$ is well defined.

If, for example, $c = [x_1]_u$ then $\mu = E_{21}$. If Φ' almost fixes E_{21} , then $\mu' = E'_{21}$ and $c' = c$. Similarly, if Φ' almost fixes E_{12} then $\Phi'_\#$ fixes $[x_2]_u$.

Corollary 5.6 implies that if c_1 and c_2 are represented by cobasis elements then so are $\Phi'_\#(c_1)$ and $\Phi'_\#(c_2)$. Thus $\Phi'_\#$ induces an automorphism of \mathcal{F} or what is the clearly the same thing, a simplicial automorphism of K .

8.4. Lemma. — *If Φ' almost fixes E_{21} and E_{12} then $\Phi'_\#$ is the identity.*

Proof. — As noted above $\Phi'_\#$ fixes $(0, 1)$ and $(1, 0)$ so $\Phi'_\#(1, 1)$ is either $(1, 1)$ or $(-1, 1)$. By construction, E_{21} corresponds to a Dehn twist about the $(1, 0)$ curve and E_{12} corresponds to a Dehn twist about the $(0, 1)$. There exist $s, t, q > 0$ so that $\Phi(E_{21}^s) = E'_{21}$ and $\Phi(E_{12}^s) = E_{12}^q$. Then $E_{21}^s E_{12}^s E_{21}^{-s}$ corresponds to a Dehn twist of order s about the $(s, 1)$ curve and

$$\Phi(E_{21}^s E_{12}^s E_{21}^{-s}) = E'_{21} E_{12}^q E_{21}^{-t}$$

corresponds to a Dehn twist of order q about the $(t, 1)$ curve. Thus $\Phi'_\#(s, 1) = (t, 1)$. Lemma 8.3 completes the proof. \square

Step 3 (Normalizing on an adjacent basis). — The above results on automorphisms of the Farey graph can be used to show how a normalization of Φ on one basis constrains the Φ -image of an “adjacent” basis, as follows.

8.5. Corollary. — *Suppose that Φ' is the unique normalization that almost fixes the basis B defined by $\{x_1, \dots, x_n\}$. If B' is the basis defined from B by replacing x_2 with $x_2 x_1$ and if μ is E_{12} defined with respect to B' then Φ' almost fixes μ up to sign.*

Proof. — Choose $s > 0$ so that $\mu' = \Phi'(\mu^s)$ is elementary. We consider $\Phi'_\#$ defined with respect to B and let $c = [x_2 x_1]_u$. Since $\hat{\mu}_1$ fixes ρ and fixes $x_2 x_1$, $\Phi'_\#(c)$ is defined to be the element of \mathcal{C} that is fixed by μ'_1 . Lemma 8.4 implies that $\Phi'_\#(c) = c$ and hence that μ'_1 fixes $[x_2 x_1]_u$. Thus μ_1 and μ'_1 belong to the same cyclic subgroup of $\text{Out}(F_2)$. Since Φ' almost fixes T_ρ up to sign, we have that μ' commutes with T_ρ which implies that $\hat{\mu}'_1$ fixes ρ . Lemma 7.2 therefore implies that $\hat{\mu}_1$ and $\hat{\mu}'_1$ belong to the same cyclic subgroup of $\text{Aut}(F_2)$. Thus $\Phi'(\mu^s) = \mu^t$ for some $t \neq 0$. \square

Step 4 (Normalizing on all bases). — In order to prove that there is a normalization of Φ which almost fixes every basis, we give the following sufficient condition for a basis to be almost fixed.

8.6. Lemma. — *Assume that definitions are made relative to a basis $\{x_1, \dots, x_n\}$ denoted B . If Φ' almost fixes \mathcal{A}_1 and $\langle_j E, E_{ij} \rangle$ for $i, j \geq 3$ and if Φ' almost fixes E_{12} up to sign, then Φ' almost fixes B .*

Proof. — Choose a normalization Φ'' that fixes B and $\psi \in \text{Out}(F_n)$ such that $\Phi' = i_\psi \circ \Phi''$. It suffices to show that ψ is the identity. Lemma 6.2 and Lemma 6.3 imply that $\psi \in \mathcal{A}_1$ and hence that ψ is UL with $[x_1]_u$ as its unique axis. They also imply, in conjunction with Lemma 2.12 and Lemma 2.17, that $[x_i]_u$ and $[\langle x_i, x_j \rangle]$ are ψ -invariant for all $i, j \geq 3$.

Let $A := \{[x_i], [x_i x_j] : i \neq j \geq 3\}$ and suppose that $F_{n-2} = F^1 * F^2$ where each element of A is carried by either F^1 or F^2 . If x_i is carried by F^1 and x_j is carried by F^2 then $[x_i x_j]$ is not carried by either F^1 or F^2 . It follows that either F^1 or F^2 carries each $[x_i]$ and so has rank at least $n - 2$. This proves that the decomposition is trivial and hence that F_{n-2} is the minimal carrier of A . Since $\psi(a)$ is carried by F_{n-2} for each $a \in A$, $\psi^{-1}[F_{n-2}]$ is also a minimal carrier of A . By uniqueness, $[F_{n-2}]$ is ψ -invariant.

The restriction $\psi|_{[F_{n-2}]}$ is trivial because F_{n-2} does not carry the unique axis of ψ . Thus there exists a representative $\hat{\psi}$ defined by $x_2 \mapsto x_1^p x_2 x_1^q$ for some p, q . Since Φ' almost fixes E_{12} up to sign, $\hat{\psi} \hat{E}_{12}^s \hat{\psi}^{-1} = \hat{E}_{12}^t$ for some $s > 0$ and some $t \neq 0$. It follows from

$$\hat{\psi} \hat{E}_{12}^s \hat{\psi}^{-1}(x_1) = \hat{\psi} \hat{E}_{12}^s(x_1) = \hat{\psi}(x_1 x_2^s) = x_1 (x_1^p x_2 x_1^q)^s$$

that $p = q = 0$ so $\hat{\psi}$ is the identity as desired. \square

With the above in hand we are now ready to prove the main result of this section.

Proof of Proposition 8.1. — By Lemma 8.2 it suffices to show that if Φ' almost fixes some basis then it almost fixes every basis. Suppose that x_1, \dots, x_n is an almost fixed basis B . It is immediate from the definitions that permuting the x_i 's or replacing some x_i with \bar{x}_i preserves the property of being an almost fixed basis. It suffices to show that the basis B' obtained from B by replacing x_2 with $x_2 x_1$ is almost fixed because these moves generate $\text{Aut}(F_n)$ and there is an automorphism carrying any one basis to any other basis.

Denote E_{12} , defined relative to B' , by μ . We have to verify the hypotheses of Lemma 8.6 with respect to B' . This is obvious except for showing that Φ' almost fixes μ up to sign, which is proved in Corollary 8.5. \square

REFERENCES

- [BFH1] M. BESTVINA, M. FEIGN and M. HANDEL, The Tits alternative for $\text{Out}(F_n)$, I: Dynamics of exponentially-growing automorphisms, *Ann. Math. (2)*, **151** (2000), 517–623.
- [BFH2] M. BESTVINA, M. FEIGN and M. HANDEL, The Tits alternative for $\text{Out}(F_n)$ II: A Kolchin type theorem, *Ann. Math. (2)*, **161** (2005), 1–59.
- [BFH3] M. BESTVINA, M. FEIGN, and M. HANDEL, Solvable subgroups of $\text{Out}(F_n)$ are virtually Abelian, *Geom. Dedicata*, **104** (2004), 71–96.
- [BH] M. BESTVINA and M. HANDEL, Train tracks and automorphisms of free groups, *Ann. Math. (2)*, **135** (1992), 1–51.
- [BuH] M. BURGER and P. DE LA HARPE, Constructing irreducible representations of discrete groups, *Proc. Indian Acad. Sci., Math. Sci.*, **107** (1997), 223–235.
- [BV] M. BRIDSON and K. VOGTMANN, Automorphisms of automorphism groups of free groups, *J. Algebra*, **229** (2000), 785–792.
- [DF] J. DYER and E. FORMANEK, The automorphism group of a free group is complete, *J. Lond. Math. Soc., II. Ser.*, **11** (1975), 181–190.
- [FH] M. FEIGN and M. HANDEL, Abelian subgroups of $\text{Out}(F_n)$, preprint, December 2006.
- [FP] E. FORMANEK and C. PROCESI, The automorphism group of a free group is not linear, *J. Algebra*, **149** (1992), 494–499.
- [IM] N. IVANOV and J. MCCARTHY, On injective homomorphisms between Teichmüller modular groups, I, *Invent. Math.*, **135** (1999), 425–486.
- [Iv] N. IVANOV, Mapping class groups, in *Handbook of Geometric Topology*, R. Daverman and R. Sher (eds.), pp. 523–633, North Holland, Amsterdam (2002).
- [Iv2] N. IVANOV, Automorphisms of complexes of curves and of Teichmüller spaces, *Int. Math. Res. Not.*, **1997** (1997), 651–666.
- [Kh] D. G. KHRAMITSOV, Completeness of groups of outer automorphisms of free groups, in *Group-Theoretic Investigations* (Russian), pp. 128–143, Akad. Nauk SSSR Ural. Otdel., Sverdlovsk (1990).
- [Ma] G. A. MARGULIS, *Discrete Subgroups of Semisimple Lie Groups*, Springer, Berlin (1990).
- [Pr] G. PRASAD, Discrete subgroups isomorphic to lattices in semisimple Lie groups, *Am. J. Math.*, **98** (1976), 241–261.
- [Vo] K. VOGTMANN, Automorphisms of free groups and outer space, *Geom. Dedicata*, **94** (2002), 1–31.
- [Zi] R. ZIMMER, *Ergodic Theory and Semisimple Groups*, Monogr. Math., vol. 81, Birkhäuser, Basel (1984).

B. F.

Department of Mathematics,
University of Chicago,
5734 University Ave.,
Chicago, IL 60637, USA
farb@math.uchicago.edu

M. H.

Department of Mathematics,
Lehman College,
Bronx, NY 10468, USA
michael.handel@lehman.cuny.edu

*Manuscrit reçu le 15 octobre 2006
publié en ligne le 14 septembre 2007.*