Monotonicity and the Construction of Quasiconformal Conjugacies in the Real Cubic Family

A Dissertation Presented

by

Christopher Arthur Heckman

to

The Graduate School in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

in

Mathematics

State University of New York

at Stony Brook

August 1996

State University of New York
at Stony Brook

The Graduate School

Christopher Arthur Heckman

We, the dissertation committee for the above candidate for the Doctorate of Philosophy degree, hereby recommend acceptance of this dissertation.

Marie-Louise Michelsohn
Professor of Mathematics
Dissertation Director

Scott Sutherland
Lecturer and Director of Computing
Chairman of Defense

Adam Epstein
Assistant Professor of Mathematics

Folkert Tangerman
Visiting Assistant Professor
Department of Applied Mathematics and Statistics
Outside Member

This dissertation is accepted by the Graduate School.

Graduate School

Abstract of the Dissertation

Monotonicity and the Construction of Quasiconformal Conjugacies in the Real Cubic Family

by

Christopher Arthur Heckman

Doctor of Philosophy

in

Mathematics

State University of New York

at Stony Brook

1996

The dependence of the dynamical behavior of polynomials on their parameters can be understood by studying topological conjugacies between the polynomials. Topological conjugacy forms an equivalence relation on polynomials, and polynomials with the same type of dynamical behavior are in the same equivalence class. In the real quadratic family, analysis of the single critical orbit, using quasiconformal pull-back techniques and a type of induced map called a box mapping, shows that topological conjugacies are quasiconformal homeomorphisms. If there is a quasiconformal conjugacy between two distinct real quadratic maps, it is always possible to deform this conjugacy into an entire family of conjugacies between real quadratic maps. This characterizes conjugacy classes of real quadratic maps, containing more

than one member, as connected and open intervals in parameter space. This provides enough information to show that topological entropy must vary monotonically in real quadratic parameter space. This is the *monotonicity* problem for real quadratic maps.

This thesis applies the same tools to the real cubic family and solves a similar monotonicity problem. We show that any set of real cubic maps with the same topological entropy forms a connected set in parameter space. If one critical orbit of a cubic map is periodic, Yoccoz partitions can be used to construct a box mapping induced from the cubic map, which then yields information about the second critical orbit. Certain topological conjugacy classes of real cubic maps are shown to be quasiconformal, and the deformation argument applied to cubic maps then yields the monotonicity result.

Table of Contents

I. Introduction

§1.1. Cubic monotonicity.

We prove a Monotonicity Conjecture for the family of real cubic maps. This conjecture states that for a suitably normalized space (parameter space) of real cubic polynomials, the points representing cubic maps with some fixed topological entropy always form a connected set.

The work of Dawson, Galeeva, Milnor, and Tresser (see [DGMT]) provides the framework and the principal reduction of the conjecture for real cubic maps. This work introduces the concept of *bones*, which are subsets of the parameter space where the dynamics of one critical point are periodic with specified order type. Bones of high enough period are shown to be one-dimensional manifolds that form a skeletal structure underlying the entire parameter space. This structure is related to a similar structure in the parameter space of certain piecewise linear maps, where the question of monotonicity can be directly addressed. This strategy hinges on the requirement that bones themselves be connected subsets of the cubic parameter space. In [DGMT], this is stated as a conjecture, the Connected Bone Conjecture, and it is shown that this implies monotonicity for real cubic maps. This paper approaches the monotonicity problem by proving the Connected Bone Conjecture. The main tools used in the proof, deformation of quasiconformal structures, Yoccoz partitions, and box mappings, are the same as those developed for studying the parameter space of quadratic polynomials.

§1.2. History.

The monotonicity problem itself was originally posed (see [MT]) for the real quadratic family. Since the parameter space of real quadratic polynomials is one-dimensional, topological entropy, considered as a continuous function on parameter space, must be monotonic if each locus of constant entropy forms a connected set. Hence the name "monotonicity."

Monotonicity of the real quadratic family is the first in a progression of results working towards the ultimate goal of proving (Fatou's Conjecture) that hyperbolic maps are dense within any family of polynomials (See the discussion of [DGMT] of *General Hyperbolicity* and how it relates to the Connected Bone Conjecture and monotonicity). The principal tools used here are the analysis of conjugacies between polynomials as quasiconformal maps, and Yoccoz partitions of the dynamical plane. The fundamental observation is that two hyperbolic quadratic maps which are in the same conjugacy class are conjugate by a quasiconformal map and, in general, such a conjugacy class must be an open set in parameter space or just a single point. Yoccoz partitions then provide a framework for studying the conjugacy.

Together these tools have been used to prove Fatou's Conjecture for the real quadratic family. It is easily seen that conjugacy classes for non-hyperbolic maps must be closed sets, and if all the conjugacy classes are quasiconformal, then non-hyperbolic classes must be single points in parameter space. It then follows that hyperbolic maps are dense. Swiatek reduces the hyperbolicity problem to the following statement about quasiconformal conjugacy classes (See [Sw]).

Theorem 1.1. (Swiatek) Let f and $\hat{f}$ be two real quadratic polynomials with a bounded forward critical orbit and no attracting or indifferent cycles. Then, if they are topologically conjugate, the conjugacy extends to a quasiconformal conjugacy between their analytic continuations to the complex plane.

This theorem and the tools used to prove it form the starting point for this work. The advantage of studying bones in cubic parameter space is that one of two critical points is effectively removed from consideration. The bulk of the tools used to prove the above theorem, which all draw conclusions about a single critical orbit of an analytic map, can be applied *directly* to cubic maps on bones. These tools come in the form of *box mappings*, *branchwise equivalences*, and Theorems [6.1] and [6.2], which are introduced in Chapter VI. The chapters following all discuss how cubic maps can be worked into the framework of box mappings. Chapter VII introduces additional tools, including the Böttcher coordinate and the λ-Lemma. In Chapter VIII, we construct box mappings and quasiconformal branchwise equivalences. Chapter IX applies Theorems [6.1] and [6.2] to box mappings constructed using cubic maps, and a result similar to Theorem [1.1] is obtained.

Chapters III through V show how quasiconformal conjugacies can be used to prove the Connected Bone Conjecture. Chapter III defines the cubic parameter space and bones. Chapter IV develops the essential quasiconformal tools for applying the box mapping results to cubic conjugacy classes, and Chapter V uses these tools to prove that bones are connected.

II. Basic Terminology

§2.1. Holomorphic dynamics.

For an introduction to holomorphic dynamics, see [Mi2] or [Bl]. The map f is always an analytic map of the complex plane, and we study its *iterates* $f^{\circ j}$, where $f^{\circ 2}(z) = f \circ f(z)$, $f^{\circ 3}(z) = f \circ f \circ f(z)$, etc. The **orbit** of a point z under f is the sequence of points $\{f^{\circ j}(z)\}_{j=0}^{\infty}$, and we are concerned with the long term behavior of this sequence and the limit, if it exists. If the sequence of points repeats itself after a finite number of iterations, we say z is **periodic**, and any points that contain a periodic point in their orbit are called **preperiodic**. Note that a periodic point z of f is **fixed** under iteration by $f^{\circ j}$, if we select j to be the size of the period of z under f. If z is fixed under $f^{\circ j}$, we call z an **attracting** periodic point, if some neighborhood of points near z have orbits which converge to z under $f^{\circ j}$. This behavior is characterized by the *multiplier*, the derivative of $f^{\circ j}$ evaluated at the fixed point, having absolute value strictly less than one. The set of points that have this same limit are called the **basin** of z, this will also be the basin of any point in the orbit of z. A basin is always an open set, and the connected components containing the finite orbit of z is called the **immediate basin** of z. A periodic point is **repelling** if the multiplier has absolute value strictly greater than one, in which case points near z will always leave small neighborhoods of z under iteration by $f^{\circ j}$. We will call a preperiodic point that contains a repelling periodic point in its orbit **repelling preperiodic**. Since the multiplier varies continuously with the parameters of the polynomial, attracting and repelling periodic points always persist with small perturbations in the map f.

For this work, the map f will always be a polynomial, and if the point z is large enough, the orbit of z will diverge to infinity. We call the set of points in $\mathbb{C}$ which diverge under f, the **basin at infinity**, and we also say these orbits *escape* to infinity. For a polynomial f, the boundary of the basin at infinity is called the **Julia Set** of f, and the complement of the Julia set is the **Fatou Set**. Since all points in the Julia set are boundary points of part of the Fatou set, the Fatou set is dense in $\mathbb{C}$. Both the Fatou and Julia sets are forward and backward invariant under f. The connected components of the Fatou set must map onto each other under the action of f. Thus it is possible to talk of a periodic or preperiodic Fatou component. The periodic Fatou components typically form the immediate basins of attracting periodic points. For real polynomials in particular, if there are no **indifferent** periodic points, points whose multiplier has absolute value equal to one, then all components of the Fatou set are part of some attracting basin.

A famous theorem of Sullivan (see [Su2]) states that all Fatou components are preperiodic. This is a form of *wandering domain* theorem. A **wandering domain** is a connected open set, which is not contained in the basin of an attracting periodic point, whose images under iterates of a map f are all disjoint. Wandering domains do not exist for polynomials. This is also true of open intervals on the real line, where the polynomial is a real map. There are no wandering intervals for real polynomials. In general, real maps have analogous definitions for attracting basins considered as intervals, although we must be careful of indifferent periodic points which may be attracting on the real line but partially repelling in the plane. If there are no indifferent periodic orbits however, connected components of real attracting basins are always the intersection of some Fatou component with the real line. A polynomial is called **hyperbolic** if every critical orbit is contained in the basin of some attracting periodic point.

§2.2. Topological conjugacy.

A **topological conjugacy** between two polynomials f and $\hat{f}$ is a homeomorphism H of either $\mathbb{R}$ or $\mathbb{C}$ onto itself which satisfies the *functional equation*

$$H \circ f = \hat{f} \circ H.$$

Conjugacies form an equivalence relation on polynomial families, and many dynamical properties are characterized up to conjugacy class. An important observation is that conjugacies must map orbits of critical points under f to orbits of critical points under $\hat{f}$. For real polynomials, the order in which a critical orbit occurs on the real line is key to determining whether a conjugacy exists between two polynomials on the real line.

Definition. (combinatorial equivalence) Two real polynomials, f and $\hat{f}$, are **combinatorially equivalent** if there is a one to one correspondence between the real critical points of each map, $\{c_n\}$ to $\{\hat{c}_n\}$, which satisfies the following condition. For any n_1, n_2, i, and j,

$$f^{\circ i}(c_{n_1}) < f^{\circ j}(c_{n_2}) \qquad \Longrightarrow \qquad \hat{f}^{\circ i}(\hat{c}_{n_1}) < \hat{f}^{\circ j}(\hat{c}_{n_2}).$$

According to de Melo and van Strien, for a large class of C^3 maps of the interval, two maps are combinatorially equivalent if and only if they are topologically conjugate. In particular, this is true for bimodal cubic maps with no indifferent periodic orbits (see [dMvS], p. 157).

§2.3. Kneading theory.

A standard method of describing qualitative information about the dynamical class of a point with respect to a real polynomial is through kneading sequences. Suppose f is a real cubic (bimodal) polynomial and x is a point on the real line. We label each point in the orbit of x under f depending on where it falls on $\mathbb{R}$ with respect to the critical points of f. The possibilities are c_R and c_L, if the point is equal to the right or left critical point, or L, M, and R, if the point falls to the left of, in the middle of, or to the right of both critical points. For work with real cubic maps, there are thus a total of five **kneading symbols**, two *critical* symbols and three *non-critical* symbols. The sequence of kneading symbols generated by the orbit of a single point x is called the **kneading sequence** of x with respect to f. The entire kneading sequence is denoted by $\theta_f(x)$, and the kneading symbol corresponding to the n-th point in the orbit is specified by $\theta_f^n(x)$. We list some of the basic properties of kneading sequences. For more details see [MT].

For a given map f, there is a standard ordering among kneading sequences, inherited from the natural ordering of the kneading symbols themselves on the real line, which satisfies

$$x < y \qquad \Longrightarrow \qquad \theta_f(x) \leq \theta_f(y).$$

A point x has a kneading sequence that is eventually periodic if and only if the orbit of x converges to a periodic orbit. If the periodic orbit is attracting and does not contain x, small changes in x will not result in a change in kneading sequence $\theta_f(x)$.

FACT. Suppose x and f change continuously such that $\theta_f^n(x)$ is represented by two different symbols. Then as x and f are varied, $\theta_f^n(x)$ must be represented by all symbols in between. In particular, the only way for $\theta_f^n(x)$ to change from one non-critical symbol to another is for it to become a critical symbol first.

§2.4. Quasiconformal maps.

A fundamental observation about conjugacies between holomorphic maps (due to Sullivan) is that they are frequently *quasiconformal.* For the complex function H, let H^z and $H^{\bar{z}}$ be the partial derivatives of H with respect to z and $\bar{z}$.

Definition. (quasiconformal maps) Let H be a homeomorphism from $\hat{\mathbb{C}}$ onto itself. The map H is K-quasiconformal if its *complex dilatation* χ_H, which is defined to be $H^{\bar{z}}/H^{z}$ satisfies

$$|\chi_H| \leq k \qquad \text{a.e. where} \quad k = \frac{K-1}{K+1},$$

and if H is absolutely continuous on almost every line parallel to the real or imaginary axis.

We list some standard properties of quasiconformal maps. Complete proofs can be found in [Ah]. The family of quasiconformal maps contains all the conformal maps, which are 1-quasiconformal, within it. In fact, all 1-quasiconformal maps are conformal. Compositions of quasiconformal maps are quasiconformal, and composing a quasiconformal map with a conformal map does not change the bounds on complex dilatation at all. This property is crucial for the refinement of partial conjugacies (see [10.1]) presented in Chapters VI and VII. If the family of K-quasiconformal maps is normalized at three points, i.e. all the maps are equal on the same three points, then the family is normal, and the limit functions of convergent sequences are also K-quasiconformal. On the existence of quasiconformal maps with a prescribed complex dilatation, we have the famous Measurable Riemann Mapping Theorem.

Theorem 2.1. Let $\mu : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ be a measurable function with $\|\mu\|_\infty < 1$. Then there exists a quasiconformal mapping $H : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ with $\chi_H = \mu$ almost everywhere. Moreover, if H is fixed at three points in $\hat{\mathbb{C}}$, it is unique, and it depends analytically on μ.

The fact that the partial derivatives of a quasiconformal map only need to be measurable makes it easy to piece together different maps into a new quasiconformal map. A quasiconformal homeomorphism defined on the unit disk can be extended to a quasiconformal homeomorphism of the entire sphere if the image of the disks boundary is a *quasicircle*. The issue of gluing together maps and extensions thus comes down to the properties of the boundaries along which the maps are cut. In the situations we need to glue maps (see Theorems [4.6] and [4.7]) the boundaries are all constructed out of smooth curves from which quasiconformal extensions exist.

§2.5. Böttcher coordinates.

The classical Theorem of Böttcher, concerning the *straightening* of analytic functions near periodic critical points, plays a central role in the dynamics of complex polynomials and is crucial for the results described in this work. We present it here in a form suitable for polynomials where it is applied to a neighborhood of infinity considered as a fixed point of the polynomial.

Theorem 2.2. Suppose we have a polynomial

$$f(z) = a_n z^n + a_{n-1} z^{n-1} + \cdots + a_0,$$

where $n \geq 2$ and $a_n \neq 0$. Then there exists a local holomorphic change of coordinate $\varphi(z)$ which conjugates f with the map $Z(z) = z^n$ throughout some neighborhood of $\varphi(\infty) = \infty$. The map φ is unique up to multiplication by an $(n-1)$-st root of unity.

See [Mi2] for instance for a proof. We will refer to φ as the **coordinate function** or **Böttcher coordinate** of f. The theorem always gives us the functional equation

$$\varphi \circ f(z) = Z \circ \varphi(z) \qquad \text{for } z \in U, \text{ some neighborhood of infinity.}$$

We will need an extension to this theorem, showing that within a polynomial family, the coordinate function varies analytically with the parameters of the family, and, if the polynomial is real, the coordinate can be chosen to be symmetric with respect to $\mathbb{R}$.

Theorem 2.3. Let $\{f_\lambda\}_{\lambda \in D}$ be a family of polynomials of degree n, which must contain at least one real polynomial, and where for fixed z, $f_\lambda(z)$ varies analytically in λ. Suppose Böttcher coordinates, φ_λ, can be defined for each f_λ on some common neighborhood of infinity, U. Then φ_λ can be chosen so that

(1) For fixed $z \in U$, $\varphi_\lambda(z)$ varies analytically with λ.
(2) Each φ_λ is univalent.
(3) If f_λ is a real polynomial then φ is symmetric, i.e. $\varphi(\bar{z}) = \bar{\varphi}_\lambda(z)$, and φ preserves the orientation of $\mathbb{R}$.

We have included a proof in the appendix. We reserve the symbol φ for the Böttcher coordinates chosen in this way. It is clear that, for real polynomials, the choice of a symmetric coordinate which preserves the orientation of $\mathbb{R}$ is unique, and within a family, since the coordinate varies continuously, the coordinates of non-real polynomials are uniquely determined. We use the Böttcher coordinate primarily for defining *rays* and *potentials*, subsets of the domain of polynomial maps.

Definition. (radial lines) A **radial line** of angle θ down to a radius of $r \geq 1$ is defined to be the set of points

$$\hat{\mathcal{R}}_\theta^r = \{se^{i\theta} \in \mathbb{C} \mid s > r\}.$$

Note that all radial lines are in $\hat{\mathbb{C}} - \bar{D}$, where $D \subset \hat{\mathbb{C}}$ is the unit disk. We refer to an *entire* radial line of angle θ as $\hat{\mathcal{R}}_\theta = \hat{\mathcal{R}}_\theta^1$.

Definition. (external rays and potentials) Let f be a polynomial, and let φ be a Böttcher coordinate which is both *well-defined* and *univalent* on some set U. Suppose the range of φ contains the radial line $\hat{\mathcal{R}}_\theta^r$. Then the **external ray** or **ray** of f of angle θ down to a radius of $r \geq 1$ is defined to be the set of points

$$\mathcal{R}_\theta^r = \varphi^{-1}(\hat{\mathcal{R}}_\theta^r).$$

Suppose the range of φ contains the circle of radius r centered at the origin. Then the **potential curve** of f of radius r is defined to be the set of points

$$G^r = \varphi^{-1}(\{z \in \hat{\mathbb{C}} - \bar{D} \mid |z| = r\}).$$

Rays and potentials are always in the basin of infinity of f and depend on the exact choice of φ. We refer to an *entire* external ray of angle θ as $\mathcal{R}_\theta = \mathcal{R}_\theta^1$.

III. Foundations

§3.1. Normalization of the cubic family.

We are concerned with the family of real cubic maps,

$$A_0x^3 + A_1x^2 + A_2x + A_3 \qquad A_0,\, A_1,\, A_2,\, A_3 \in \mathbb{R}.$$

Standard normalizations restrict this family so that we need not consider two maps in the same conformal conjugacy class. I.e. no two maps in the restricted family, f and $\hat{f}$, should satisfy the equation

$$f = L^{-1} \circ \hat{f} \circ L \qquad L(x) = ax + b.$$

We effectively lose two "degrees of freedom" from the original cubic family and end up with a family in two independent variables, which take values in some sub-manifold of $\mathbb{R} \times \mathbb{R}$ (or $\mathbb{C} \times \mathbb{C}$). This manifold will be referred to as a **parameter space** of the family, and it will, of course, depend on exactly how we choose to parameterize the family.

We define our restricted family, $\mathcal{F}$, by requiring that each cubic polynomial be a bimodal map of the unit interval. This normalization follows [DGMT] in particular.

Definition. (the family $\mathcal{F}$) $f \in \mathcal{F}$ satisfies

(1) $f : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is a real cubic map.

(2) f maps $[0,1]$ into itself and $f|_{[0,1]}$ is bimodal, i.e. has two distinct critical points in $(0,1)$.

(3) The boundary of $[0,1]$ is mapped onto itself by f.

This restriction requires the coefficient A_0 to be non-zero. A major consequence is that any parameterization that varies continuously with the original cubic coefficients requires a split between maps with a positive A_0 or a negative A_0. Consequently, there are two fundamentally different pieces of the parameter space.

Real parameterization. Still following [DGMT], we consider a parameterization of $\mathcal{F}$. Being bimodal, each map f has exactly two critical points, both real and non-degenerate, which we label c_1 and c_2, with $c_1 < c_2$. The corresponding critical values, v_1 and v_2, when considered as a pair (v_1, v_2), uniquely determine a map in $\mathcal{F}$ and its dynamics up to conformal conjugacy. (See [DGMT], Lemma 1) So we consider the pair $(v_1, v_2) \in \mathbb{R} \times \mathbb{R}$ as our parameterization. As noted above, the set of points in $\mathbb{R} \times \mathbb{R}$ for which we have a corresponding map, $f \in \mathcal{F}$, comes naturally in two pieces. One piece corresponds to maps for which $1 \geq v_1 \geq v_2 \geq 0$, where the boundary points of $[0, 1]$ are fixed by f, and the other corresponds to maps for which $0 \leq v_1 \leq v_2 \leq 1$, where the boundary points of $[0, 1]$ form an orbit of period two.

Restricting our attention to only one such piece of the parameter space, we can consider a map, f_λ, with $\lambda = (v_1, v_2)$, to be *continuously varying* with respect to this parameter λ. We refer to such a piece as P.

Complex parameterization. We frequently need to consider the maps in $\mathcal{F}$ as complex maps, and sometimes we will need to consider these maps within a larger family whose corresponding parameter space is an open subset of $\mathbb{C} \times \mathbb{C}$. In particular, we will need such a family to apply the λ-Lemma (See Chapter VII part 3), as it requires that parameters be allowed to take values in an open disk of $\mathbb{C}$. This will require our family to contain cubic maps which are not real. As in $\mathcal{F}$, we will require two pieces to in order to take into account both types of maps. The definition of this extended family is as follows.

Definition. (the extended family $\mathcal{G}$) A map f is in $\mathcal{G}$ if it can be written as either

$$z(z-1)(Az-B)+z \qquad A \neq 0 \text{ and } A, B \in \mathbb{C} \quad \text{or}$$

$$(z(Az-B)+1)(1-z) \qquad A \neq 0 \text{ and } A, B \in \mathbb{C}.$$

Both pieces are parameterized by the pair $(A, B) \in \mathbb{C} \times \mathbb{C}$, and together they contain all of $\mathcal{F}$.

§3.2. Bones.

Let P be one piece of the parameter space representing maps in $\mathcal{F}$, as parameterized by the critical values.

Definition. (bones) The **left bone** $B_-(\hat{o})$ is the set of parameter values for which the left hand critical point is periodic with **order type** $\hat{o}$. By definition, this means that the points of the orbit, numbered as $x_1 < \cdots < x_p$, satisfy $x_i \mapsto x_{\hat{o}(i)}$ where $\hat{o}$ is some given cyclic permutation of $\{1, \ldots, p\}$. The **dual right bone** $B_+(\hat{o})$ is the set of parameter values for which the right critical point is periodic with this same order type.

A **bone** can be either a left bone or a right bone, and we know, in P, that almost all bones, those where the fixed critical point has period three or more, are smooth one-dimensional manifolds with exactly two boundary points (See [DGMT], Lemma 4). We ultimately show that all bones are connected and form simple arcs. A one-dimensional manifold that is disconnected and has exactly two boundary points must contain a connected component which forms a simple closed curve. We call such a component a **bone-loop**.

§3.3. Special quasiconformal conjugacies.

Definition. As in the study of quadratic polynomials, we show (see Chapter V part 1) that given the existence of a topological conjugacy between two real cubic maps, we can continuously deform this conjugacy into a whole family of conjugacies and conjugate maps. This construction assumes that we can show topological conjugacies between maps in $\mathcal{F}$ have certain properties. Here is a list of the required properties.

(1) $H : \hat{\mathbb{C}} \to \hat{\mathbb{C}}$ is a quasiconformal homeomorphism with $H(\infty) = \infty$.

(2) H preserves the real line and its orientation.

(3) H preserves $[0, 1]$ so we must have $H(0) = 0$ and $H(1) = 1$

(4) Suppose H conjugates f_1 and f_2, with $f_1, f_2 \in \mathcal{F}$. Also suppose F is the Fatou set of f_1. Then $H|_F$ is conformal. In other words, $\chi_H(z)$ is zero on F.

Conjugacies which satisfy these properties will be called **special quasiconformal** or **special qc-conjugacies.**

Properties. The next lemma shows a special qc-conjugacy between two given maps, if it exists at all, is unique.

Lemma 3.1. Suppose H is a special qc-conjugacy that conjugates $f \in \mathcal{F}$ with itself. Then H is the identity map on $\hat{\mathbb{C}}$.

PROOF. The first step is to show that H is the identity map on the Julia set of f, J. The two critical points of f have bounded orbits, they are contained in $[0,1]$, so the basin at infinity of f, U, is simply connected. Furthermore H, as a conjugacy between f and f, must map connected components of the Fatou set onto other components, and since $H(\infty) = \infty$, U must be mapped onto itself.

Let $\varphi : U \to D$ be the Riemann map that takes U conformally onto the unit disk, mapping infinity to zero. We specify φ uniquely by picking a positive real point p that we require φ to map onto the positive real line. The composed map $\varphi \circ H \circ \varphi^{-1}$ is a conformal homeomorphism from D onto itself, fixing zero. By the Schwarz Lemma, this map must be a rotation. Furthermore this map does not change the argument of $\varphi(p)$, since φ maps p to the positive real line and H preserves the positive real line. So the rotation, in fact, must be the identity map. The map H then must also be the identity on U. Since H is continuous, we must have $H|_{\partial U}$ is also the identity.

To show that H is the identity everywhere, we can apply a lemma due to Bers.

Lemma 3.2. Let $F \subset \hat{\mathbb{C}}$ be an open set, and let H be a topological automorphism of $\hat{\mathbb{C}}$ such that $H|_F$ is quasiconformal and $H|_{\hat{\mathbb{C}}-F}$ is the identity. Then H is quasiconformal and $\chi_H|_{\hat{\mathbb{C}}-F} = 0$.

For our purposes, we let F be the Fatou set of f, and then the Julia set is $J = \hat{\mathbb{C}} - F$. By the assumption that H is special quasiconformal, we know χ_H restricted to F is zero. Combining this with the second conclusion of the Bers Lemma, we have $\chi_H = 0$ a.e., and thus H is conformal. Since H is the identity on the basin of infinity, H must be the identity everywhere. Q.E.D.

Corollary. Suppose H is a special qc-conjugacy between f_1 and $f_2 \in \mathcal{F}$. Then H is unique.

PROOF. Suppose H_* is another special qc-conjugacy between f_1 and f_2, then $H_*^{-1} \circ H$ is a special qc-conjugacy between f_1 and itself. Thus $H = H_*$. Q.E.D.

Corollary. Suppose H is a special qc-conjugacy between f_1 and $f_2 \in \mathcal{F}$. Then H is symmetric.

PROOF. Define a new map

$$H_* = \bar{H}(\bar{z}).$$

The map H_* preserves the real line and its order, fixes zero and one, and is quasiconformal with the same bound on dilatation as H. Since H is conformal on the Fatou set of f_1, H_* is as well. Therefore, H_* is also a special qc-conjugacy between f_1 and f_2. So by the uniqueness of special qc-conjugacies, we have $H(z) = \bar{H}(\bar{z})$. This is precisely the definition of symmetry. Q.E.D.

§3.4. Statement of results.

For a piecewise monotone map f on the interval, the topological entropy of f can be defined in terms of the number of **laps** $\ell(f^{\circ j})$ of the iterate $f^{\circ j}$, that is the number maximal intervals on which $f^{\circ j}$ is monotone. The **topological entropy** of f is then

$$h(f) = \lim_{j\to\infty} \frac{\log \ell(f^{\circ j})}{j}.$$

We can now state the monotonicity result for cubic maps. Recall that P is a connected component of the parameter space of $\mathcal{F}$.

Theorem 3.3. (Cubic Monotonicity) For each $h_0 \in [0, \log 3]$, the set

$$\{\lambda = (v_1, v_2) \in P \mid h(f_\lambda) = h_0\}$$

is connected.

The principal reduction of the Cubic Monotonicity Theorem is given in [DGMT].

Connected Bone Conjecture. Every bone for the cubic family $\mathcal{F}$ is a simple arc.

The work of Dawson, Galeeva, Milnor, and Tresser shows that proving this conjecture is enough to prove Theorem [3.3]. We show that bone-loops do not exist, and so both the Connected Bone Conjecture and Theorem [3.3] are true.

IV. Quasiconformal pull-back

Let f and $\hat{f}$ be maps in the normalized family $\mathcal{F}$. A standard argument due to Sullivan (see [Su1]) shows that in order to construct a quasiconformal conjugacy between real polynomials it is enough to construct a symmetric quasiconformal homeomorphism that maps the critical orbits of f onto those of $\hat{f}$, i.e. the homeomorphism is equal to the conjugacy on the critical orbits. The basic procedure used in this argument is a construction called quasiconformal *pull-back.*

We provide a complete proof that pull-back works in the specific case of real cubic maps. This amounts to a short-cut in the construction of a special qc-conjugacy between f and $\hat{f}$. An initial homeomorphism H_0 will be constructed which satisfies the functional equation of a conjugacy only on a subset U_0 of the Riemann sphere. Then we attempt to build the topological lift of this map through the dynamical maps f and $\hat{f}$. As we will see, the new map H_1 satisfies the conjugacy equation on a subset of $\hat{\mathbb{C}}$ which is exactly the preimage, under f, of the original U_0 for H_0. We can then repeat this process and construct an entire sequence of maps H_n. The idea is to choose the initial U_0 so that repeated liftings will always be possible and, with each lift, the subset on which the conjugacy is satisfied, U_n, expands to fill all of $\hat{\mathbb{C}}$.

The main obstruction to building H_{n+1} from H_n is that H_n must map the critical values of f onto the corresponding critical values of $\hat{f}$ so that standard lifting theorems on regular covering maps apply. We must temporarily remove the critical values in the image space of f and $\hat{f}$ before they can be considered covering maps, and H_n must *respect these punctures.* Let c_R, c_L, $\hat{c}_R$, and $\hat{c}_L$ be the left and right critical points of f and $\hat{f}$ respectively, and let v_R, v_L, $\hat{v}_R$, and $\hat{v}_L$ be right and left critical values for each map. Define the punctured spaces

$$C_I = \hat{\mathbb{C}} - \{\infty, v_L, v_R\} \qquad \hat{C_I} = \hat{\mathbb{C}} - \{\infty, \hat{v_L}, \hat{v_R}\}$$
$$C_D = \hat{\mathbb{C}} - \{\infty, c_L, c_R, w_L, w_R\} \quad \hat{C_D} = \hat{\mathbb{C}} - \{\infty, \hat{c_L}, \hat{c_R}, \hat{w_L}, \hat{w_R}\}.$$

Assuming H_n respects this structure we will be able to construct H_{n+1}, and in general we will have the following commutative diagram:

$$\begin{array}{ccc} C_D & \xrightarrow{H_{n+1}} & \hat{C_D} \\ \big\downarrow f & & \big\downarrow \hat{f} \\ C_I & \xrightarrow{H_n} & \hat{C_I} \end{array}$$

Clearing the obstruction to lifting for the entire sequence translates to the condition that the initial H_0 must map the entire forward orbit of both critical points of f onto the corresponding orbits of $\hat{f}$. This condition will be propagated by the lifting construction.

Any bound on the complex dilatation of H_n, χ_{H_n}, is passed on to H_{n+1}. As is clear from the commutative diagram, H_{n+1} is the composition of H_n with two conformal maps. Since conformal maps have zero complex dilatation, the upper bound of the dilatation of H_{n+1} is exactly equal to the upper bound of H_n (See [10.1]). Special qc-conjugacies must also be conformal on the Fatou set of f, and quasiconformal on the Julia set. All these requirements can be viewed as local conditions on our maps, which will propagate easily through the construction since we are lifting through conformal maps.

Assuming we can construct this initial map, we will show that the sequence of lifts have a well-defined limiting map, which necessarily must conjugate f and $\hat{f}$. Furthermore, all the additional properties of a special qc-conjugacy will be inherited from the initial H_0, and the new map will be a special qc-conjugacy.

§4.1. Induction hypotheses.

For a fixed f and $\hat{f}$ in $\mathcal{F}$, the map H_n and the set U_n must satisfy these properties during induction. Let F be the Fatou set of f.

(a) H_n is quasiconformal with complex dilatation $|\chi_{H_n}(z)| \leq K < 1$ for all z.

(b) H_n maps $\mathbb{R}$ onto itself and preserves the orientation.

(c) H_n maps the interval $[0, 1]$ onto itself, fixing 0 and 1.

(d) H_n maps v_L to $\hat{v_L}$ and v_R to $\hat{v_R}$.

(e) H_n is conformal for $z \in U_n \cap F$, i.e., $|\chi_{H_n}(z)| = 0$

(f) U_n contains 0, 1, c_L, and c_R.

(g) $f(U_n) \subset U_n$.

(h) For $z \in U_n$, we have $H_n \circ f(z) = \hat{f} \circ H_n(z)$.

(i) H_n maps $f^{-1}(\mathbb{R})$ onto $\hat{f}^{-1}(\mathbb{R})$.

We also have an additional property that is easier to verify than property (i) for the initial induction step.

(j) $U_0 \cup \mathbb{R}$ is path-connected.

For the initial induction step, property (j) will replace (i), and for all subsequent steps only property (i) will be required.

Relationships between the maps and the sets are given by

$$\hat{f} \circ H_{n+1}(z) = H_n \circ f(z) \qquad \text{for all } z \in \hat{\mathbb{C}}, \tag{1}$$

$$U_{n+1} = f^{-1}(U_n), \qquad \text{and} \tag{2}$$

$$H_{n+1}(z) = H_n(z) \qquad \text{for all } z \in U_n. \tag{3}$$

Note: The assumption that H_n maps forward critical orbits onto forward critical orbits is encoded in conditions (d), (f), (g), and (h). These conditions also hide a strong condition on f and $\hat{f}$, which is why there is no explicit requirement on these maps. Their critical orbits must come in the same order, which means they are, in fact, combinatorially equivalent.

§4.2. Induction step.

The induction argument is broken up between three lemmas which we are now ready to state. The bulk of the induction is proved in this first lemma, while the next two lemmas deal with the initial induction step.Let (H_n, U_n) be an associated pair of one homeomorphism and one subset of $\hat{\mathbb{C}}$.

Lemma 4.1. If (H_n, U_n) satisfies properties (a)-(h), then there exists a unique pair (H_{n+1}, U_{n+1}) satisfying (a)-(g) and (i). Moreover, the pairs together satisfy equations (1) and (2).

PROOF. As mentioned earlier, $f : C_D \mapsto C_I$ and $\hat{f} : \hat{C}_D \mapsto \hat{C}_I$ are covering maps as we have removed the only points without proper three to one covers. We still need to make some additional normalizations. Define the affine map G taking v_L to $\hat{v}_L$ and v_R to $\hat{v}_R$ by

$$G(z) = \left(\frac{\hat{v}_L - \hat{v}_R}{v_L - v_R}\right)(z - v_L) + \hat{v}_L.$$

Notice that G is also well-defined on the space C_I. Now consider lifting the map $H_n \circ G^{-1}$, through the maps $G \circ f$ and $\hat{f}$. If we label the lift H_{n+1}, it will satisfy the functional equation

$$H_n \circ G^{-1} \circ G \circ f = \hat{f} \circ H_{n+1},$$

which is the same as the desired equation (1).

To define this lift, we construct a homotopy from $H_n \circ G^{-1}$ to the identity map, and first lift the identity map. The homotopy $I_t(z) : \hat{C}_I \to \hat{C}_I$ is defined as

$$I_t(z) = I(z,t) = z(1-t) + t(H_n \circ G^{-1}(z)).$$

Clearly $I_0(z) = z$ and $I_1(z) = H_n \circ G^{-1}(z)$, but we need to check that I_t respects the structure of $\hat{C}_I$. It is quick to verify that $I_t(\hat{v}_L) = \hat{v}_L$ and $I_t(\hat{v}_R) = \hat{v}_R$ for all t, thus I_t maps punctures to punctures. The only thing left that can go wrong is if there exists some other point $y \neq \{\hat{v}_L, \hat{v}_R\}$ and t with $I_t(y) = \hat{v}_L$ or $\hat{v}_R$. By property (b), H_n preserves the real line, and

G clearly does, so I_t must as well. Therefore y must be real. Without loss of generality, suppose $I_t(y) = \hat{v}_L$. If $t = 0$, then we have $y = \hat{v}_L$, contrary to the definition of y. If $t = 1$, then $H_n \circ G^{-1}(y) = \hat{v}_L$, which implies that $G^{-1}(y) = v_L$ and $y = \hat{v}_L$. This is again a contradiction, so we must have $0 < t < 1$. Since I deforms y linearly into $H_n \circ G^{-1}(y)$, these two points must be on either side of $\hat{v}_L$ on the real line. So $H_n \circ G^{-1}$ must reverse the ordering of y and $\hat{v}_L$. In other words, if $y < \hat{v}_L$ then $H_n(y) > H_n(\hat{v}_L)$, but this violates property (b). So $I_t : \hat{C}_I \to \hat{C}_I$ is well-defined.

The lifting of the identity map is done algebraically. Define

$$G'(z) = \left(\frac{\hat{c}_L - \hat{c}_R}{c_L - c_R}\right)(z - c_L) + \hat{c}_L.$$

The map G' is an affine map taking c_R to $\hat{c}_R$ and c_L to $\hat{c}_L$. Now we examine the map $\hat{f} \circ G' - G \circ f$. It is a cubic map with critical points (zero derivatives) at c_L and c_R, and the corresponding critical values are both zero. The whole map must be identically zero, and we have the equation

$$\hat{f} \circ G' = G \circ f.$$

We consider G' as a lift of the identity map through $\hat{f}$ and $G \circ f$. We can use G' to define $H_{n+1}(z)$.

Let z be any point in C_D. We define a path $\alpha : [0,1] \to \hat{C}_I$ as

$$\alpha(t) = I(G \circ f(z), t).$$

Since G' is a lift, $G'(z)$ is in $\hat{f}^{-1}(G \circ f(z))$. By the standard path lifting theorem, α has a unique lift with endpoint $G'(z)$. Let $H_{n+1}(z)$ be the other end of this path. By construction, H_{n+1} satisfies the functional equation (1) for $z \in C_D$. Continuity follows from the continuity of our choice of endpoints, $G \circ f(z)$ and $H_n \circ f(z)$, and the continuity of the path-lifting construction.

We extend H_{n+1} from C_D to $\hat{\mathbb{C}}$ by continuity as well. We can consider locally $H_{n+1}(z) = \hat{f}^{-1} \circ H_n \circ f(z)$. Continuity insists that $\{c_L, c_R, w_L, w_R\}$ is mapped onto $\{\hat{c}_L, \hat{c}_R, \hat{w}_L, \hat{w}_R\}$ in some fashion. At any rate, H_{n+1} is well-defined on $\hat{\mathbb{C}}$ and satisfies equation (1).

Since the lift of a homeomorphism is also a homeomorphism, H_{n+1} is injective on C_D, and given that $\{c_L, c_R, w_L, w_R\}$ are isolated points, H_{n+1} is a homeomorphism of $\hat{\mathbb{C}}$.

Let us examine how H_{n+1} behaves on the real line. Let r be an element of $\mathbb{R}$. Since G, f, and I_t all preserve $\mathbb{R}$, the path we lift, $\alpha(t) = I_t(G \circ f(r))$ is entirely contained in $\mathbb{R}$. Since $\alpha([0,1])$ is a connected set in $\hat{C}_I$, the set is either entirely contained in the interval $(\hat{v}_L, \hat{v}_R)$ or does not intersect it at all. In the first case, all points in $\alpha([0,1])$ have three real preimages under $\hat{f}$, and in the second case, exactly one preimage of each point in the path is real. Thus the set $\hat{f}^{-1}(\alpha([0,1]))$ consists of three paths, at least one of which is entirely contained in $\mathbb{R}$, the other two paths are either both contained in $\mathbb{R}$ or do not intersect it at all. But the path lifting specified by the construction of H_{n+1} has $G'(r)$ as an endpoint, a real number. Thus $H_{n+1}(r)$ must be real.

So H_{n+1} maps $\mathbb{R}$ onto itself. The map G' preserves the orientation of the real line, so H_{n+1} inherits this from G'. To see this, choose z smaller than all points $\{c_L, c_R, w_L, w_R\}$, then $G'(z)$ will be smaller than the corresponding set

$\{\hat{c}_L, \hat{c}_R, \hat{w}_L, \hat{w}_R\}$. In the construction of $H_{n+1}(z)$, we get a path lifting entirely contained in $\mathbb{R}$, which must be entirely to the left of the same points, since it is continuously connected to $G'(z)$. Thus $H_{n+1}(z)$ must be less than $H_{n+1}(c_L) \in \{\hat{c}_L, \hat{c}_R, \hat{w}_L, \hat{w}_R\}$, and H_{n+1} preserves the ordering of $z < c_L$. Consequently the order of $\{c_L, c_R, w_L, w_R\}$ is preserved under H_{n+1}, and property (b) is satisfied. We also obtain

$$H_{n+1}(c_L) = \hat{c}_L \qquad H_{n+1}(c_R) = \hat{c}_R$$
$$H_{n+1}(w_L) = \hat{w}_L \quad H_{n+1}(w_R) = \hat{w}_R.$$

Equation (1) is satisfied by construction, and we use equation (2) as the definition of U_{n+1}. From equation (1) and property (c) and (f) for H_n, we get

$$H_{n+1}(0) = \hat{f}^{-1} \circ H_n \circ f(0) = \hat{f}^{-1} \circ \hat{f}(0).$$

The only real point in the set $\hat{f}^{-1} \circ \hat{f}(0)$ is 0, since neither 0 nor 1 is in $(\hat{v}_L, \hat{v}_R)$, except in the exceptional case where $\hat{f}(0)$ is a critical value itself. In this case, the only real points in the set $\hat{f}^{-1} \circ \hat{f}(0)$ are 0 and a critical point. But we already know, $H_{n+1}(0)$ cannot be critical. In any case we must have $H_{n+1}(0) = 0$. A similar proof shows $H_{n+1}(1) = 1$. This verifies property (c).

As we have already mentioned, $|\chi_{H_{n+1}}(z)|$ and $|\chi_{H_n}(z)|$ share the same bound K, so property (a) is satisfied for H_{n+1}. If z is in $U_{n+1} \cap F$, then $f(z)$

must be in $U_n \cap F$. Locally we know $H_{n+1} = \hat{f}_{-1} \circ H_n \circ f$, so on $U_{n+1} \cap F$, H_{n+1} is the lift of a conformal map and is thus conformal itself. This proves property (e).

To show property (g), let z be an element of U_{n+1}. By definition, $f(z) \in U_n$. By property (g) for H_n, $f \circ f(z)$ is in U_n, therefore $f(z)$ is in U_{n+1}. As a tiny corollary to this, we know $U_n \subset U_{n+1}$, which immediately gives us property (f).

Property (d) is split into cases depending on how $\{c_L, c_R\}$ is mapped into $\{v_L, v_R\}$ by f. Suppose without loss of generality that $f(c_L) = v_R$. We know already that c_L is in U_{n+1}, so we must have $H_{n+1}(v_R) = \hat{f} \circ H_{n+1}(c_L)$. But as is shown above $H_{n+1}(c_L) = \hat{c}_L$. So we get $H_{n+1}(v_R) = \hat{v}_R$ as required.

Let z be in the set $f^{-1}(\mathbb{R})$, so $f(z) \in \mathbb{R}$. We know from equation (1), $\hat{f} \circ H_{n+1}(z) = H_n(f(z))$. By property (b) for H_n, $\hat{f} \circ H_{n+1}(z)$ must be real, so $H_{n+1}(z)$ itself must be in $\hat{f}^{-1}(\mathbb{R})$. This finishes the final property (i). Q.E.D.

We are ready to prove a full lifting lemma now. By verifying equation (3), we show that the sequence of maps stabilizes for points in the Fatou set. This will be a crucial step in showing that the sequence converges.

Lemma 4.2. Suppose (H_n, U_n) satisfies properties (a)-(i), then there exists a unique pair (H_{n+1}, U_{n+1}) satisfying (a)-(i) and together with (H_n, U_n) satisfying equations (1)-(3).

Proof. Because of Lemma [4.1], the existence of (H_{n+1}, U_{n+1}) and most of the properties have been verified. We need only check property (h) and equation (3). We consider the set $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$ and the corresponding set $\hat{\mathbb{C}} - \hat{f}^{-1}(\mathbb{R})$. Evidently $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$ consists of several pieces. We start by showing that f maps a connected component univalently onto one of the standard half-planes.

Let V be a connected component of $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$. Suppose f is not univalent, and V contains two different points, z and w, with $f(z) = f(w)$. Since they are in a connected set there exists a path in V connecting z and w. Such a path obviously does not intersect $f^{-1}(\mathbb{R})$, but we claim its image under f intersects $\mathbb{R}$, a contradiction. The image of the path is a closed loop by assumption. The region bounded by this loop must contain a critical value of f, a real point. For the loop to bound a real point it must intersect $\mathbb{R}$ itself. This contradiction shows $f|_V$ is univalent.

Since V contains no point in $f^{-1}(\mathbb{R})$, $f(V)$ contains no point in $\mathbb{R}$, and since $f^{-1}(\mathbb{R})$ contains the boundary of V, the boundary of $f(V)$ must be contained in $\mathbb{R}$. Thus $f(V)$ must be either $\mathbb{H}^+$ or $\mathbb{H}^-$, an entire half-plane, with boundary the entire real line. We can conclude there must be six components to $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$, three preimages of each half-plane.

Both half-planes have the interval (v_L, v_R) as part of their boundary, so all components of $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$ have a preimage of (v_L, v_R), either (w_L, c_L), (c_L, c_R), or (c_R, w_R), as part of their boundary. We can characterize each piece of $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$ uniquely, then, by which half-plane it is in and which preimage of (v_L, v_R) is contained in its boundary.

We know H_{n+1} satisfies property (i). Thus H_{n+1} must map the pieces of $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$ onto the corresponding pieces of $\hat{\mathbb{C}} - \hat{f}^{-1}(\mathbb{R})$. We know H_{n+1}, by properties (a) and (b), maps the half-planes $\mathbb{H}^+$ and $\mathbb{H}^-$ onto themselves. Furthermore (w_L, c_L), (c_L, c_R), and (c_R, w_R) map onto $(\hat{w}_L, \hat{c}_L)$, $(\hat{c}_L, \hat{c}_R)$, and $(\hat{c}_R, \hat{w}_R)$ respectively. Thus the pieces of $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$ are mapped in a canonical way onto the pieces of $\hat{\mathbb{C}} - \hat{f}^{-1}(\mathbb{R})$. The map H_n must map the pieces in the same canonical way.

Let z be a point in U_n. From equation (1) and property (h) for H_n we have

$$H_{n+1}(z) \quad = \quad \hat{f}^{-1} \circ H_n \circ f(z) \quad = \quad \hat{f}^{-1} \circ \hat{f} \circ H_n(z).$$

Suppose $z \in f^{-1}(\mathbb{R})$, then $\hat{f} \circ H_n(z)$ is real by property (i). We divide this case into three sub-cases: z is in $\mathbb{R}$, $\mathbb{H}^+$, or $\mathbb{H}^-$. Suppose z is in the upper half-plane, then $H_n(z)$ is also in the upper half-plane and $\hat{f} \circ H_n(z)$ is real as noted. Being real, $\hat{f} \circ H_n(z)$ can have at most one preimage under $\hat{f}^{-1}$ in the upper half-plane, $H_n(z)$. Since H_{n+1} must also map z into the upper half-plane, we must have $H_{n+1}(z) = H_n(z)$ as required by equation (3). The other two sub-cases have similar proofs.

Now suppose our z is in $\hat{\mathbb{C}} - f^{-1}(\mathbb{R})$, then $H_n(z)$ and $H_{n+1}(z)$ must be contained in the same piece of $\hat{\mathbb{C}} - \hat{f}^{-1}(\mathbb{R})$ as noted above. But again from equation (1) and property (h) we have $\hat{f} \circ H_{n+1}(z) = \hat{f} \circ H_n(z)$. $H_{n+1}(z)$ and $H_n(z)$ have the same image, but $\hat{f}$ maps each piece of $\hat{\mathbb{C}} - \hat{f}^{-1}(\mathbb{R})$ univalently. Again we must have $H_{n+1}(z) = H_n(z)$.

This proves equation (3). We get property (h) for H_{n+1} immediately from this and equation (1). Suppose $z \in U_{n+1} = f^{-1}(U_n)$, then $f(z) \in U_n$. We get

$$H_{n+1} \circ f(z) \quad = \quad H_n \circ f(z) \quad = \quad \hat{f} \circ H_{n+1}(z).$$

Q.E.D.

As mentioned before, in constructing the initial pair (H_0, U_0), it is easier to start with property (j) than with property (i). The next lemma justifies this change of requirements.

Lemma 4.3. Suppose (H_n, U_n) satisfies properties (a)-(h) and (j). Then there exists a unique pair (H_{n+1}, U_{n+1}) satisfying properties (a)-(i) and, together with (H_n, U_n), satisfying equations (1)-(3).

PROOF. Again existence of (H_{n+1}, U_{n+1}) and most of the properties are given by Lemma [4.1]. We need to verify property (h) and equation (3). Again we start with equation (3). Let z be a point in U_n and assume $z \notin \mathbb{R}$. By property (j) we can find a path, $\alpha : [0,1] \to U_n \cup \mathbb{R}$, connecting $z = \alpha(0)$ with a real point $\alpha(1) \in \mathbb{R}$. We normalize things by assuming that $\alpha(1)$ is real and $\alpha(t) \notin \mathbb{R}$ for $t \neq 1$. If $z \in \mathbb{R}$, we define a constant path $\alpha(t) = z$. In either case, we have a path connecting z and $\mathbb{R}$ with $\alpha(t) \in U_n$ for $t \neq 1$.

From equation (1) and property (h) for H_n, we have for $t \neq 1$

$$H_{n+1}(\alpha(t)) = \hat{f}^{-1} \circ H_n \circ f(\alpha(t)) = \hat{f}^{-1} \circ \hat{f} \circ H_n(\alpha(t)).$$

By continuity, this equation holds for $t = 1$ as well; however, $\alpha(1)$ is real, so H_{n+1} must map it onto the real line. Suppose $\alpha(1)$ is in the interval (c_L, c_R), then $\hat{f} \circ H_n(\alpha(1))$ is in $(\hat{v}_L, \hat{v}_R)$, and the possibilities for $H_{n+1}(\alpha(1))$, the three points in the inverse image under $\hat{f}$ of $\hat{f} \circ H_n(\alpha(1))$, include exactly one point in $(\hat{c}_L, \hat{c}_R)$, $H_n(\alpha(1))$. Other cases with $\alpha(1)$ not in (c_L, c_R) similarly prove that $H_{n+1}(\alpha(1)) = H_n(\alpha(1))$.

Consider $H_{n+1}(\alpha(1))$ as a base point as we construct H_{n+1} as a lift of H_n. Then $H_{n+1}(\alpha(0))$ would be defined as the other endpoint of the lift of $H_n \circ f(\alpha)$ with endpoint $H_{n+1}(\alpha(1))$. But clearly $H_n(\alpha)$ is this lift because $\hat{f} \circ H_n(\alpha) = H_n \circ f(\alpha)$ by property (h), and this path has endpoint $H_n(\alpha(1)) = H_{n+1}(\alpha(1))$. Therefore $H_{n+1}(z) = H_n(\alpha(0)) = H_n(z)$, proving equation (3). Property (h) is proved as in Lemma [4.2] from equations (1) and (3). Q.E.D.

§4.3. Limit of the induction.

The induction process is now completely defined, and we now show that the limit of the H_n is a special qc-conjugacy.

Theorem 4.4. Let f and $\hat{f}$ in $\mathcal{F}$ be maps with no indifferent periodic cycles, and suppose we can find a map H_0 and set U_0 satisfying properties (a)-(h) and (j). Additionally suppose U_0 contains a neighborhood of every attracting periodic point of f in $\hat{\mathbb{C}}$, including infinity. Then there exists a special qc-conjugacy between f and $\hat{f}$.

PROOF. By Lemma [4.3], we can construct the first lift (H_1, U_1) in a sequence. The rest of the sequence, (H_n, U_n) is constructed using Lemma [4.2]. All pairs share at least properties (a)-(h), and any two successive pairs satisfy equations (1)-(3). By equation (2) and property (g) we have a chain of sets $U_0 \subset U_1 \subset \cdots \subset U_n \subset \cdots$. This chain eventually contains all of the Fatou set F. For if $z \in F$, then for some n, $f^{\circ n}(z) \in U_0$, since U_0 contains a neighborhood of all attracting periodic points. Therefore z is in U_n. For any integer $i > n$, $H_i(z) = H_{i-1}(z) = \cdots = H_n(z)$, by applying equation (3) repeatedly. Thus for any point z of the Fatou set, the sequence $\{H_i(z)\}_{i=1}^{\infty}$ not only converges but becomes fixed.

For $z \in F$, the sequence $\{H_n(z)\}_{n=0}^{\infty}$ converges, but we need to show it converges everywhere and to a continuous function. The Fatou set of a polynomial is dense in $\hat{\mathbb{C}}$, so the sequence $\{H_n\}$ converges pointwise on a dense set. This is not quite enough, we still need some uniform convergence conditions. By property (a), all our maps are quasiconformal with the same bound on the dilatation and fix the points 0, 1, and ∞, so these maps form a normal family.

Let $\{H_{n_k}\}_{k=0}^{\infty}$ be some subsequence that converges uniformly on all of $\hat{\mathbb{C}}$ to a continuous function H. We want to show that the full sequence has this same limit. Let z be a point in the Julia set $J = \hat{\mathbb{C}} - F$. Since the Fatou set F is dense, we can find a sequence $\{z_i\}_{i=0}^{\infty}$, contained in F, converging to our point z. The sequence of functions $\{H_n\}_{n=0}^{\infty}$ becomes fixed for our points $\{z_i\}$. We define integers N_i so that if $k \geq N_i$, then $H_k(z_i)$ equals the limit of the sequence $\{H_n(z_i)\}_{n=0}^{\infty}$, which we know exists and label ω_i. Since ω_i is a limit for the whole sequence, we know $H(z_i) = \omega_i$, and since H is continuous, $\lim_{i\to\infty} H(z_i) = H(z)$. We need to show that $\lim_{n\to\infty} H_n(z) = H(z)$.

Let $\epsilon > 0$ be given. By Arzela's Theorem, functions in a normal family are equicontinuous, so we can find a $\delta > 0$ so that if $|z_i - z| < \delta$, then

$$|H_n(z_i) - H_n(z)| < \epsilon/2 \qquad \text{for all } n \geq 0.$$

Choose the integer i_0 so that two things happen. First let $|z_i - z|$ actually be smaller than δ so that the above equation is true for $i \geq i_0$, and second let

$$|H(z_{i_0}) - H(z)| \;=\; |\omega_{i_0} - H(z)| \;<\; \epsilon/2$$

be true as well. Now finally if $n > N_{i_0}$, then we have

$$\begin{aligned} |H_n(z) - H(z)| &\leq |H_n(z) - \omega_{i_0}| + |\omega_{i_0} - H(z)| \\ &= |H_n(z) - H_n(z_{i_0})| + |\omega_{i_0} - H(z)|. \end{aligned}$$

By the previous inequalities we then get

$$|H_n(z) - H(z)| < \epsilon \qquad \text{for } n > N_{i_0}.$$

This proves $\{H_n(z)\}_{n=0}^{\infty}$ converges to $H(z)$.

The map H, being the limit of quasiconformal maps with a single bound on dilatation, is quasiconformal with the same bound. Suppose z is in the Fatou set. As has been noted before, the sequence $\{H_n(z)\}_{n=0}^{\infty}$ becomes fixed when $f^{\circ n}(z) \in U_0$. The part of F contained in U_0 is, by construction, a neighborhood of attracting periodic points. Thus there is a neighborhood of z in F, which becomes fixed by the sequence at the same time. On this neighborhood, $H_n = H$, but H_n, for n greater than the fixing value, by property (e) is conformal on this neighborhood. The uniform limit of conformal functions is conformal, so H is conformal on the Fatou set. Clearly since all the $\{H_n\}$ fix 0, 1, and ∞, H fixes these points as well. For

$z \in \mathbb{R}$, $\{H_n(z)\}_{n=0}^{\infty}$ are all real numbers, so H inherits property (b) from the sequence also. Finally taking the limit of both sides of equation (1), we get

$$H \circ f(z) = \lim_{n\to\infty} H_n \circ f(z) = \lim_{n\to\infty} \hat{f} \circ H_{n+1}(z) = \hat{f} \circ H(z).$$

Thus $H(z)$ is a special qc-conjugacy between f and $\hat{f}$. Q.E.D.

§4.4. Conjugacies between hyperbolic maps.

We continue to ease the conditions for which special qc-conjugacies exist. The requirement that the conjugacy be conformal on the Fatou set is very restrictive. The conditions under which a conformal conjugacy exists between attracting basins of f and $\hat{f}$ are related to the multiplier of an attracting orbit. For instance, a conformal conjugacy can only be defined locally around a periodic point of one map if the multiplier of that point is equal to the multiplier of the corresponding periodic point of the other map. We generalize this idea to construct a measure of when conjugacies can be defined on entire basin, and then use this result to show the only special qc-conjugacy between hyperbolic maps is the identity map.

Characterizing attractive basins. Recall that special qc-conjugacies are symmetric by the second corollary to Lemma [3.1] and conformal on the Fatou set. Concentrating on a connected component of the Fatou set, notice that a special qc-conjugacy is nearly determined by these properties because of the Riemann Mapping Theorem. This specifies how H_0 should be defined on Fatou components to satisfy Theorem [4.4]. We will initially construct H_0 only on Fatou components containing at least one point in either critical orbit of the map f. We will call these components **post-critical**. We now start with a preliminary lemma that describes post-critical components. Note that post-critical components must intersect the real line.

Lemma 4.5. Let B be a connected component of the Fatou set of $f \in \mathcal{F}$, and let $\hat{B}$ be any simply connected open set, symmetric with respect to $\mathbb{R}$. Suppose both B and $\hat{B}$ contain a specific real marked point, p and $\hat{p}$ respectively. Then B is symmetric, and there exists a unique conformal map which is symmetric, preserves the orientation of the real line, maps B onto $\hat{B}$, and maps p and $\hat{p}$.

PROOF. First note that f is a symmetric map, so if z escapes to infinity under f, its conjugate $\bar{z}$ escapes as well. Thus the basin at infinity of f is symmetric with respect to $\mathbb{R}$, and therefore the Julia set and Fatou set of either map are symmetric with respect to the real line. The connected component B, which must intersect the real line, is symmetric as well. We define the Riemann map from B to $\hat{B}$ taking p to $\hat{p}$, making it unique be requiring the derivative at p to be a positive real number. Since the mapping is unique and the sets are symmetric, the Riemann map must itself be symmetric and preserve the orientation of $\mathbb{R}$. Q.E.D.

We can now almost completely determine how H_0 must be defined on the post-critical Fatou components of f. Once we establish that B and $\hat{B}$ correspond and pick two marked points, p and $\hat{p}$, we can immediately apply

Lemma [4.5]. Correspondence will come from the combinatorial equivalence of f and $\hat{f}$, but we assume we have correspondence for now and concentrate on finding a marked point for each component, which we will call the *center* of that component. The center will almost always be the unique periodic or preperiodic point in the component, but we need to take into account a number of special cases.

For the following definitions, we restrict our attention to hyperbolic maps. Let B be any post-critical component of the Fatou set of some map $f \in \mathcal{F}$. By the Fatou-Julia wandering domain theorem of Sullivan, the component B must map into some periodic cycle of components each of which contain a unique periodic point on the real line. We define the unique *primary point.*

Definition. (primary point) If B itself contains a real periodic point, then this is the primary point. The other possibility is that B contains only preimages of this periodic cycle, and in this case, B cannot be a periodic component itself. Since there are only two critical points and one must be in the periodic components, at most one critical point can come between B and the periodic components. Therefore B must contain either one or two preperiodic points. If there is only one, this is the primary point, which necessarily must be real. If there are two preperiodic points, either they are both real, or they occur as a symmetric pair on either side of $\mathbb{R}$. In this case the primary point is either the point farthest right on the real line or the point in the upper half-plane.

It is not quite enough to mark only the periodic points, H_0 must make the critical orbits correspond, so we must also mark the critical point, if there is one, in each component. There may be zero, one or two critical points in any given post-critical component.

We refer to the primary point of B as α. A critical point in B will be labeled c, but we will not try to distinguish between two critical points in the same component. We are now ready to define the center of B.

Definition. (the center of B) If the primary point of B is real, it defines the center of B. In the special case where α is not real, B must contain a single critical point, which then defines the center of B.

Definition. (normalizing maps) By Lemma [4.5], for each B there exists a unique Riemann map, symmetric and preserving the orientation of the real line, ψ_B, which takes the center of B to zero and maps B onto the unit disk. We refer to the image of B under ψ_B as a **normalized component**.

The definition of H_0 considered as a map from B to $\hat{B}$ is clear. $H_0|_B$ is the unique map taking the center of B to the center of $\hat{B}$. Unfortunately, this construction makes no guarantee that critical points will correspond under H_0, which is crucial for Theorem [4.4]. We will need an additional hypothesis on B and $\hat{B}$, namely that any critical points occur in the same place with respect to the normalized components.

Definition. (critical factors) For each critical point c in B, its corresponding **critical factor** is

$$\psi_B(c) - \psi_B(\alpha),$$

where α is the primary point of B.

This can be considered the distance between the primary point α and c in the normalized component.

Starting the induction. When f and $\hat{f}$ are hyperbolic, in several cases we can now state when the pull-back induction can be started.

Theorem 4.6. Assume f and $\hat{f}$ are combinatorially equivalent and are both hyperbolic. Suppose the corresponding leftmost and corresponding rightmost critical points, i.e. c_L and $\hat{c}_L$, and c_R or $\hat{c}_R$, have equal critical factors. Then there exists a special qc-conjugacy between f and $\hat{f}$.

PROOF. Since f and $\hat{f}$ are hyperbolic and combinatorially equivalent, they are topologically conjugate on the real line. There are a finite number of post-critical Fatou components for both f and $\hat{f}$, by the Fatou-Julia wandering domain theorem. The conjugacy establishes a correspondence between the components of f and those of $\hat{f}$.

Define H_0 as follows. On each of the post-critical Fatou components, B_i, we pair it with its corresponding component, $\hat{B}_i$. Then $H_0|_{B_i}$ is the symmetric Riemann map preserving the order of the real line, taking the center of B_i to the center of $\hat{B}_i$, and mapping B_i onto $\hat{B}_i$. This map exists and is unique by Lemma [4.5]. We define H_0 between the components containing infinity in the same way. Note that infinity is the primary point, critical point, and center of this component. Because of the combinatorial equivalence between f and $\hat{f}$, as far as we have defined H_0, it preserves the order of the real line. H_0 is also symmetric and quasiconformal (actually conformal).

We need to extend this definition to the entire complex sphere and preserve these three properties. There is no obstruction to a quasiconformal extension of H_0 to the upper half-plane because, thus far, H_0 has been defined to be conformal on a finite number of connected domains. We then define H_0 on the entire sphere by reflection.

Let U_0 be the union of the post-critical Fatou components, the component containing infinity, and the two points 0 and 1.

We verify the hypotheses of Theorem [4.4] one by one. The map H_0 is conformal on certain regions and quasiconformal everywhere else, so H_0 satisfies property (a). Property (b), symmetry, is satisfied by the definition of H_0 and Lemma [4.5]. Properties (c), (e), and (f) are satisfied by construction. It is also easy to see that f maps U_0 into itself. Under f, the union of the post-critical components is forward invariant, the component containing infinity is fixed, and $\{0, 1\}$ is also fixed, so property (g) is satisfied. Every attracting basin contains a critical point, so U_0, consisting of post-critical components, must contain a neighborhood of each attracting periodic point.

We still must verify the functional equation, property (h), and show that critical values are mapped properly, property (d). We verify these two properties on each Fatou component in U_0. Let B_1 and $\hat{B}_1$ be corresponding components and let B_2 and $\hat{B}_2$ be the corresponding image components. The two maps $\psi_{\hat{B}_1} \circ H_0 \circ \psi_{B_1}^{-1}$ and $\psi_{\hat{B}_2} \circ H_0 \circ \psi_{B_2}^{-1}$ are both conformal homeomorphisms that map the unit disk onto itself. They both fix zero, preserve the order of the real line, and are symmetric. By Lemma [4.5], both maps must be the identity map, so

$$H_0|_{B_1} = \psi_{\hat{B}_1}^{-1} \circ \psi_{B_1} \quad \text{and} \quad H_0|_{B_2} = \psi_{\hat{B}_2}^{-1} \circ \psi_{B_2}. \tag{1}$$

Now look at the normalized dynamic maps, $\psi_{B_2} \circ f \circ \psi_{B_1}^{-1}$ and $\psi_{\hat{B}_2} \circ \hat{f} \circ \psi_{\hat{B}_1}^{-1}$. They also map the unit disk onto itself. It is well known that such maps can be represented by Blaschke products, with factors of the form $\frac{z-a}{1-\bar{a}z}$ and an additional rotation factor. Both maps are also symmetric, so we know the rotation factors are plus or minus one and equal to each other, depending on the orientation of f on the right real boundary point of B_1, which is the same as the orientation of $\hat{f}$ on the right boundary of $\hat{B}_1$. We break into cases based on the number of critical points in B_1 (and $\hat{B}_1$).

If there are no critical points, the primary points of B_1 and B_2 are also the centers, so the normalized dynamic map is one-to-one and fixes zero. The single Blaschke factor must be z. The same holds for the corresponding normalized dynamic map between $\hat{B}_1$ and $\hat{B}_2$.

If there are two critical points, the component must be periodic, and again, the primary points of B_1 and B_2 form the centers. The normalized dynamic map must be three-to-one and maps zero to zero, so z is one of the three Blaschke factors. The two critical factors of B_1 are exactly enough to determine the remaining two factors. To be precise, if λ_1 and λ_2 are the two critical factors, then they are the critical points of $\psi_{B_2} \circ f \circ \psi_{B_1}^{-1}$. This is because $\lambda_1 = \psi_{B_1}(\beta_1)$, and similarly $\lambda_2 = \psi_{B_1}(\beta_2)$, where β_1 and β_2 are the critical points of B_1. The other normalized map, $\psi_{\hat{B}_2} \circ \hat{f} \circ \psi_{\hat{B}_1}^{-1}$, must have precisely the same form since the critical factors are equal by hypothesis.

If we have only one critical point, we need to split into two subcases based on whether or not the primary point of B_1 is the center or not. If it is the center, then we have a similar situation to the two critical point case. Each normalized dynamic map fixes zero and has a critical point at λ, the critical factor associated with either corresponding critical point. The Blaschke product must have the form

$$\frac{z-a}{1-az}z \qquad \text{where } a = \tfrac{2\lambda}{1+\lambda^2}.$$

If the primary point is not the center, then the critical point is the center. Both normalized dynamic maps must have a critical point at zero. As noted in the definition of primary points, B_1 must have two symmetric precritical points, α and $\bar{\alpha}$. These must both map to the primary point of B_2 under f and then to 0 under ψ_{B_2}. The Blaschke product of $\psi_{B_2} \circ f \circ \psi_{B_1}^{-1}$ must then have the form

$$\frac{(z-\psi_{B_1}(\alpha))}{(1-\bar{\psi}_{B_1}(\alpha)z)}\frac{(z-\psi_{B_1}(\bar{\alpha}))}{(1-\bar{\psi}_{B_1}(\bar{\alpha})z)}.$$

Note that $\psi_{B_1}(\alpha)$ is purely imaginary. So the Blaschke product has the simpler form

$$\frac{z^2+a^2}{1+a^2z^2}, \qquad \text{where } \psi_{B_1}(\alpha) = ai.$$

This is more clearly symmetric. The other normalized map, $\psi_{\hat{B}_2} \circ \hat{f} \circ \psi_{\hat{B}_1}^{-1}$, must have the same form since the critical factors associated with the critical points in B_1 and $\hat{B}_1$ respectively are equal. This shows that the image of the primary points of B_1 and $\hat{B}_1$ under ψ_{B_1} and $\psi_{\hat{B}_1}$ respectively are equal.

At any rate, in all three cases $\psi_{B_2} \circ f \circ \psi_{B_1}^{-1} = \psi_{\hat{B}_2} \circ \hat{f} \circ \psi_{\hat{B}_1}^{-1}$. Combining this with the identities in (1), we immediately get the functional equation $H_0 \circ f(z) = \hat{f} \circ H_0(z)$ for $z \in B_1$. As a side effect, we also get that any critical points in B_1 are mapped to critical points in $\hat{B}_1$. The same arguments work for all the post-critical components and the component containing infinity. The functional equation also holds trivially for the two points 0 and 1. This concludes the verification of property (h).

Since the functional equation holds on neighborhoods of both critical orbits of f, we can conclude immediately that since the critical points map to critical points under H_0, the critical values, v_L and v_R, must map to the critical values of $\hat{f}$. Since H_0 preserves the order of $\mathbb{R}$, v_R must map to $\hat{v}_R$, and v_L must map to $\hat{v}_L$, verifying property (d). All the hypotheses of Theorem [4.4] have been satisfied, so we may invoke it and construct a special qc-conjugacy between f and $\hat{f}$. Q.E.D.

Corollary. Suppose f and $\hat{f}$ in $\mathcal{F}$ satisfy the hypotheses of Theorem [4.6]. Then $f = \hat{f}$.

PROOF. By Theorem [4.6], there exists a quasiconformal conjugacy, H, between f and $\hat{f}$. Furthermore, H is conformal on the Fatou set of f. But f is hyperbolic, so the Julia set has measure zero in $\hat{\mathbb{C}}$. Therefore H is in fact 1-quasiconformal, so H is a conformal map. By definition of $\mathcal{F}$, however, the only conformal conjugacy between maps in $\mathcal{F}$ is the identity map. Q.E.D.

For conjugacies between cubic maps with critical points in the Julia set, we cannot immediately construct a special qc-conjugacy. However, we have reduced the requirements for when a special qc-conjugacy exists. In particular for maps on a bone-loop, we may (see proposition [5.3]) have a critical point in the Julia set.

Theorem 4.7. Let f and $\hat{f}$ in $\mathcal{F}$ be combinatorially equivalent. Suppose each map has one periodic critical point and one chaotic critical point, in the Julia set. Suppose there exists a quasiconformal homeomorphism H that preserves the real line and its orientation and that H satisfies the functional equation

$$H \circ f(z) = \hat{f} \circ H(z),$$

for z in the forward orbit of the chaotic critical point of f. Then there exists a special qc-conjugacy between f and $\hat{f}$.

Proof. We construct a new quasiconformal map making sure the proper set of points correspond under the new map. The chaotic critical orbit of f must correspond with the chaotic critical orbit of $\hat{f}$. A neighborhood of each point in the periodic critical orbit of f must correspond with a neighborhood of each point in the periodic critical orbit of $\hat{f}$. Also 0 and 1 must be fixed. Our hypotheses have taken care of the most difficult task, making the chaotic critical orbits correspond.

Without loss of generality, we assume that the periodic critical points are actually a fixed points of f and $\hat{f}$. Let c_* and $\hat{c}_*$ be the fixed critical point of f and $\hat{f}$ respectively.

We claim there exists an open neighborhood V of c_* so that $H(V)$ contains $\hat{c}_*$ and has no point in the chaotic critical orbit in the closure of

$H(V)$. Furthermore, V can be choosen to be a topological disk symmetric with respect to the real line. Such a set must contain the real interval bounded by c_* and $H^{-1}(\hat{c}_*)$. If this set does not exist, then this interval must contain a point x in the chaotic critical orbit. We must have either that $x < c_*$ and $\hat{c}_* < H(x)$, or we have $x > c_*$ and $\hat{c}_* > H(x)$. But $H(x)$, by hypothesis, is the point in the chaotic critical orbit of $\hat{f}$ corresponding to x. The inequalities then violate the condition of combinatorial equivalence. Therefore our neighborhood V exists.

If the chaotic critical orbit of f contains the periodic point 0 (or 1), then combinatorial equivalence requires that the corresponding orbit of $\hat{f}$ does as well. In this case, we do not need to modify H on 0 (or 1) because H fixes the point as required for Theorem [4.4]. Otherwise, the chaotic critical orbit falls to the right of 0 and to the left of 1 for both f and $\hat{f}$ by definition of $\mathcal{F}$. A similar argument to that used to construct V shows that a symmetric open set V_0 containing 0 exists so that $H(V_0)$ also contains 0 and none of the chaotic critical orbit of $\hat{f}$. A set V_1 containing 1 and mapping to 1 and a neighborhood of infinity V_∞ with the same property also exist. We have that $V \cup V_0 \cup V_1 \cup V_\infty$ contains none of the chaotic critical orbit of f and its image under H contains no point in the chaotic critical orbit of $\hat{f}$.

We can find a symmetric orientation preserving Böttcher coordinate φ defined on a neighborhood of c_* U and conjugating f with the map $z \mapsto z^2$ defined on a disk centered at 0. Similarly, there exists a $\hat{\varphi}$ for $\hat{f}$ defined on neighborhoods $\hat{U}$ of $\hat{c}_*$. Then $\hat{\varphi}^{-1} \circ \varphi$ conjugates f and $\hat{f}$ on small U. We choose U small enough so that U is contained in V and $\hat{\varphi}^{-1} \circ \varphi(U)$ is contained in $H(V)$. Similarly Böttcher coordinates for the fixed point at infinity exist, φ_∞ and $\hat{\varphi}_\infty$, defined on neighborhoods U_∞ and $\hat{U}_\infty$ contained in V_∞ and $\hat{V}_\infty$ respectively. Notice that all these sets are open and can be chosen with smooth boundaries, so there is no obstruction to quasiconformal extensions.

We can now build a new quasiconformal homeomorphism H_0. The map H_0 is defined as fixing 0 and 1, equal to $\hat{\varphi}^{-1} \circ \varphi$ on U, and equal to $\hat{\varphi}_\infty^{-1} \circ \varphi_\infty$ on U_∞. Outside of $V \cup V_0 \cup V_1 \cup V_\infty$, H_0 is defined to be equal to H, and we extend H_0 quasiconformally everywhere else.

In preparation for invoking Theorem [4.4], we define the set U_0 to be the union of U, U_∞, $\{0, 1\}$, and the entire chaotic critical orbit of f. The properties (a) through (f) of Theorem [4.4] are satisfied by the construction of H_0 and U_0. The set U_0 maps into itself under f because each set in the union defining it also does, satisfying property (g). The functional equation, by hypothesis, is satisfied on the chaotic critical orbit. It is also satisfied by construction on U and U_∞ because of the property of Böttcher coordinates. It is satisfied on $\{0, 1\}$ because these points are fixed by H_0 and either fixed or flipped by f and $\hat{f}$. Thus property (h) is satisfied. Property (j) is easily verified because all the relevant points are on the real line. The set U_0 contains a neighborhood of c_* and infinity, and Theorem [4.4] can now be applied. Therefore there exists a special qc-conjugacy between f and $\hat{f}$. Q.E.D.

V. Applications of pull-back

§5.1. Characterizing bone-loops.

Conjugacy classes on bone-loops. The goal of this section is to prove that all points in a bone-loop represent maps in the same conjugacy class. We first show that the critical kneading sequences are the same for maps represented by points in a bone-loop. If λ and μ are in the same bone-loop, we know from the definition of the bone containing the bone-loop, that both maps f_λ and f_μ have a periodic critical orbit whose order on the real line is determined by the bone. We label these critical points c_λ^* and c_μ^*. The other two uncontrolled critical points are labeled c_λ and c_μ.

Lemma 5.1. Let L be a bone-loop in the parameter space of maps in $\mathcal{F}$, and let $\mathcal{L} = \{f_\lambda\}_{\lambda \in L}$. Then any two maps f_λ and f_μ in $\mathcal{L}$, share the same kneading data. In other words,

$$\theta_{f_\lambda}(c_\lambda^*) = \theta_{f_\mu}(c_\mu^*) \quad \text{and} \quad \theta_{f_\lambda}(c_\lambda) = \theta_{f_\mu}(c_\mu).$$

Proof. All maps we consider here have a periodic critical point c_λ^* with a fixed order and a fixed period. We will make crucial use of the pullback construction on hyperbolic maps, so we will need to worry about critical factors. The point c_λ^* is attracting and periodic, clearly in some periodic Fatou component B. By definition, c_λ^* is the primary point of B and the center, so it is mapped by ψ_B to zero. But c_λ^* is also a critical point of B, therefore the associated critical factor is zero. This holds true for all maps in $\mathcal{L}$.

We start by showing that $\theta_{f_\lambda}(c^*_\lambda)$ is constant, for $\lambda \in L$. Remember that as we continuously deform a map f_λ, the k-th kneading symbol of f_λ, $\theta^k_{f_\lambda}(c^*_\lambda)$, only changes when $f^{\circ k}_\lambda(c^*_\lambda)$ becomes a critical point. We define n to be the largest integer, if it exists, so that $\theta^k_{f_\lambda}(c^*_\lambda)$ is constant for all $\lambda \in L$ and $k < n$ (Note: n may be zero). If no such n exists then we are done, so we assume it does exist and try to arrive at a contradiction. We follow $f^{\circ n}_\lambda(c^*_\lambda)$ as λ is varied around the bone-loop L. By definition of n, there must be some map f_λ where the kneading symbol changes, and the new symbol $\theta^n_{f_\lambda}(c^*_\lambda)$ is critical. The number n must be less than the fixed period of c^*_λ in $\mathcal{L}$, or else $\theta^n_{f_\lambda}(c^*_\lambda)$ would be determined. Furthermore $\theta^n_{f_\lambda}(c^*_\lambda)$ cannot be c^*_λ, as this would mean the fixed period of c^*_λ in $\mathcal{L}$ had changed. Thus $\theta^n_{f_\lambda}(c^*_\lambda) = c_\lambda$. In other words, c^*_λ and c_λ, for f_λ, have the same periodic orbit.

The point c_λ falls somewhere on the orbit of c^*_λ and must be in some component B_1. We may have $B = B_1$. But c_λ is also a critical point, and again since c_λ is periodic, it is also the center of B_1. The associated critical factor of c_λ is therefore zero.

In a neighborhood of λ in L, the attracting periodic orbit of f_λ must persist, and the critical factors associated with c_λ and c^*_λ vary continuously. For f_λ itself, $f^{\circ n}_\lambda(c^*_\lambda)$ and c_λ are the same point, the primary point of B_1. In some neighborhood of λ on L, for a different map f_μ, $f^{\circ n}_\mu(c^*_\mu)$ and c_μ may no longer be equal, but $f^{\circ n}_\mu(c^*_\mu)$ is still the primary point of B_1, and c_μ is still in B_1. Let $\beta_\mu = \psi_{B_1}(c_\mu)$ and $0 = \psi_{B_1}(f^{\circ n}_\mu(c^*_\mu))$ be the corresponding critical point and primary point in the normalized component of any f_μ in this neighborhood of f_λ. We may thus consider the normalized component, the unit disk, as a model of the relative positions of c_μ and $f_\mu^{\circ n}(c_\mu^*)$ in B_1.

Following a map, f_μ, as μ is continuously deformed through λ in L, the point β_μ starts as a real value either to the right or left of zero. The point β_μ moves continuously until $\mu = \lambda$, where $f_\lambda^{\circ n}(c_\lambda^*) = c_\lambda$ and $\beta_\lambda = 0$. The question is what happens as we perturb μ further around L. The point β_μ may cross zero, either positive to negative or negative to positive. In this case, the symbol $\theta^n_{f_\mu}(c_\mu^*)$ is no longer critical, nor is it the same symbol as before when β_μ was on the other side of zero. As we continue to follow μ all the way around L, back to where we started, $\theta^n_{f_\mu}(c_\mu^*)$ must cross back to the original symbol. At the point where it crosses back, we find a parameter, $\mu \neq \lambda$, with $f_\mu^{\circ n}(c_\mu^*) = c_\mu$. For both maps, the original f_λ and the new f_μ, both critical points share an orbit which is periodic and has an order determined by the bone. We see that f_λ and f_μ are combinatorially equivalent, and furthermore f_λ and f_μ are both hyperbolic. The critical factors for both maps are all zero, so the hypotheses of Theorem [4.6] are satisfied. By the corollary, f_λ and f_μ are the same map. This is a contradiction, since μ and λ are defined as distinct points in parameter space.

In the other case, where β_μ does not cross zero, we also get a contradiction. If β_μ does not cross zero, it must stop altogether or turn around as μ is varied through λ. In either case, we can find two distinct parameters, μ_1 and μ_2, in a neighborhood of λ, where β_{μ_1} and β_{μ_2} are equal. Note that in this neighborhood the critical factor associated with c_μ is equal to β_μ, so the critical factors of f_{μ_1} and f_{μ_2} associated with the uncontrolled critical point are equal. As mentioned at the beginning of the proof, the critical factors associated with the periodic critical point are zero. If the critical factors associated with c_{μ_1} and c_{μ_2} are zero, we have the same argument as before. If not, both maps must have the uncontrolled critical point attracted to the orbit of the periodic critical point in the same way, and again we must have combinatorial equivalence between f_{μ_1} and f_{μ_2}, giving us the same contradiction.

We prove that $\theta_{f_\lambda}(c_\lambda)$ is constant for $\lambda \in L$ in nearly the same way. The number n is defined in the same manner with respect to $\theta_{f_\lambda}(c_\lambda)$, and we find a map f_λ where $\theta^n_{f_\lambda}(c_\lambda)$ is critical. The case where $\theta^n_{f_\lambda}(c_\lambda) = c_\lambda$ occurs when f_λ has two distinct periodic critical orbits. When we perturb λ in L, the orbit of c_λ will no longer be periodic but the periodic orbit to which c_λ is attracted will persist. We proceed exactly as before, by finding two maps sharing the same critical factors. The case where $\theta^n_{f_\lambda}(c_\lambda) = c^*_\lambda$ also proceeds in nearly the same way. This corresponds to the case where f_λ maps c_λ into the periodic orbit of c^*_λ, but c_λ is not periodic itself. The only problem may occur if the preimage of c^*_λ in the Fatou component with c_λ, the primary point, does not stay real when we perturb λ. The primary point will not be the center of the component in this case, and β_λ will always be zero. We define α_λ as the normalized image of the primary point and follow it as it crosses zero instead. All other steps should follow in the same way. Q.E.D.

We are now show that all maps in a bone-loop are combinatorially equivalent.

Theorem 5.2. Let L be a bone-loop. Then if f and $\hat{f}$ are in $\mathcal{L} = \{f_\lambda\}_{\lambda \in L}$, they are combinatorially equivalent.

PROOF. We know f and $\hat{f}$ share the same kneading data, by Lemma [5.1]. This is almost enough to show combinatorial equivalence by itself, but not quite. Consider two points, $x = f^{\circ j}(c)$ and $y = f^{\circ k}(c)$, where c is a critical point of f. Define the corresponding points $\hat{x} = \hat{f}^{\circ j}(\hat{c})$ and $\hat{y} = \hat{f}^{\circ k}(\hat{c})$. We need to show that if $x < y$ then $\hat{x} < \hat{y}$. Kneading theory gives us this immediately in many cases.

The kneading sequences of x and y are just shifts of $\theta_f(c)$. The sequences of $\hat{x}$ and $\hat{y}$ are the same shifts of $\theta_{\hat{f}}(\hat{c})$. Since the kneading data of f and $\hat{f}$ are the same, we always have $\theta_f(x) = \theta_{\hat{f}}(\hat{x})$ and $\theta_f(y) = \theta_{\hat{f}}(\hat{y})$. By the ordering property of kneading sequences, if $x < y$ and $\theta_f(x) \neq \theta_f(y)$, then we must have $\theta_f(x) < \theta_f(y)$. Since the corresponding sequences are equal, we get $\theta_{\hat{f}}(\hat{x}) < \theta_{\hat{f}}(\hat{y})$, which implies that $\hat{x} < \hat{y}$.

Kneading theory says nothing if $\theta_f(x) = \theta_f(y)$, however. So to finish the proof, we assume x and y have the same kneading sequence. We can also assume that this sequence does not contain a critical symbol. If it does, then x and y are precritical, and two precritical points with the same kneading sequence must be equal. Note that x and y must be in the forward orbit of the uncontrolled critical point c because the kneading sequence of the periodic critical point contains an infinite number of critical symbols.

Suppose $x < y$ and $\hat{x} > \hat{y}$ for f and $\hat{f}$, respectively. Then somewhere in between f and $\hat{f}$ in $\mathcal{L}$, there must be a map g where the corresponding points cross, i.e. a map where there exists a point $z = g^{\circ j}(c) = g^{\circ k}(c)$. Therefore c is a strictly preperiodic critical point of g. Let λ be the multiplier of the periodic orbit of z. We cannot have $\lambda = 0$ or the orbit of z would contain a critical symbol. Suppose $0 < |\lambda| \leq 1$, where the periodic orbit is attracting or indifferent. In general, we must have some sort of attracting basin, which may not even intersect the real line. In any case, we know that the immediate basin of this periodic orbit must contain a critical point. In our case, this point must be c, as the periodic critical point is in its own attracting basin. But this is a contradiction since c is strictly preperiodic, and immediate basins do not contain preperiodic critical points. So $|\lambda| > 1$, and thus z and c must be in the Julia set of g.

We return to studying f. The points x and y share the same kneading sequence and by the ordering property, so do all points in between. As noted, this sequence is the same as that of $\theta_g(z)$, which is periodic. Let I be the set of points in $[0, 1]$ under f that share this sequence. Since $x < y$, the set I must be a non-trivial interval. Since the map f acts as a shift operator on kneading sequences, and since the kneading sequence of points in I is periodic, I must be mapped into itself by some iterate of f. Furthermore, the sequence contains no critical symbols, so I is mapped monotonically into itself. There then must be an attracting fixed point for this iterate of f, whose multiplier μ must satisfy $|\mu| \leq 1$. This attracting orbit will persist as we deform f into g along $\mathcal{L}$, and μ will also vary continuously. Throughout this deformation $|\mu|$ is always bounded by 1 as the periodic point is always attracting, but this implies the multiplier for g, λ, must also satisfy $|\lambda| \leq 1$.

This is a contradiction that rules out the original assumption that $f^{\circ j}(c)$ and $f^{\circ k}(c)$ crossed when f was deformed into $\hat{f}$. Thus combinatorial equivalence is proved. Q.E.D.

Critical points of maps in bone-loops. Theorem [5.2] combined with Theorem [4.6] can be used to sharply restrict the type of map represented by a bone-loop. We show that any map represented by a point in a bone-loop always has exactly one periodic critical point and one critical point in the Julia set.

Proposition 5.3. Let L be a bone-loop, and let f_λ be in $\mathcal{L} = \{f_\lambda\}_{\lambda \in L}$. Then one critical point of f_λ is chaotic, i.e. in the Julia set of f_λ.

PROOF. If there exists even one f_λ in $\mathcal{L}$ with a chaotic critical point, then all maps in $\mathcal{L}$ must have a chaotic critical point. This follows from combinatorial equivalence. For instance, if f_λ has a preperiodic critical point then every other map in $\mathcal{L}$ must also have a preperiodic critical point. If f_λ has a critical point which is not preperiodic, the kneading sequence of every point in the forward orbit must be aperiodic or else the critical point is not in the Julia set, and all other maps in $\mathcal{L}$ must have the same aperiodic critical kneading sequence. A chaotic critical point cannot be periodic.

Suppose there is no map in $\mathcal{L}$ with a chaotic critical point. Let c_λ be the uncontrolled critical point of maps in $\mathcal{L}$. Then all points in L represent maps for which c_λ is attracted to some periodic orbit. The points in L representing hyperbolic maps, where the periodic orbit is not indifferent, must be dense in L because indifferent periodic orbits are unstable. For the hyperbolic maps, the critical factor of c_λ, which we label β_λ, varies continuously with $\lambda \in L$. We claim that we can always find two hyperbolic maps with the same critical factor β_λ.

Let I be some connected component of the points in L representing hyperbolic maps. If I is all of L, then our claim is immediately true because L is a closed loop, and β_λ cannot be a continuous injective map from I into $\mathbb{R}$. So suppose that I is some interval of points in L. In the limit as we perturb λ along I, f_λ changes from being a hyperbolic map to the limiting case where f_λ has an indifferent cycle, the critical factor β_λ must approach either -1 or 1 in the limit. If no two points on I have the same critical factor, then β_λ forms an injective function from I *onto* $(-1, 1)$. But if I is not all of L, then there must exist a separate connected component of representing hyperbolic maps whose critical factors also take all values in $(-1, 1)$. Therefore, in all cases, we can find two hyperbolic maps with the same β_λ.

The critical factor associated with the other critical point, the periodic one, must be zero. So we can apply Theorem [4.6] and its corollary. These show that the two hyperbolic maps, which are constructed as distinct maps, are actually the same map. This contradiction rules out the possibility that there are no maps with chaotic critical points in $\mathcal{L}$. Therefore, all maps have a chaotic critical point. Q.E.D.

Corollary. Between any two maps in the same bone-loop, there exists a topological conjugacy on the real line.

PROOF. This follows because points on a bone-loop represent maps that are combinatorially equivalent and have no indifferent periodic cycles, as neither critical orbit can be in the basin of attraction of such a cycle. Q.E.D.

§5.2. Non-existence of bone-loops.

For this section, we make the following claim.

Theorem 5.4. Let f and $\hat{f}$ in $\mathcal{F}$ be two maps represented by points on the same bone-loop. Then there exists a special qc-conjugacy between them.

The proof of this theorem involves the machinery of box mappings and is the subject of Chapters VI through IX. Using this theorem now however, we can prove that bone-loops do not exist. In this way we reduce the Connected Bone Conjecture (and monotonicity) to proving Theorem [5.4].

Paths of conjugate maps. Quasiconformal conjugacy classes, families of maps that are conjugate to each other by quasiconformal homeomorphisms, have been used effectively in the study of the quadratic family. Quasiconformal conjugacy classes of quadratic maps, if they have more than one member, always form *open* sets in parameter space (See [MSS]). The argument used to prove this is very general, and we reproduce it here in the context of the real cubic family. The quasiconformal conjugacy classes of cubic maps form open paths in real parameter space.

Recall that χ_H is the complex dilatation of the map H. We use the convention that H^z and $H^{\bar{z}}$ are the partial derivatives of H with respect to z and $\bar{z}$. So $\chi_H = H^{\bar{z}}/H^z$.

Theorem 5.5. Let f, $\hat{f} \in \mathcal{F}$ be conjugated by a special qc-conjugacy H, with dilatation $|\chi_H|$ bounded by $0 < K < 1$. Then there is a *path* of maps $f_t \in \mathcal{F}$, for $t \in (-1, 1)$, forming a continuous deformation, *within* $\mathcal{F}$, between $f_0 = f$ and $f_K = \hat{f}$. All these maps, f_t, are conjugate to each other by special qc-conjugacies.

PROOF. Given a form

$$\mu dz + \lambda d\bar{z} \qquad \text{with } \left|\frac{\lambda}{\mu}\right| \leq K < 1,$$

the Measurable Mapping Theorem, Theorem [2.1], states there exists a quasi-conformal map H, unique up to composition with a Moebius transformation, that solves the equation $\chi_H = \frac{\lambda}{\mu}$. Moreover, the complex dilatation of H is bounded by K. Given a specific H, we can work in the other direction and obtain a form

$$H^z dz + H^{\bar{z}} d\bar{z} \qquad \text{with } \left|\frac{H^{\bar{z}}}{H^z}\right| \leq K < 1.$$

Notice that H solves the equation $\chi_H = H^{\bar{z}}/H^z$.

We build our path by perturbing this form. To be precise, we define

$$\omega_t = H^z dz + \frac{t}{K} H^{\bar{z}} d\bar{z}.$$

and the complex dilatation associated with this form as

$$\chi_t = \frac{t}{K} \frac{H^{\bar{z}}}{H^z} = \frac{t}{K} \chi_H.$$

Clearly for $|t| < 1$, the value $|\chi_t(z)|$ is bounded by $|t|$ itself as required by the hypothesis of the Mapping Theorem. To eliminate consideration of a Moebius transformation, we require that any solution to these forms fix the points 0, 1, and ∞. Then the Measurable Mapping Theorem gives us, for each t in $(-1, 1)$, a unique quasiconformal map $H_t(z)$, and we can consider its attributes as a definition.

Definition. (H_t) H_t is quasiconformal, fixing $\{0, 1, \infty\}$, with $\chi_{H_t} = \chi_t$.

From these maps, we can define our path

$$f_t = H_t^{-1} \circ f \circ H_t. \tag{1}$$

Some facts are immediately obvious. Each f_t is conjugate to f and thus they are conjugate to each other. The map H_0 solves the equation $\chi_{H_0} = 0$, thus $H_0^{\bar{z}} = 0$ and H_0 must be conformal. The identity map, $H_0(z) = z$, is the unique conformal homeomorphism on $\hat{\mathbb{C}}$ fixing 0, 1, and ∞, and we get $f_0 = f$. The map H_K is the original special qc-conjugacy H, so $f_K = \hat{f}$ as required.

The Measurable Mapping Theorem also states that the solution H_t depends analytically on the original form ω_t, and this then holds true for f_t as well. So the f_t, for $t \in (-1, 1)$ does form a continuous path connecting f and $\hat{f}$. We still need to check that, for each t, f_t is in $\mathcal{F}$.

We start by showing that f_t is conformal. Since H_t fixes infinity by definition, f_t is then a conformal map of the sphere which also fixes infinity.

The only such maps are polynomials, and we get immediately that f_t is a cubic polynomial.

From the functional equation on H_t, we must have $\chi_{f_t \circ H_t} = \chi_{H_t \circ f}$, and we start by evaluating these dilatations. After expanding the partial derivatives of compositions (See [10.1]) we get

$$\frac{(f_t^z \circ H_t)H_t^{\bar z} + (f_t^{\bar z} \circ H_t)\bar H_t^{\bar z}}{(f_t^z \circ H_t)H_t^z + (f_t^{\bar z} \circ H_t)\bar H_t^z} = \frac{(H_t^{\bar z} \circ f)\bar f^{\bar z}}{(H_t^z \circ f)f^z}.$$

The simplifications in the expansion of $\chi_{H_t \circ f}$ are due to the fact that $\bar f^z$ and $f^{\bar z}$ are zero. By the definition of H_t

$$\frac{(H_t^{\bar z} \circ f)\bar f^{\bar z}}{(H_t^z \circ f)f^z} = \frac{t}{K}\frac{(H^{\bar z} \circ f)\bar f^{\bar z}}{(H^z \circ f)f^z} = \frac{t}{K}\chi_{H \circ f} = \frac{t}{K}\frac{H^{\bar z}}{H^z} = \frac{H_t^{\bar z}}{H_t^z}.$$

Cross multiplying this version of $\chi_{H_t \circ f}$ with the expansion of $\chi_{f_t \circ H_t}$ and cancelling the common term gives

$$(f_t^{\bar z} \circ H_t)\bar H_t^{\bar z} H_t^z = (f_t^{\bar z} \circ H_t)\bar H_t^z H_t^{\bar z}.$$

If $f_t^{\bar z} \circ H_t$ is not identically zero, we cancel it and obtain

$$\frac{\bar H_t^{\bar z}}{\bar H_t^z} = \frac{H_t^{\bar z}}{H_t^z} \quad \text{or} \quad \chi_{\bar H_t} = \chi_{H_t}.$$

Note in general, however, that

$$\chi_{\bar H_t} = \frac{\bar H_t^{\bar z}}{\bar H_t^z} = \frac{\overline{H_t^z}}{H_t^{\bar z}} = 1/\bar\chi_{H_t}.$$

So we would have $1/\bar{\chi}_{H_t} = \chi_{H_t}$. However, this is impossible because H_t is quasiconformal and $|\chi_{H_t}| < 1$. Therefore $f_t^{\bar{z}} \circ H_t = 0$. The map H_t is a homeomorphism and takes on zero at only one point, so $f_t^{\bar{z}}$ must be identically zero. Therefore f_t is conformal.

If, for each t, H_t is a special qc-conjugacy, then the second corollary to Lemma [3.1] shows that H_t is symmetric and preserves the order of $\mathbb{R}$. Since the original f also has these properties, f_t will also be symmetric and preserve the order of $\mathbb{R}$. Moreover, f permutes the set $\{0, 1\}$ in some way, and since H_t fixes these points, f_t must permute them in the same way as f. We conclude that $f_t \in \mathcal{F}$, proving the last statement in our theorem. All we have left to do is show that H_t is special quasiconformal.

We know already that H_t is quasiconformal and fixes $\{0, 1, \infty\}$. The original H is special quasiconformal, so $H^{\bar{z}}(z) = 0$ for z in the Fatou set of f. For the same z,

$$\chi_{H_t}(z) = \chi_t(z) = \frac{t}{K}\frac{H^{\bar{z}}(z)}{H^{z}(z)} = 0.$$

But if $\chi_{H_t} = 0$, then we must have $H_t^{\bar{z}}(z) = 0$ on the Fatou set of f, as required for a special qc-conjugacy.

We need to show that H_t is symmetric. The original map H is symmetric, so we know $H^{\bar{z}}$, H^{z}, and $\chi_H = H^{\bar{z}}/H^{z}$ are also symmetric functions. By definition,

$$\chi_{H_t} = \frac{t}{K}\chi_H.$$

Since t is always chosen to be real, χ_{H_t} is symmetric as well. If we define

$H_*(z) = \bar{H}_t(\bar{z})$, then

$$\chi_{H_*}(z) = \frac{H_*^{\bar{z}}}{H_*^{z}} = \frac{\bar{H}_t^{\bar{z}}(\bar{z})}{\bar{H}_t^{z}(\bar{z})} = \bar{\chi}_{H_t}(\bar{z}) = \chi_{H_t}(z).$$

Therefore H_* is a quasiconformal map fixing $\{0, 1, \infty\}$ and it solves the same Beltrami equation as H_t. The Measurable Mapping Theorem says that such a map is unique, so we must have $H_t(z) = H_*(z) = \bar{H}_t(\bar{z})$. Therefore H_t is symmetric.

As an immediate consequence of symmetry, H_t must preserve $\mathbb{R}$, and because zero and one are fixed, H_t must preserve the order of $\mathbb{R}$ as well. Therefore we have determined that H_t is a special qc-conjugacy, concluding the proof of the theorem. Q.E.D.

We did not make any assumptions about K, the bound on the dilatation of H, in this construction. However, we need to do so in the next section. Assuming a special qc-conjugacy H exists between f and $\hat{f}$, it is quasiconformal and its complex dilatation χ_H is bounded on $\hat{\mathbb{C}}$. Certainly there exists a least upper bound on $|\chi_H|$, and since special qc-conjugacies are unique, the least upper bound is a well-defined operator on the pair of functions, f and $\hat{f}$. We normalize the path construction by requiring the bound, K, used in Theorem [5.5] to be minimal.

Definition. (maximal path) Given f and $\hat{f} \in \mathcal{F}$, and given a special qc-conjugacy, H, between them, the *maximal path* connecting f and $\hat{f}$ is the set of points in parameter space representing the path of maps in $\mathcal{F}$ given by Theorem [5.5] using the minimal dilatation bound on H. The maximal path, if it exists at all is unique, since special qc-conjugacies are unique.

Lemma 5.6. Suppose L is the maximal path of parameters deforming f_{λ_0} into f_{λ_1} in $\mathcal{F}$. Then the least upper bound operator, $K_{\lambda_0,\lambda_1}(\lambda)$, equal to the minimum complex dilatation bound on the conjugacy between f_λ and f_{λ_0}, is a well-defined and continuous map on L.

PROOF. The uniqueness of special qc-conjugacies shows that the least upper bound operator is well-defined. By definition $K_{\lambda_0,\lambda_1}(\lambda_1)$ is the minimal bound on the dilatation of H conjugating f_{λ_0} and f_{λ_1}, i.e. $|\chi_H(z)|$. But as in Theorem [5.5], we have $\chi_{H_t} = (t/K)\chi_H$, where $K = K_{\lambda_0,\lambda_1}(\lambda_1)$. Therefore the minimal bound on $|\chi_{H_t}(z)|$ is just $|t|$. But H_t conjugates f_{λ_0} and some map $f_{\lambda_t} \in \mathcal{F}$, so

$$K_{\lambda_0,\lambda_1}(\lambda_t) = |t|.$$

Since both λ_t and $|t|$ vary continuously with respect to t, we are done. Q.E.D.

Contradiction. Suppose that two maps, f_{λ_0} and f_{λ_1}, are quasiconformally conjugate and λ_0 is in some bone in parameter space. We know f_{λ_0} must have a periodic critical orbit occuring in some fixed order on the real line. There is, of course, a path connecting the points λ_0 and λ_1 in parameter space. Moreover, every point along this path must represent a map having a periodic critical orbit with the same order and period because special qc-conjugacies preserve these properties of f_{λ_0}. So the entire path is contained in the same bone as λ_0 We can use this observation to reduce the Connected Bone Conjecture to a simpler form.

The maximal path connecting the points λ_0 and λ_1 forms an *open* path, and it this openness that is incompatible with the closed and compact bone-loop.

Lemma 5.7. A bone-loop L cannot consist of points representing maps which are all conjugate by special qc-conjugacies.

PROOF. Assume we can find a special qc-conjugacy between any two maps in $\mathcal{L} = \{f_\lambda\}_{\lambda \in L}$. We pick an arbitrary map $f_{\lambda_0} \in \mathcal{L}$. For any other map f_{λ_1} in $\mathcal{L}$, there exists a maximal path between λ_0 and λ_1. This path must be contained in the same bone, and since L is the connected component of the bone containing λ_0, the path must be contained in L. Since we assume special qc-conjugacies always exist, we can cover L with maximal paths, which form open sets in the subset topology on L. We define a new least upper bound operator on L by

$$K_{\lambda_0}(\mu) = K_{\lambda_0,\lambda_1}(\mu) \qquad \text{for any } \lambda_1 \in L.$$

Once again, this is well-defined since special qc-conjugacies are unique. It does not matter what λ_1 we pick, since we always end up taking the bound on the same special qc-conjugacy between f_{λ_0} and f_μ. By Lemma [5.6], K_{λ_0} is a continuous function on L.

But from any maximal path constructed around λ_0, we can construct a sequence of maps, $\{f_{\lambda_k}\}_{k=1}^{\infty}$, where the dilatation bound on the conjugacy between f_{λ_0} and f_{λ_k} is arbitrarily close to one. Since L is compact, there must be an accumulation point, λ_∞, of this sequence in L. By continuity, we must have $K_{\lambda_0}(\lambda_\infty) = 1$. But this is a contradiction. By definition, the special qc-conjugacy between f_{λ_0} and f_{λ_∞} is quasiconformal, and there must exist a bound on the complex dilatation which is strictly less than one, thus $K_{\lambda_0}(\lambda_\infty)$ is always less than one. Q.E.D.

This result and Theorem [5.4] imply the Connected Bone Conjecture. All that remains to be proven is Theorem [5.4], for which we will need elaborate machinery, originally developed to study quadratic maps.

VI. Box mappings and branchwise equivalences

To prove Theorem [5.4], we will transfer into a cubic setting tools designed originally for analyzing real quadratic maps. These objects, called box mappings, are part of the class of functions known as **induced** mappings, functions which are formed piecewise from the iterates of some continuous base function. The primary example of an induced map (and a box mapping as well) is a first return map, formed from a continuous function and an interval. Induced maps in various different forms have been used frequently to study metric and expansion properties of dynamical systems, particularly unimodal maps of the interval. For example, see [GJ].

§6.1. Standard box mappings.

The type of induced box mapping used here has continually evolved through the work of Graczyk, Jakobson, and Swiatek. In particular see [JS1], [JS2], and [GS2]. Later versions of this type of box mapping have been used to prove that hyperbolic maps are dense in the real quadratic family, by first proving Theorem [1.1] (See [Sw] and [GS1]). It provides a complete machinery for constructing quasiconformal conjugacies between conjugate maps. The definition presented here is taken directly from [GS1].

Although the box mappings we use will always be induced maps, induced from cubic maps in fact, the definition itself does not require the map to be induced from some other map. Actually, the manner in which box mappings are defined is crucial. Since the definition makes no mention of how the box mapping is constructed, we are able to transfer theorems originally designed for quadratic or unimodal applications to cubic maps.

First return maps. The inspiration for box mappings comes from the so-called *first return map* induced from some other function f.

Definition. (first return maps) Let f be some function from a space V, often taken to be $\mathbb{R}$, into itself. Some initial subset, U, of the domain of f is selected, usually an interval for real f, and for each $x \in U$, the orbit under f of x, $\{f(x), f^{\circ 2}(x), \dots\}$, is examined to see if it intersects U. The first return map of f to the set U is defined to be the map taking each point to the first point in its orbit inside U, if this point exists. The first return map is undefined elsewhere.

Suppose f is a continuous real map, the set U is an interval of $\mathbb{R}$, and Φ is the resulting first return map of f to U. Then we say that Φ has been **induced** from f, and we call U the **inducing interval**.

The first return map Φ is a map *induced* from f. The domain of Φ is in general disconnected, and on connected components, Φ is equal to an iterate of f. The map restricted to these components is called a **branch** of Φ. A primary feature here is that by post composing a branch of Φ with Φ itself, a new induced map, consisting of higher iterates of f, is created.

Real box mappings. The **Schwarzian derivative** of a map f is defined to be

$$Sf = \left(\frac{f'''}{f'}\right) - \frac{3}{2}\left(\frac{f''}{f'}\right)^2.$$

Definition. (real box mappings) Let Φ be a function defined on an open subset U of the real line into the real line. Restrictions of Φ to the connected components of U will be referred to as branches of Φ. For Φ to be a real box mapping, it must have the following properties.

(1) $B = (-a, a)$ for some $a > 0$ is a connected component of U called the **central domain**, while $\psi = \Phi|_B$ is the **central branch**.

(2) There exists a set B' chosen to be the smallest interval symmetric with respect to zero that contains the range of ψ. B' must contain B.

(3) $\psi = h(x^2)$ where h is a diffeomorphism onto its image B' with non-positive Schwarzian derivative.

(4) All branches of Φ different from the central branch are diffeomorphisms onto their respective images and have non-positive Schwarzian derivative.

(5) If Δ is any connected component of U, then Δ is disjoint from the border of B'.

(6) If V is the range of some monotone (non-central) branch of Φ and Δ is a connected component of U, then $\Delta \cap \partial V = \emptyset$.

The requirement on the Schwarzian derivative comes from the application of induced maps to the study of smooth dynamics on the interval, where a negative Schwarzian derivative is a hypothesis of the Koebe Distortion Theorem for instance(see [JS1]). For cubic maps, this property comes comes from using strictly bimodal maps (See the corollary to Theorem [8.7]).

Some of the analysis involving the Schwarzian derivative, in particular finding bounds on quasisymmetric distortion, has now been replaced by the analysis of quasiconformal maps (see Lemma [10.1]). The remaining properties derived from the Schwarzian are hidden in Theorems [6.2] and [6.1], so we will make little further reference to it.

Holomorphic box mappings. The definition of a *holomorphic* box mapping is the same conceptually as the real case and is formally obtained by changing a few words.

Definition. (holomorphic box mappings) Let Φ be a function defined on an open subset U of the complex plane into the complex plane. Restrictions of Φ to the connected components of U will be referred to as branches of Φ. For Φ to be a holomorphic box mapping, it must have the following properties.

(1) There exists an open topological disk B, a connected component of U, that is mapped onto itself by the transformation $z \mapsto -z$. B is called the **central domain**, and $\psi = \Phi|_B$ is called the **central branch**.

(2) Let B' be the range of ψ.

(3) $\psi = h(z^2)$ where h is univalent onto its image B'.

(4) All branches of Φ different from the central branch are holomorphic and one to one onto their respective images.

(5) If Δ is any connected component of U, then Δ is disjoint from the border of B'.

(6) If V is the range of some univalent (non-central) branch of Φ and

Δ is a connected component of U, then $\Delta \cap \partial V = \emptyset$.

Even though the domain of a holomorphic box mapping is open in $\mathbb{C}$, holomorphic box mappings should still be considered a tool for studying the real line. We will always use symmetric box mappings, whose restriction to $\mathbb{R}$ is a real box mapping. The conformal properties will always be used in conjunction with the properties of (symmetric) quasiconformal maps to draw conclusions about the underlying real maps.

We will use the expression *box mapping* where both real and holomorphic box mapping can be substituted.

Definition. (box mapping types) We distinguish two special types of box mappings, both real and holomorphic. A **type I** box mapping is determined by the condition that all non-central branches have range B. A **type II** box mapping is characterized by the property that all non-central branches have range B'.

§6.2. Branchwise equivalences.

Box mappings will be used as a tool for showing that conjugacies between cubic maps are quasiconformal. This is accomplished by constructing a sequence of quasiconformal maps which are *partial* conjugacies, or maps which satisfy the necessary functional equation on only part of the complex plane. We then show that a partial conjugacy suitable for Theorem [4.7] exists in the closure of the family of partial conjugacies. Branchwise equivalences play the role of these partial conjugacies. The all important functional equation for conjugate maps

$$H \circ f = \hat{f} \circ H,$$

where f and $\hat{f}$ are the conjugate maps and H is the conjugacy, plays a crucial role in the following definition. A branchwise equivalence will be a partial conjugacy between f and $\hat{f}$ that satisfies the functional equation only on the domain boundary of the box mapping Φ induced from f.

The notions of branchwise equivalence for holomorphic and real box mappings are completely equivalent. We will express the definition for holomorphic box mappings, giving in brackets the few changes needed to produce the definition for real box mappings.

Definition. (branchwise equivalences) Let Φ and $\hat{\Phi}$ be two holomorphic (real) box mappings. An orientation preserving homeomorphism H of the complex plane (real line) onto itself is considered a branchwise equivalence between Φ and $\hat{\Phi}$ provided that H maps the domain of Φ onto the domain of $\hat{\Phi}$ and satisfies the following dynamical condition. If V is a connected component of the domain of Φ, ζ is Φ restricted to V and extended continuously to the closure of V, while $\hat{\zeta}$ is $\hat{\Phi}$ restricted to $H(V)$ and continuously extended to the closure, then

$$H \circ \zeta = \hat{\zeta} \circ H$$

on the boundary of V.

A branchwise equivalence between two box mappings is not unique. We have the following similarity relation, which stresses that the important part of a branchwise equivalence is where is satisfies the functional equation.

Definition. (similarity between branchwise equivalences) Let Φ and $\hat{\Phi}$ be box mappings, real or holomorphic, with branchwise equivalences H_1 and H_2. Let D be the boundary of the domain of Φ. We say that H_1 and H_2 are *similar* if they coincide on D and are homotopic to each other in the complement of D. We will use the notation $[H]$ for the similarity class of H.

§6.3. Inducing algorithm.

The motivations for introducing box mappings are Theorems [6.2] and [6.1]. These results are the culmination of research (see [Sw], [GS1], and also [GS2]) on constructing quasiconformal conjugacies between real quadratic polynomials. It is crucial to this work and a testimonial to the power of the tools being used that these theorems can be immediately applied to cubic maps. These tools include Yoccoz partitions, pull-back and inducing constructions using quasiconformal maps, and the λ-Lemma. All these tools are essential for this work as well and will be presented later in detail. In fact, the construction of box mappings induced from cubic polynomials and the construction of quasiconformal branchwise equivalences between these mappings can be viewed as a miniature version of the inducing construction used in [GS1]. In particular, see Lemma [7.1]. The delicate moduli estimates for nested annuli around the chaotic critical point and renormalization are the only significant tools used in the quadratic case that are not directly used in the cubic case. Although, we inherit these tools in the new form of Theorems [6.2] and [6.1].

The central construction in the proofs of Theorem [6.2] and [6.1] is the inducing algorithm on box mappings. In short, the branches of a box mapping are composed with each other to produce new box mappings whose disconnected domains form finer partitions of the dynamical plane. Branch-wise equivalences, which can be shown to be quasiconformal, defined with respect to these finer partitions approach a partial conjugacy, suitable for Theorem [5.4], in the limit. We present an outline of the inducing algorithm as it applies to Theorems [6.2] and [6.1], taken directly from [GS1].

Step A - filling in.. Suppose that ϕ is a real or holomorphic box mapping. Choose a set S of monotone (univalent) branches of ϕ all with the same range R which must contain the closure of B.

Then define a sequence of (holomorphic) box mappings ϕ_i as follows. ϕ_0 is equal to ϕ outside of B and the identity on B. S_0 is S. Given ϕ_i, $i \geq 0$, and a set S_i of monotone (univalent) branches construct ϕ_{i+1} as follows. Set $\phi_{i+1} = \phi_i$ except on the union of domains of branches of S_i, and $\phi_{i+1} = \phi_0 \circ \phi_i$ on the union of domains of branches from S_i. At the same time, S_{i+1} becomes the set of all branches of ϕ_{i+1} in the form $\zeta_1 \circ \zeta_2$ where ζ_1, ζ_2 belong to S_i.

The (holomorphic) box mapping Φ which is the outcome of Step A is defined on the set of points z such that the sequence $\phi_i(z)$ is defined for all i and eventually constant. Then $\Phi(z) = \lim_{i\to\infty} \phi_i(z)$ if $z \notin B$ and $\Phi(z) = \phi(z)$ if $z \in B$.

When Φ is compared with ϕ, we see that all branches except those with range R have been left undisturbed, while those branches onto R have all vanished and been replaced with compositions among themselves and with other branches with different images.

A typical example of filling-in occurs if ϕ is a type II (holomorphic) box mapping. In that case there is only one possibility for R, namely $R = B'$ and the outcome is a type I (holomorphic) box mapping with the same B and B'. Note that this is just constructs the first return of all monotone branches to B leaving the central branch unchanged.

Step B - critical filling. Now suppose that a (holomorphic) box mapping ϕ is given. For Step B to be feasible, the critical value of ϕ has to belong to the domain of ϕ. Construct ϕ_0 by changing ϕ on the central domain only, and making it the identity there. Then define Φ again by changing ϕ on the central domain only, where we set $\Phi = \phi_0 \circ \phi$. This Φ is the outcome of Step B applied to ϕ.

Observe that for Φ the range B' is the central domain of ϕ. The central branch of Φ has the form $\zeta \circ \psi$ where ψ is the central branch of ϕ and ζ is either a monotone (univalent) branch of ϕ, or the identity restricted to B.

Again, the particular case of most interest to us is when ϕ is a type I (holomorphic) box mapping. In that case, Φ is a type II (holomorphic) box mapping. According to whether the critical value of ϕ is in the central domain of ϕ or not, we describe the situation as either a *close* or *non-close return*.

Inducing steps. We will now define a **simple inducing step** for type I (holomorphic) box mappings. If ϕ is such a mapping, the simple inducing step is defined to be Step B followed by Step A. As remarked above, the outcome will be a type I (holomorphic) box mapping. We make a distinction between a *close* and *non-close* return for ϕ, depending on whether the critical value of ϕ is in the central domain of ϕ. The simple inducing step is defined provided that Step B is defined, i.e. the critical value of ϕ is in the domain of ϕ.

Let us now define a type of box mapping which is an obstacle to continuing an inducing construction.

Definition. (terminal box mappings) A box mapping is called terminal if there is an open interval $I \subset B$ containing the critical point of ϕ so that $\phi(I) \subset I$ and $\phi(\partial I) \subset \partial I$. The interval I (which must be unique) will then be called the restrictive interval of ϕ.

Now we define the **type I inducing step**. It takes a (holomorphic) type I box mapping ϕ. The type I inducing step is defined recursively so that it is equal to the simple inducing step if ϕ makes a non-close return, and is equal to the type I inducing step applied to ϕ_1 obtained by the simple inducing step for ϕ otherwise. In other words, the type I inducing step is an iteration of simple inducing steps continued until the first non-close return occurs. This definition may fail if at some point the simple inducing step is no longer defined, or, more interestingly, if a non-close return is never achieved.

If the last possibility occurs for a box mapping ϕ, and $\phi(0) \neq 0$, then it must be *terminal*. Indeed, let ψ denote the central branch of ϕ, and B its central domain. If a non-close return never occurs, then the critical value must be contained in $\psi^{-n}(B)$ for any $n \geq 0$. These intervals form a descending sequence, and the intersection must be more than a point, since otherwise 0 would be fixed by ψ. So the intersection is a non-degenerate interval symmetric with respect to 0 and invariant under ψ which meets the definition of a terminal box mapping.

The **inducing algorithm** is applied to box mappings as follows. We begin with the initial type II holomorphic box mapping ϕ_0 and first fill it in to get a type I map. Then we proceed by a sequence of type I inducing steps. There are two obstacles to performing a type I inducing step. One is if the critical value escapes from the domain of the box mapping. The other possibility is that a terminal mapping is obtained. Assuming our box mappings are all induced from some polynomial f, the first difficulty occurs only if the chaotic critical orbit of f is *non-recurrent*. In this case, the inducing algorithm will not proceed to a limit. As described in Theorem [9.2] however, the desired partial conjugacy will still be constructed before the algorithm fails. The second obstacle can also occur and leads to *renormalization*. In Theorem [9.5], we show that the problem of constructing the necessary partial conjugacy in this case actually reduces to the same problem for *quadratic* maps, where it follows from the original quadratic results

(Theorem [1.1]). Issues of renormalization in the cubic case, in contrast to the quadratic case, are of little concern in the construction of conjugacies. This is because any renormalization of a cubic map, in the situations we must deal with, is quadratic-like (in the sense of Douady and Hubbard) and can be dealt with using previous results.

Pull-back. We show how the inducing algorithm is used to produce a sequence partial conjugacies, closer and closer to a full conjugacy. We construct two initial box mappings ϕ_0 and $\hat{\phi}_0$, induced from f and $\hat{f}$ respectively, and a quasiconformal branchwise equivalence between them. The inducing algorithm is applied to the two box mappings in parallel. For each step of the algorithm, an operation called *pull-back* is applied to the branchwise equivalence, producing a new quasiconformal branchwise equivalence between the new box mappings ϕ_i and $\hat{\phi}_i$. This new branchwise equivalence satisfies the functional equation on a larger set.

The basic step involved in pull-back is as follows. Suppose Δ_1 is a component of the domain of the box mapping ϕ, and let $\zeta_1 = \phi|_{\Delta_1}$ be the branch. Suppose the image of Δ_1 under ζ_1 contains another component Δ_2, with branch ζ_2. If h is a branchwise equivalence between ϕ and $\hat{\phi}$, there exists corresponding components $\hat{\Delta}_1 = h(\Delta_1)$ and $\hat{\Delta}_2 = h(\Delta_2)$, with branches $\hat{\zeta}_1$ and $\hat{\zeta}_2$ respectively. Furthermore, since h satisfies the functional equation on the boundary of Δ_1, the images of Δ_1 and $\hat{\Delta}_1$ correspond under h, and they contain Δ_2 and $\hat{\Delta}_2$ respectively.

A new box mapping is created from ϕ by replacing ζ_1 on Δ_1 with $\zeta_2 \circ \zeta_1$. Similarly, a new box mapping is created from $\hat{\phi}$ by replacing $\hat{\zeta}_1$ with

$\hat{\zeta}_2 \circ \hat{\zeta}_1$. The method for creating a new branchwise equivalence between the new box mappings is to redefine h on Δ_1 with

$$h(z) = \hat{\zeta}_1^{-1} \circ h \circ \zeta_1(z).$$

Inspection of this equation and the definition of branchwise equivalences reveals that this is the *only* possible definition of the branchwise equivalence on the boundary of Δ_2. Recall that the similarity class of a branchwise equivalence is determined by how it maps the boundaries of domain components, so we have the following fact.

Fact. The pull-back construction operates on similarity classes. Pull-back of any branchwise equivalence in a class, since the mapping is determined on the boundary of domain components, always produces another branchwise equivalence in a similarity class independent of the map chosen within the original similarity class.

A complete description of the pull-back operation, but without the complication of a folding branch, is contained in the proof of Lemma [7.1].

Critical Consistence and Inducing in Parallel. Referring back to the previous example, the redefinition of h is well-defined if ζ_1 and $\hat{\zeta}_1$ are injective but involves a two to one topological lifting if ζ_1 and $\hat{\zeta}_1$ are central branches. There is an extra requirement on the branchwise equivalence for this lifting to exist. The critical value of ζ_1, $\zeta_1(0)$ must be mapped to the critical value of $\hat{\zeta}_1$, $\hat{\zeta}_1(0)$. This is the basic idea behind *critical consistence.*

Definition. (critically consistent branchwise equivalences) A branchwise equivalence is **critically consistent** if some member of its similarity class maps the critical value of one box mapping, $\phi_i(0)$, to the critical value of the other box mapping, $\hat{\phi}(0)$.

It turns that critical consistence is never an issue in the construction of branchwise equivalences, in situations where the box mappings are induced from polynomials which are topologically conjugate on the real line. It is a fundamental principal of the inducing algorithm that the construction is topologically invariant.

Fact. Let f and $\hat{f}$ be two real polynomials that are topologically conjugate on the real line by a conjugacy H. Suppose $\{\phi_i\}$ is a series of box mappings induced from f constructed from ϕ_0 using the inducing algorithm. Then box mappings, $\{\hat{\phi}_i\}$, induced from $\hat{f}$ using the inducing algorithm, will be constructed *in parallel* with the $\{\phi_i\}$. In other words, at any point of the construction, H will form a branchwise equivalence between ϕ_i and $\hat{\phi}_i$.

For example, if ϕ is any box mapping induced from f, by virtue of the functional equation $H \circ f^{\circ j} = \hat{f}^{\circ j} \circ H$, which holds for any j, the map $\hat{\phi} = H \circ \phi \circ H^{-1}$ is a box mapping induced from $\hat{f}$. Clearly H is branchwise equivalence between ϕ and $\hat{\phi}$. Furthermore, performing a pull-back operation, by replacing H with $(\hat{f}^{\circ j})^{-1} \circ H \circ f^{\circ j}$, where $\hat{f}^{\circ j}$ and $f^{\circ j}$ are branches of ϕ and $\hat{\phi}$, is a null operation, merely reproducing H.

There is a corollary to this principle which states that critical consistence is not an issue.

Corollary. Let f and $\hat{f}$ be two real polynomials that are topologically conjugate on the real line, by a conjugacy H. Let ϕ and $\hat{\phi}$ be box mappings induced, *in parallel*, from f and $\hat{f}$, and suppose $[h]$ is the similarity class of branchwise equivalences which contains H.

Then any h in $[h]$ is critically consistent. Performing an inducing step on ϕ and $\hat{\phi}$, then pulling back h, produces a branchwise equivalence which is critically consistent and is in the same similarity class as H, the full conjugacy.

PROOF. Let ψ and $\hat{\psi}$ be the central branches of ϕ and $\hat{\phi}$, which must equal $f^{\circ j}$ and $\hat{f}^{\circ j}$ respectively for some j. Since H conjugates f and $\hat{f}$, we must have $H \circ f^{\circ j}(0) = \hat{f}^{\circ j} \circ H(0)$, and $H(0) = 0$. Therefore H maps the critical value of ψ to the critical value of $\hat{\psi}$, and so by definition, any h in $[h]$ must be critically consistent. The pull-back of the conjugacy H is just H. So the pull-back, considered as an operation on similarity class, produces a similarity class containing H. Q.E.D.

§6.4. Box mapping results.

We can now state the central results for box mappings and branchwise equivalences. For an annulus A in $\mathbb{C}$, let $\operatorname{mod}(A)$ be the modulus of A.

Proposition 6.1. Let ϕ and $\hat{\phi}$ be type I holomorphic box mappings. Suppose that H is a K-quasiconformal branchwise equivalence between ϕ and $\hat{\phi}$. Assume that the similarity class $[H]$ is critically consistent for each simple inducing step composing a type I inducing step and then again for one more simple inducing step. Assume further that whenever D is the domain of a branch of $\hat{\phi}$ and $D \subset \hat{B}'$, then $\bmod(\hat{B}' - D) \geq \epsilon$.

Then, for every $\epsilon > 0$ there is a Q so that the similarity class obtained from $[H]$ by the pull-back associated to a type I inducing step has a QK-quasiconformal representative. Moreover, if $\epsilon \geq 4\log 8$, then one can take $Q = \exp(Q' \exp(-\frac{\epsilon}{4}))$ where Q' is a constant, independent of all parameters.

Theorem 6.2. Let ϕ be a type II holomorphic box mapping, ϕ_0 be the type I mapping obtained from ϕ by filling-in, and ϕ_i a sequence, finite or not, of holomorphic box mappings set up so that ϕ_{i+1} is derived from ϕ_i by the type I inducing step for $i \geq 0$. Suppose that ϕ restricted to the real line is a real type II box mapping. Let B_i and B_i' denote the central domain and central range of each ϕ_i.

Suppose that $\bmod(B_0' - B_0) \geq \beta_0$. For every $\beta_0 > 0$ there is a number $C > 0$ with the property that for every i

$$\bmod(B_i' - B_i) \geq C \cdot i.$$

For complete proofs see [GS1] (also see [Sw] and [JS2]).

§6.5. Generalized box mappings.

For purposes of applying box mappings to cubic polynomials, we make two extensions to the family of holomorphic box mappings as originally defined. The second extension is really a separate family but is similar enough to still be called a box mapping. Both extensions are concerned with the central branch. For the most part, the term *box mapping* will refer to this new extended family. Where there is a possible source of confusion, we will use the words *standard* and *general* to refer to box mappings satisfying the original or extended definition.

Quadratic symmetry. The first extension eliminates the need for quadratic symmetry in a box mapping. By quadratic symmetry, we mean $f(a) = f(-a)$, as is the case with the quadratic family $x^2 + c$.

The original definitions of real and holomorphic box mappings require that the branch of the map defined on the central domain be *symmetric*, in the above sense (properties (1) and (3) of either definition respectively). This is a natural property which is easy to satisfy when inducing box mappings from *quadratic* maps. But in fact, it is not a necessary assumption, and we really only need the fact that the central branch is a conformal two-to-one covering map. To be specific we have the following more general definition.

Definition. (general holomorphic box mappings) The definition is identical to that of standard holomorphic box mappings, except properties (1) and (3) of the original definition are replaced with the following requirement.

There exists a set B, the central domain, that is an open disk and a component of the domain of Φ, U. B contains precisely one critical point of Φ, c, and Φ restricted to $B - \{c\}$ is a two-to-one covering map onto its image.

Box maps are a tool for building conjugacies between maps, and it is for these applications that the symmetry is not necessary because we have the following lemma.

Lemma 6.3. Suppose Φ is an extended holomorphic box mapping, whose domain is completely contained in some topological disk V. Suppose Φ is induced from a single real analytic map f, i.e. each branch of Φ is of the form, $f^{\circ k}$. Finally suppose that V contains one and only one critical point of f, c, which must be the same unique critical point of Φ in B, and that f restricted to $V - \{c\}$ is a two-to-one covering map onto its image.

Then Φ is conformally conjugate, by a map defined on V, to a standard holomorphic box mapping.

PROOF. Define two new maps, $F(z) = f(z) - f(c)$ and $Q(z) = z^2$. Also define $\hat{V} = Q^{-1}(F(V))$. Note that both F and Q are two-to-one covers, from $V - \{c\}$ and $\hat{V} - \{0\}$ respectively, onto $F(V) - \{0\} = Q(\hat{V}) - \{0\}$. We are in a simple lifting situation. Thus we can find a lift of F through Q that is conformal in this case, symmetric, preserves the orientation of $\mathbb{R}$, and maps V univalently onto $\hat{V}$. We label this lift H. We of course have the functional equation $Q \circ H = F$. The map H is our desired conjugacy. We define $\hat{\Phi} = H \circ \Phi \circ H^{-1}$ on domain $\hat{U} = H(U)$, where U is the domain of Φ, and we claim that $\hat{\Phi}$ is a standard holomorphic box mapping.

Suppose the central branch $\psi = \Phi|_B = f^{\circ k}$, where k is some positive integer. The central domain of $\hat{\Phi}$, $\hat{B} = H(B)$, is an open disk and a connected component of $\hat{U}$. Suppose $z \in \hat{B} - \{H(c)\}$, then $H^{-1}(z)$ is in B, and by definition of general box mappings, there is a second point in B, y, with the same image as $H^{-1}(z)$ under $\psi = f^{\circ k}$. So we have $\psi(H^{-1}(z)) = \psi(y)$.

We want to show that the orbits of y and $H^{-1}(z)$ are equal after the first iteration of f. Consider f as a map defined on $\mathbb{R}$ for a moment. The set B contains one critical point of f, which is the same as the critical point of ψ. So $f^{\circ(k-1)}$ must be monotone on the set $f(B)$. In other words, if $f(y) \neq f(H^{-1}(z))$ then $f^{\circ k}(y) \neq f^{\circ k}(H^{-1}(z))$, which we know are equal. So we must have that $y \in B$ and

$$f(y) = f(H^{-1}(z)). \tag{1}$$

Rearranging the functional equation, we get $F(H^{-1}(z)) = Q(z)$. From this equation, equation (1), and the definition of F, we get $F(y) = Q(z)$, which implies $Q(H(y)) = Q(z)$, or $[H(y)]^2 = z^2$. By definition, $y \neq H^{-1}(z)$ or $H(y) \neq z$, so we must have $H(y) = -z$. However since $y \in B$, $H(y) \in \hat{B}$, and we have shown that the map $z \mapsto -z$ maps $\hat{B}$ onto itself, property (1) of holomorphic box mappings.

By definition, the central branch of $\hat{\Phi}$ is $\hat{\psi} = H \circ \psi \circ H^{-1}$, which is equal to $H \circ f^{\circ k} \circ H^{-1}$. From the functional equation and the definitions of F and Q notice,

$$f \circ H^{-1}(z) = z^2 + f(c).$$

The map $f \circ H^{-1}$ takes $\hat{B}$ to $f(B)$, and $f^{\circ(k-1)}$ is univalent on $f(B)$. Thus

$$h(z) = H \circ f^{\circ(k-1)}(z + f(c))$$

is a diffeomorphism on $f(B) - f(c) = F(B)$. For $z \in \hat{B}$, which is equal to $H(B) = Q^{-1} \circ F(B)$, we get

$$h(z^2) = H \circ f^{\circ(k-1)} \circ f \circ H^{-1}(z) = \hat{\psi},$$

as required for property (3).

All other properties of holomorphic box mappings follow immediately from the fact that Φ is conjugate to $\hat{\Phi}$ by a conformal map. Q.E.D.

Box mappings with no central branch. We can extend the family of holomorphic box mappings further by allowing the central branch to be optional. We allow box mappings *with no central branch.* Such a function only needs to satisfy properties (4) and (6) in the definition of standard holomorphic box mappings. Note that the definition of *branchwise equivalences* does not need to be modified in any way to be used with this new type of box mapping. Use of just the term *box mapping* will refer to a box mapping with a proper central branch. A box mapping that does not have a central branch will always be referred to as a *box mapping with no central branch.*

VII. Polynomial tools

We need some additional tools before attempting to construct box mappings induced from cubic maps. First, we present a detailed pull-back construction in the special case of box mappings with no central branch. This result will be used as a bridge between the Theorems [6.2] and [6.1], cubic maps, and quasiconformal conjugacies between cubic maps.

The λ-Lemma, presented in Chapter VII part 3, is the primary tool used to construct initial quasiconformal branchwise equivalences, which are converted by pull-back into full quasiconformal conjugacies. The hypotheses of the λ-Lemma will require a closer examination of the Böttcher coordinate. The coordinate needs to be defined as a univalent map, even in situations where the Julia set of the underlying map is disconnected.

§7.1. Monotone pull-back.

We present a more rigorous definition of the pull-back construction in an environment without critical points, specifically on box mappings with no central branch. The induction presented here will be useful in two situations. First, by simply ignoring the critical points of two cubic polynomials, objects can be constructed using monotone pull-back that look *almost* like type II box mappings induced from the polynomials and a branchwise equivalence between them. The only pieces missing are the two central folding branches of the box mappings which can be added at the end. In this way an initial branchwise equivalence and box mappings can be constructed, which will then fall immediately into the inducing algorithm described above. This is described in Chapter VIII. If a non-recurrent or renormalizable situation is not encountered, the central branches of the box mappings will shrink to a point during the inducing algorithm, producing, in the limit, another pair of box mappings with no central branches. The limit of the quasiconformal branchwise equivalences will be a partial conjugacy that conjugates part of the forward orbits of the chaotic critical points of the original cubic polynomials. In order to apply Theorem [5.4] and obtain a full conjugacy, the partial conjugacy must conjugate the entire forward orbit of the critical point, and to achieve this condition, an additional application of the monotone pull-back will be required. This is described in Chapter IX.

Monotone pull-back will be performed exclusively on holomorphic box mappings, with no central branch.

Lemma 7.1. Let ϕ_0 and ϕ be a pair of holomorphic box mappings with no central branches, which are symmetric with respect to $\mathbb{R}$. Let $\hat{\phi}_0$ and $\hat{\phi}$ be another such pair. Suppose there exists a map h that is a branchwise equivalence between ϕ_0 and $\hat{\phi}_0$ and also between ϕ and $\hat{\phi}$. Suppose h is K-quasiconformal, symmetric, and preserves the orientation of $\mathbb{R}$.

Let a sequence of maps $\{\phi_i\}$ and $\{\hat{\phi}_i\}$ be defined inductively as follows, with $\phi_1 = \phi$, $\hat{\phi}_1 = \hat{\phi}$, and $h_1 = h$.

$$\phi_{i+1}(z) = \phi_0 \circ \phi_i(z).$$

$$\hat{\phi}_{i+1}(z) = \hat{\phi}_0 \circ \hat{\phi}_i(z).$$

Then each pair $(\phi_{i+1}, \hat{\phi}_{i+1})$ are symmetric box mappings with no central branches, and there exists a branchwise equivalence, h_{i+1}, between them, which is K-quasiconformal, symmetric, and preserves the orientation of $\mathbb{R}$.

PROOF. By induction we assume ϕ_i, $\hat{\phi}_i$, and h_i already exist with the desired properties. Throughout the proof we will use the following notation.

Let Δ_i be any connected component of the domain of ϕ_i, and let $\hat{\Delta}_i$ be the corresponding component $h_i(\Delta_i)$ of $\hat{\phi}_i$. Let V_i and $\hat{V}_i$ be the images of Δ_i and $\hat{\Delta}_i$ under ϕ_i and $\hat{\phi}_i$ respectively. Let ζ_i and $\hat{\zeta}_i$ be the branches of ϕ_i and $\hat{\phi}_i$ restricted to Δ_i and $\hat{\Delta}_i$. We will also use Δ_0, V_0, and ζ_0 to refer to domains, images, and branches of ϕ_0. Note that the component of $\hat{\phi}_0$ corresponding to Δ_0 is $\hat{\Delta}_0 = h(\Delta_0)$.

We also make the following additional induction hypotheses.

(1) Assume any domain component of ϕ_0, Δ_0, is either contained in or disjoint from any V_i. Similarly, any $\hat{\Delta}_0$ is either contained in or disjoint from $\hat{V}_i$.

(2) Any Δ_i is contained in some Δ_0, and similarly $\hat{\Delta}_i$ is contained in some $\hat{\Delta}_0$.

(3) For each V_i, $\partial V_i \subset \partial V_0$ for some V_0, and $\partial \hat{V}_i \subset \partial \hat{V}_0$.

(4) For any z in ∂V_0, $h_i(z) = h(z)$.

We of course must verify these for the first induction step. With $i = 1$, since ϕ and ϕ_0 must have the same domain components, (1) amounts to the Markov property required of all box mappings. Property (2) and (4) are trivially true. Property (3) may not be true for the first step, $i = 1$, but we will show that it is true for every step after the first. At any rate, we claim for $z \in \partial V_i$, $h_i(z) = h(z)$. This clearly follows from (3) and (4) for $i > 1$, and for $i = 1$, the claim is trivial.

Define h_{i+1} as follows.

$$h_{i+1}(z) = \begin{cases} h_i(z), & \text{for } z \notin \Delta_i \text{ a component of } \phi_i \\ \hat{\zeta}_i^{-1} \circ h \circ \zeta_i(z), & \text{for } z \in \Delta_i. \end{cases}$$

Fix a domain component of ϕ_0, Δ_0. For arbitrary V_i, either Δ_0 is contained in it or is totally disjoint by (1). Notice that any domain component of ϕ_{i+1} arises when $\Delta_0 \subset V_i$, in which case ϕ_{i+1} has component $\Delta_{i+1} = \zeta_i^{-1}(\Delta_0)$. Let us fix a Δ_i and V_i where this is the case. Since h_i is a branchwise equivalence

$$\hat{\zeta}_i^{-1} \circ h_i \circ \zeta_i(z) = h_i(z) \qquad \text{for } z \in \partial\Delta_i.$$

This also implies that $h_i(\partial V_i) = \partial \hat{V}_i$, and from (4) and the claim stated earlier, this gives $h(\partial V_i) = \partial \hat{V}_i$. Since Δ_0 is contained in V_i, this shows that $\hat{\Delta}_0$ is contained in $\hat{V}_i$. The set $\hat{\Delta}_{i+1} = \hat{\zeta}_i^{-1}(\hat{\Delta}_0)$ is now defined as a component of $\hat{\phi}_{i+1}$, which we will show naturally corresponds to Δ_{i+1}. We have $h_{i+1}(\Delta_{i+1}) = h_{i+1}(\zeta_i^{-1}(\Delta_0))$. By definition of h_{i+1} on Δ_i, we get $h_{i+1}(\zeta_i^{-1}(\Delta_0)) = \hat{\zeta}_i^{-1}(h(\Delta_0)) = \hat{\zeta}_i^{-1}(\hat{\Delta}_0)$. Putting this all together gives $h_{i+1}(\Delta_{i+1}) = \hat{\Delta}_{i+1}$ as expected.

Any connected component of the domain of ϕ_{i+1} will be a Δ_{i+1} as constructed above, and ϕ_{i+1} restricted to Δ_{i+1} is equal to $\zeta_0 \circ \zeta_i$, the composition of univalent and symmetric maps. Therefore ϕ_{i+1} is a symmetric map whose branches are univalent. Let V_{i+1} be the image of Δ_{i+1} under $\zeta_{i+1} = \zeta_0 \circ \zeta_i$.

Thus $V_{i+1} = \zeta_0 \circ \zeta_i(\Delta_{i+1}) = \zeta_0(\Delta_0) = V_0$, and similarly, $\hat{V}_{i+1} = \hat{V}_0$. Note this verifies property (3) for step $i+1$. Let Δ'_{i+1} be any other component of ϕ_{i+1}, which must be contained in some Δ'_i by construction. By (2), Δ'_i is contained in some Δ'_0. But ϕ_0 is a box mapping, so V_0, the image of Δ_0, is disjoint from or contains Δ'_0. Therefore Δ'_{i+1} is disjoint from or contained V_{i+1}. This finishes the proof that ϕ_{i+1} is a symmetric box mapping with no central branch. A similar proof shows $\hat{\phi}_{i+1}$ is a box mapping. Along the way we have also verified property (1) for step $i+1$.

Turning to the map h_{i+1}, we show first that h_{i+1} is a homeomorphism. Keeping the same notation, let us examine how h_{i+1} is redefined on Δ_i. Since h_i is a branchwise equivalence, for $z \in \partial\Delta_i$, we have $h_i(z) = \hat{\zeta}_i^{-1} \circ h_i \circ \zeta_i(z)$. As shown before, $h_i = h$ on the boundary of V_i, so this equation becomes

$$h_i(z) = \hat{\zeta}_i^{-1} \circ h \circ \zeta_i(z) \qquad \text{for } z \in \partial\Delta_i.$$

This matches the definition of h_{i+1} on Δ_i. Therefore the two pieces glued together to form h_{i+1} agree on the boundary of Δ_i, so h_{i+1} must be a continuous map. Furthermore, h_{i+1} takes Δ_i to $\hat{\Delta}_i$ and the complement of Δ_i to the complement of $\hat{\Delta}_i$. Since the pieces of h_{i+1} are both injective, we have shown that h_{i+1} is a homeomorphism.

It follows immediately that h_{i+1} is K-quasiconformal because, for all points in $\mathbb{C}$, h_{i+1} either equals h_i or is a composition by conformal maps with h. Both h and h_i are K-quasiconformal. Also h_{i+1} is either equal to h_i, which is symmetric, or it is redefined on a domain (which must be symmetric with respect to $\mathbb{R}$) as the composition of symmetric maps. Therefore h_{i+1} is symmetric. Symmetric homeomorphisms, can be either orientation preserving or reversing, and since h_{i+1} is partly equal to h_i, which is orientation preserving, h_{i+1} must be orientation preserving as well.

By construction, each Δ_{i+1} is contained in some Δ_i, which by (2) for step i is contained in a component Δ_0. Similarly $\hat{\Delta}_{i+1} \subset \hat{\Delta}_0$. This verifies (2) for step $i+1$. To verify (4) for step $i+1$, let z be a point in ∂V_0. We claim z is not contained in any Δ_{i+1}, or else by (2), z would be contained in some Δ_0. The boundary of V_0 would intersect Δ_0, which cannot happen if ϕ_0 is a box mapping. So by definition, $h_{i+1}(z) = h_i(z)$. Therefore by (4) for step i, $h_{i+1}(z) = h_i(z) = h(z)$.

All we have left to verify is that h_{i+1} is a branchwise equivalence. We have already shown $h_{i+1}(\Delta_{i+1}) = \hat{\Delta}_{i+1}$, so h_{i+1} maps domain components onto domain components. For z in $\partial\Delta_{i+1}$, we have by previous analysis shown

$$\hat{\zeta}_{i+1}^{-1} \circ h_{i+1} \circ \zeta_{i+1}(z) = \hat{\zeta}_i^{-1} \circ \hat{\zeta}_0^{-1} \circ h \circ \zeta_0 \circ \zeta_i(z).$$

We know for $z \in \partial\Delta_{i+1}$, $\zeta_i(z) \in \partial\Delta_0$. Therefore since h is a branchwise equivalence, $h(z) = \hat{\zeta}_0^{-1} \circ h \circ \zeta_0(z)$, and we get

$$\hat{\zeta}_{i+1}^{-1} \circ h_{i+1} \circ \zeta_{i+1}(z) = \hat{\zeta}_i^{-1} \circ h \circ \zeta_i(z).$$

Notice $z \in \partial\Delta_{i+1}$ is contained in some Δ_i, so the right hand side is the definition of $h_{i+1}(z)$. Thus we get

$$\hat{\zeta}_{i+1}^{-1} \circ h_{i+1} \circ \zeta_{i+1}(z) = h_{i+1}(z) \qquad \text{for } z \in \partial\Delta_{i+1},$$

verifying that h_{i+1} is a branchwise equivalence. Q.E.D.

Monotone pull-back is really just a filling in step (Step A). It can be used to build a first return map induced from ϕ by redefining ϕ_0 from ϕ by replacing ϕ with the identity map on the desired inducing interval.

§7.2. Extending the Böttcher coordinate.

Traditionally rays and potentials are defined for polynomials with a connected Julia set, where it is well known that the Böttcher coordinate extends as a univalent function throughout the entire basin at infinity. Unfortunately we will also be trying to use rays and potentials in situations where the Julia set may not be connected. This will be necessary for the application of the λ-Lemma. Rays and potentials must be defined with a little more care, or to be precise, the domain of the Böttcher coordinate must be explicitly constructed because the coordinate can no longer be extended to the entire basin. The extended domain must be large enough to contain all the rays we will need but must still allow the coordinate to be a well-defined and univalent function.

Domains of the Böttcher coordinate. The Böttcher coordinate is always defined for any polynomial in a neighborhood of infinity. To fix our notation, let f be a polynomial of degree n, and let φ be the univalent coordinate function defined on U, the neighborhood of infinity. We can normalize U somewhat by requiring that its image under φ be a disk in $\hat{\mathbb{C}}$ centered at infinity, i.e. its boundary has constant radius in $\mathbb{C}$. This can be accomplished because the image of an arbitrary U must contain such disks. Also define $Z(z) = z^n$.

Suppose z is any point in the basin at infinity for f. Since U is a neighborhood of infinity, if we iterate f enough, the orbit of z must enter U. We then have $f^{\circ j}(z) \in U$. Any extension of φ must still satisfy the functional equation $\varphi \circ f = Z \circ \varphi$. By repeatedly applying this and then rearranging we obtain

$$\varphi(z) = (Z^{\circ j})^{-1} \circ \varphi \circ f^{\circ j}(z).$$

By choosing appropriate inverse branches of Z, we have a suitable local definition of φ at z in terms of the original φ defined on U. From this we see that the only obstructions to extending φ locally as a univalent map are the critical and precritical points of f. This next lemma starts to address the problems of defining φ globally.

Lemma 7.2. Let φ be a Böttcher coordinate, for the polynomial f of degree n, defined on the set U, a neighborhood of infinity. Then φ can be extended as a well-defined analytic function to any path connected set V containing U, provided V is simply connected, does not contain any critical or precritical points of f, and is contained in the basin at infinity of f. Furthermore, the extended φ is locally one to one.

PROOF. We show that φ can be analytically continued to any point in V from U, and then the Monodromy Theorem guarantees that a well-defined extension of φ to the simply connected V exists.

Let x be a point in $V - \{\infty\}$. Since V is path connected and avoids precritical points we can construct a closed path segment in V with one endpoint at x and the other endpoint in U. Since the path segment is compact, we can cover it with a finite number of open sets A_1 through A_m, where we require for each $1 \leq i \leq m$

(1) There exists an integer j so that $f^{\circ j}(A_i) \subset U$

(2) $f^{\circ j}$ is univalent on A_i

(3) $f^{\circ j}(A_i)$ is simply connected and does not contain infinity

Property (1) can be satisfied because V is contained in the attracting basin of infinity. Property (2) can be satisfied because a critical point of $f^{\circ j}$ is a critical or precritical point of f, avoided by V. Since $f^{\circ j}$ is not critical at x, it looks locally like a univalent map, so we can satisfy property (3) by taking the A_i to be very small.

Inductively we extend φ to each of the A_i one by one, starting with a set that intersects U. Suppose A_1 intersects U. As noted before, we can use the equation

$$\varphi(z) = (Z^{\circ j})^{-1} \circ \varphi \circ f^{\circ j}(z) \qquad \text{for } z \in A_1 \cap U. \tag{1}$$

as a local definition of φ. Define $B = f^{\circ j}(A_1)$ and $\hat{B} = \varphi(B)$. Since B and thus $\hat{B}$ do not contain infinity and are simply connected, the inverse image of $\hat{B}$ through $Z^{\circ j}$ consists of exactly n^j connected components. Let y be some point in $A_1 \cap U$, which by equation (1) must be mapped by φ into one of these n^j components. We label this component $\hat{A}_1$. Clearly $Z^{\circ j}$ maps $\hat{A}_1$

univalently onto $\hat{B}$. We have a lifting situation where A_1 and $\hat{A}_1$ are simply connected covering spaces, through $f^{\circ j}$ and $Z^{\circ j}$ respectively, of B and $\hat{B}$. The lifting criteria for φ is trivially satisfied as $(Z^{\circ j})^{-1} : \hat{B} \to \hat{A}_1$ is well defined. Furthermore a base point has been established by the requirement that $\varphi(y)$ is in $\hat{A}_1$, so φ can be lifted to a univalent function L on A_1, satisfying $L \circ f^{\circ j} = Z^{\circ j} \circ L$ and taking y into $\hat{A}_1$. Since φ itself satisfies this equation on $A_1 \cap U$, by the uniqueness of the lifting, this new map L must be an extension of φ.

We extend φ to x by using the same argument to extend φ to each of the A_i, eventually connecting x to U. Note that the extended φ looks locally like the composition of three univalent functions and is thus locally one to one. Q.E.D.

The definition of a ray reflects the idea that a ray should always be a path with one endpoint at infinity. We can extend this idea to the Böttcher coordinate itself. Intuitively, we associate the fact that φ is well-defined and univalent on its initial domain U with the fact that U contains no *partial* rays, rays that do not extend to infinity. To extend φ, we analytically continue it along each ray, and stop if it hits a precritical point of f. The resulting domain should be simply connected, as we can connect any two points by traveling up one ray and down another, and by Lemma [7.2] the extended φ would be well-defined. Also the extended φ should be univalent because we have carefully avoided critical points. We encompass the idea of domains containing no partial rays in the *radially convex* property.

Definition. (radially convex sets) Let D be the unit disk. A set $V \subset \hat{\mathbb{C}} - \bar{D}$

is **radially convex**, if V contains infinity, and for every $z \in V - \{\infty\}$ and $r > 1$, rz is also contained in V. Note a radially convex set must be simply connected, and the union of radially convex sets is always a radially convex set.

Now we define the corresponding concepts in the dynamical plane of a specific f a coordinate φ, and its original domain U.

Definition. (ray convex sets) Suppose the set $O \subset \hat{\mathbb{C}}$ is simply connected, contains U, and contains no critical or precritical points of f. There is a well-defined extension of φ to O, by Lemma [7.2]. We will call the set O **ray convex** if it satisfies all these properties and its image under the extended φ is a radially convex set.

We claim that φ can always be extended as a well-defined and univalent map to any domain constructed as the union of ray convex sets, which then must itself be ray convex. To construct a proof we will need maximal versions of ray convex sets.

Definition. (maximal ray convex sets) Fix a radius $r \geq 1$. Suppose that $\{O_\alpha\}_{\alpha \in \mathcal{A}}$ is the collection of all possible ray convex sets, and $\{\varphi_\alpha\}$ the corresponding extensions of φ, such that $\varphi_\alpha(O_\alpha)$ does not intersect the closed disk of radius r in $\mathbb{C}$ centered at the origin. Then define the set

$$M_r = \bigcup_{\alpha \in \mathcal{A}} O_\alpha.$$

For $s < r$, we clearly have $M_r \subset M_s$. We know that at least one ray convex set, U itself, exists. So for r small enough the set M_r is non-empty.

The coordinate φ can certainly be locally extended to every point in M_r, but we cannot say yet that φ is well-defined on M_r. We can say something about the existence of an inverse map.

Lemma 7.3. Let $\{O_\alpha\}_{\alpha \in \mathcal{A}}$ be a collection of simply connected sets in $\hat{\mathbb{C}}$, none of which contain a critical or precritical point of the polynomial f, and all of which contain a neighborhood of infinity U on which a univalent Böttcher coordinate φ is defined for f. By Lemma [7.2], we can extend φ onto each O_α, as φ_α, a well-defined analytic function. Let $V = \bigcup_{\alpha \in \mathcal{A}} \varphi_\alpha(O_\alpha)$. If V is simply connected, then there exists a single analytic map φ^{-1} defined on V satisfying $\varphi^{-1} \circ \varphi_\alpha(z) = z$ for all $z \in O_\alpha$ and $\alpha \in \mathcal{A}$.

PROOF. Since φ is univalent on U, the map φ^{-1} is defined on $\varphi(U)$. Suppose that we could show that, from $\varphi(U)$, φ^{-1} can be extended to V. It is then easy to check that φ^{-1} is actually the inverse of every φ_α because every φ_α is extended from φ, of which φ^{-1} is the inverse by definition. Then since O_α is connected, the functional equation $\varphi^{-1} \circ \varphi_\alpha(z) = z$ must hold throughout O_α. In other words, $\varphi^{-1} \circ \varphi_\alpha(z) - z$ is analytic and identically zero on U, so it must be on O_α.

To show that φ^{-1} can be extended to V we first show that φ^{-1} can be analytically continued throughout V. To show this, let x be in V, and let $\varphi_\alpha(O_\alpha)$ be a set containing x. The set $\varphi_\alpha(O_\alpha)$ is connected and contains $\varphi(U)$. By Lemma [7.2], φ_α is locally one to one on O_α. Therefore φ^{-1} will be locally well-defined if it is analytically continued from $\varphi(U)$ through $\varphi_\alpha(O_\alpha)$ to x. The Monodromy Theorem says that φ^{-1} can be extended as a single valued analytic function throughout the simply connected V. Q.E.D.

Using this well-defined inverse, we can show that φ can be extended as a well-defined function to M_r for any $r \geq 1$.

Proposition 7.4. For $r \geq 1$, the Böttcher coordinate φ of the polynomial f, of degree n, can be extended as a univalent function to the maximal ray convex set M_r.

PROOF. By applying Lemma [7.3] to the collection of sets $\{O_\alpha\}$ used to build M_r, there exists a single inverse map φ^{-1} for any extension of φ to a subset of M_r. So all we have to do to prove the lemma is show that a well-defined extension of φ exists. Such an extension has to be univalent because its inverse map exists.

If r_0 is the constant radius of the boundary of $\varphi(U)$, then clearly the maximal ray convex set M_{r_0} is equal to U, and φ is well-defined on M_r, for any $r \geq r_0$. Let b be the greatest lower bound of all numbers x such that $r > x$ implies φ is well-defined when analytically continued to M_r. Our goal then is show that $b = 1$. Suppose $b > 1$. Then we can find an $s > 1$ such that $\sqrt[n]{b} < s \leq b$ and φ is not well-defined when continued to M_s.

If φ is not well-defined, we can find two points that φ^{-1}, which is well-defined, maps to the same point in M_s. In other words, we can find x, y, and z such that $\varphi^{-1}(x) = \varphi^{-1}(y) = z \in M_s$. Let O_x be a ray convex set containing U whose image under φ_x, the extension of φ to O_x, contains x. Similarly we define O_y and φ_y for the point y. These sets and extensions must exist by the definition of M_r. Since $\varphi_x(O_x)$ is radially convex and contains x, it must contain the radial line of angle equal to the argument of x down to the radius of $|x|$. If α is the argument of x. This line is denoted

$\hat{\mathcal{R}}_{\alpha}^{|x|}$, and let $\mathcal{R}_{\alpha}^{|x|} \subset O_x$ be its image under φ^{-1}. Similarly define the radial line hitting y, $\hat{\mathcal{R}}_{\beta}^{|y|}$, and its image $\mathcal{R}_{\beta}^{|y|} \subset O_y$.

We now examine the two sets $f(\mathcal{R}_{\alpha}^{|x|})$ and $f(\mathcal{R}_{\beta}^{|x|})$. Intuitively, these must represent rays where φ is again not well-defined but, this time, in a maximal ray convex set of radius greater than b. We claim that $f(\mathcal{R}_{\alpha}^{|x|})$ and $f(\mathcal{R}_{\beta}^{|x|})$ are contained in M_{s^n}, to which φ extends as a well-defined map because $s^n > b$. The set $f(\mathcal{R}_{\alpha}^{|x|}) \cup U$ is simply connected and contains no critical or precritical points of f because neither U nor $\mathcal{R}_{\alpha}^{|x|}$ do. Thus φ can be extended to this set by Lemma [7.2], and from the functional equation of f and φ, we get

$$\varphi(f(\mathcal{R}_{\alpha}^{|x|})) = Z \circ \varphi_x(\mathcal{R}_{\alpha}^{|x|}).$$

Since φ^{-1} is the inverse of φ_x, $\varphi_x(\mathcal{R}_{\alpha}^{|x|})$ is the radial line $\hat{\mathcal{R}}_{\alpha}^{|x|}$ and applying $Z(z) = z^n$ to it gives us another radial line that extends down to the point $Z(x)$. We have now shown that $\varphi(f(\mathcal{R}_{\alpha}^{|x|}) \cup U)$ is radially convex. Notice $|Z(x)| > s^n$, so $Z(\hat{\mathcal{R}}_{\alpha}^{|x|})$ is outside the closed disk of radius s^n. Therefore $f(\mathcal{R}_{\alpha}^{|x|})$ is in M_{s^n}. The same proof also works for $f(\mathcal{R}_{\beta}^{|y|})$.

But $f(\mathcal{R}_{\alpha}^{|x|})$ and $f(\mathcal{R}_{\beta}^{|y|})$ intersect at $f(z)$. The continuation of φ to M_{s^n} is well-defined by definition of b, so φ^{-1} is univalent on $\varphi(M_{s^n})$. We must have $Z(x) = Z(y)$. This implies

$$Z \circ \varphi_x(\mathcal{R}_{\alpha}^{|x|}) = Z \circ \varphi_y(\mathcal{R}_{\beta}^{|y|}).$$

This in turn gives

$$\varphi(f(\mathcal{R}_{\alpha}^{|x|})) = \varphi(f(\mathcal{R}_{\beta}^{|y|})),$$

and finally

$$f(\mathcal{R}_\alpha^{|x|}) = f(\mathcal{R}_\alpha^{|y|}).$$

But $\mathcal{R}_\alpha^{|x|}$ and $\mathcal{R}_\beta^{|y|}$ were constructed to be disjoint sets. We must conclude that f is locally two or more to one near z, and z is a critical point of f. This is a contradiction because z is in M_s, which contains no critical points. Therefore b cannot be a lower bound, and φ can always be extended to any M_r for $r \geq 1$. Q.E.D.

Additional properties of rays. We present some additional properties of rays that are well-known in the case of a connected Julia set but now have a precise meaning in situations where the Julia set is disconnected.

Corollary. Suppose a Böttcher coordinate of the polynomial f, on any domain of infinity, is specified. Then the ray $\mathcal{R}_\theta^r$ is unique, and it is always defined unless there exists a smaller ray of the same angle, $\mathcal{R}_\theta^s$, for $s > r$, containing a critical or precritical point of f as a limit point.

PROOF. The extension of φ to M_1 is unique because any two such extensions are analytic and equal on an open set. So if $\mathcal{R}_\theta^r$ is defined at all it is unique. Let $\hat{\mathcal{R}}_\theta^r$ be a radial line in $\varphi(M_1)$, where r is chosen to be minimal. Therefore $\mathcal{R}_\theta^r$ is defined. If $r \neq 1$ then $U \cup \mathcal{R}_\theta^r$ has a radially convex image $\varphi(U) \cup \hat{\mathcal{R}}_\theta^r$. Look at φ restricted to $U \cup \mathcal{R}_\theta^r$. Suppose we can locally continue this restricted φ around the limit point of $\mathcal{R}_\theta^r$ in a univalent manner. Then M_1 would have to contain a neighborhood of this limit point, and the choice of r would not be minimal. Therefore φ cannot be extended in a one to one manner, so the limit point of $\mathcal{R}_\theta^r$ must be a critical or precritical point of f. Q.E.D.

Corollary. Any ray $\mathcal{R}_\theta^r$ is a connected set forming a simple path extending to infinity, and any two rays of different angles are disjoint.

PROOF. This follows immediately from the fact that rays are images, under a single diffeomorphism, of radial lines, which are disjoint from each other and extend to infinity. Q.E.D.

Corollary. If the basin of infinity for the polynomial f contains no critical point of f, the Böttcher coordinate φ extends as a well-defined and univalent function mapping the entire basin onto $\hat{\mathbb{C}} - \bar{D}$.

PROOF. If the basin contains no critical point, it certainly contains no precritical points of f. Thus there is no obstruction to analytically continuing φ throughout the entire basin, which is simply connected because its complement, the filled-in Julia set is connected. The range of the extended φ must, by the first corollary, be equal to all of $\hat{\mathbb{C}} - \bar{D}$. Thus the maximal ray convex set M_1 must be equal to the basin, so φ can be extended univalently to it. Q.E.D.

Note that f acting on a ray corresponds to $Z(z) = z^n$ acting on a radial line. Therefore the action of f tends to map one ray onto another.

Lemma 7.5. Suppose $\mathcal{R}_\theta^r$ is an external ray of the polynomial f. Let A be either the forward image of the ray, $f(\mathcal{R}_\theta^r)$, or one of the connected components of the inverse image, $f^{-1}(\mathcal{R}_\theta^r)$, that does not contain a critical or precritical point of f. Then A is also an external ray of f.

PROOF. By the functional equation for φ and f, we have

$$\varphi \circ f(\mathcal{R}_\theta) = Z \circ \varphi(\mathcal{R}_\theta) = Z(\hat{\mathcal{R}}_\theta).$$

The image of a radial line under Z is always another radial line. Therefore the ray $\varphi^{-1}(Z(\hat{\mathcal{R}}_\theta))$, if it is defined at all, must be $f(\mathcal{R}_\theta)$. By the first corollary, we know that an entire ray is always defined unless part of it hits a critical or precritical point of f. But $f(\mathcal{R}_\theta)$ cannot hit such points because we know $\mathcal{R}_\theta$ does not hit such points. A similar proof works for $f^{-1}(\mathcal{R}_\theta)$ where again we know that it does not hit critical or precritical points. Q.E.D.

We can thus think of f as a map on rays. We can speak of a ray $\mathcal{R}_\theta$ as being periodic or having a period, if $f^{\circ k}(\mathcal{R}_\theta) = \mathcal{R}_\theta$ for some integer k, or fixed if $k = 1$. Notice that the points in a ray are parameterized by the corresponding radius of the potential curve they intersect. We say a ray $\mathcal{R}_\theta$ of f **lands** at x, the **landing point** of $\mathcal{R}_\theta$, if $\lim_{r \to 1^+} \varphi^{-1}(re^{i\theta})$ exists and is equal to x. Note that the orbit of such a limit point does not escape to infinity and thus must be a boundary point of the basin of infinity. Therefore the landing points of rays are always in the Julia set of f. The theory of when landing points exist still has many open questions, but we have the following standard result due to Douady and Yoccoz.

Theorem 7.6. Suppose f is a polynomial whose corresponding Julia set is connected, and suppose x is a repelling periodic point of f. Then there are at least one, but only a finite number, of rays landing at x, all of which have the same period.

We also have the converse result due to Sullivan, Douady, and Hubbard.

Theorem 7.7. Suppose f is a polynomial whose corresponding Julia set is connected. Then every periodic external ray lands at a periodic point which is either repelling or parabolic.

These facts can easily be extended to the corresponding preperiodic cases.

§7.3. Lambda lemma.

The primary tool of this section is an extension theorem due to Sullivan and Thurston, originally known as the λ-Lemma.

Theorem 7.8. Let $D \subset \hat{\mathbb{C}}$ be the open unit disk, and let E be any subset of the sphere containing at least four points. Suppose $i_\lambda : E \to \hat{\mathbb{C}}$ is a family of injections, for $\lambda \in D$, with i_0 equal to the identity map on E. Also suppose that for fixed $z \in E$, $i_\lambda(z)$ is holomorphic in λ (the family forms a holomorphic motion). Then for fixed $\lambda \in D$, i_λ can be extended to a quasiconformal homeomorphism of the $\hat{\mathbb{C}}$ onto itself.

The original λ-Lemma stated only that such motions could be extended to the closure of the set E (see [MSS]). A later version ([ST]) showed that the extension could be made to the entire Riemann sphere, and the form presented here is taken from [BR].

As this result, shows quasiconformal maps can be constructed out of analytically varying families of functions, like the Böttcher coordinate. By Theorem [2.3], rays and potentials vary analytically as the underlying map, f, is perturbed. Thus the Böttcher coordinate, or more precisely the inverse, restricted to pieces of radial lines and circles, forms a holomorphic motion. Branchwise equivalences are constructed by extending these motions using the λ-Lemma.

The Theorem is stated using families of one complex parameter, but the Böttcher coordinates of maps in $\mathcal{F}$ vary in two parameters. Thus for coordinates of cubic maps, we must use the λ-Lemma twice and compose the results, a minor technical problem. The injective requirement of the theorem is not a problem because two different rays are naturally disjoint objects. The Böttcher coordinate has been carefully defined so that it is well-defined, univalent (injective), and varies analytically in a neighborhood of the underlying map f, so long as we can show the ray avoids precritical points of f. This is the most difficult part, constructing a ray, for maps within an open set in parameter space, which still avoids precritical points when the Julia set becomes disconnected. The strategy here, starting with an existing ray, is to keep track of the critical points and the *end* of the ray, near its landing point, and show they stay apart if the perturbation is small enough.

Universal bounds. We start our study of holomorphic motions by showing that the potential curve G^r, for r larger than some constant, always exists in a neighborhood of parameter space, provided the neighborhood is bounded, as is the case with a bone-loop for instance. This will be of use immediately, when we construct the holomorphic motion on rays, and later, by providing a universal curve on which to build box mapping domains.

Lemma 7.9. Suppose $\{f_\lambda\}_{\lambda \in V}$ is a sub-family of $\mathcal{G}$, where $\mathcal{G}$ is parameterized by (A, B). Furthermore, suppose there exists constants $K_1, K_2 > 0$ so that for $\lambda = (A, B) \in V$

$$\frac{1}{K_1} \leq |A| \leq K_1 \quad \text{and} \quad |B| \leq K_2.$$

Then there exists a radius $R > 0$ so that for $\lambda \in V$

$$|f_\lambda(z)| > |z| \qquad \text{for } |z| > R.$$

PROOF. Choose $R = K_1(1 + K_2) + 3$. We certainly have the following inequalities, if $|z| > R$.

$$|z - 1| > 2 \qquad \text{and} \tag{1}$$

$$|Az - B| \geq |A||z| - |B| > \frac{1}{K_1}|z| - K_2 > 1. \tag{2}$$

Suppose $|z| > R$ and $f_\lambda(z)$ is of the form $z(z-1)(Az - B) + z$. Then we must show

$$|(z)(z-1)(Az - B) + z| > |z|,$$

or equivalently

$$|(z-1)(Az - B) + 1| > 1.$$

But this follows easily from the triangle inequality and by multiplying equations (1) and (2). A similar proof works for the other possible form of f. Q.E.D.

Lemma 7.10. Suppose $\{f_\lambda\}_{\lambda \in V}$ is a sub-family of $\mathcal{G}$, where $\mathcal{G}$ is parameterized by $(A, B) \in \mathbb{C} \times \mathbb{C}$. Furthermore, suppose there exists constants $K_1, K_2 > 0$ so that for $\lambda = (A, B) \in V$

$$\frac{1}{K_1} \leq |A| \leq K_1 \quad \text{and} \quad |B| \leq K_2.$$

Then there exists a radius $R > 0$ and potential radius $r > 1$ so that the potential curve, G^r of f_λ, is contained in the set $\{z \in \mathbb{C} \mid |z| > R\}$, which itself is contained in the basin at infinity of f_λ for every $\lambda \in V$.

PROOF. The radius R constructed in Lemma [7.9] is precisely what we need. The set $\{z \in \mathbb{C} \mid |z| > R\}$ cannot have any periodic points under any f_λ for $\lambda \in V$, as the modulus of every point is strictly increased by the action of f_λ. Therefore the complement of $\{z \in \mathbb{C} \mid |z| > R\}$ must contain the Julia set of every f_λ. Now we attempt to construct r.

Fixing f_λ, notice if G^r is not contained in $\{z \in \mathbb{C} \mid |z| > R\}$, then no potential curve G^s, with $s < r$, is contained it. Because, for every point in G^r, there is a point in G^s with smaller modulus. So suppose for each integer $i > 1$, we can find a $\lambda_i \in V$ so that the potential curve G^i of f_{λ_i} is not contained in $\{z \in \mathbb{C} \mid |z| > R\}$. The closure of V is a bounded set in $\mathbb{C} \times \mathbb{C}$ and is thus compact. Considering the sequence $\{\lambda_i\}$ as points in this compact set, there must be an accumulation point, μ. The map f_μ in $\mathcal{G}$ evidently has no potential curve at all in $\{z \in \mathbb{C} \mid |z| > R\}$. The coefficient A, of maps represented by the closure of V, is non-zero, so f_μ is a strictly cubic map. But this is a contradiction then as the potential curves G^i of any cubic map always bound a nested sequence of neighborhoods of infinity, whose intersection is equal to infinity. In other words, the G^i are eventually outside a disk of any radius. Therefore we can find an r for which G^r of f_λ is contained in $\{z \in \mathbb{C} \mid |z| > R\}$ as required, for every $\lambda \in V$. Q.E.D.

It is convenient having such universal sets which have the same dynamical behavior under every map in a family. We now prove a similar result about neighborhoods of a repelling point. This lemma is used to control the *end* of a ray, while the underlying map is varied.

Lemma 7.11. Suppose $\{f_\lambda\}_{\lambda \in V}$, for $V \subset \mathbb{C} \times \mathbb{C}$, is a family of conformal maps, depending analytically on $(A, B) = \lambda \in \mathbb{C} \times \mathbb{C}$. Also suppose for some $\lambda_0 \in V$, x_{λ_0} is a fixed point of f_{λ_0}, and suppose there exists a constant $K > 1$ so that the absolute value of the multiplier of the fixed point is bounded from below by K.

$$|M_{\lambda_0}| = |f'_{\lambda_0}(x_{\lambda_0})| \geq K > 1.$$

Then there exists a universal radius R and an open neighborhood of λ_0, $U \subset \mathbb{C} \times \mathbb{C}$, so that the fixed point x_{λ_0} persists as a function x_λ, analytic in $(A, B) = \lambda$. Moreover, with $D_\lambda(R)$ defined as the open disk of radius R centered at x_λ, for every $\lambda \in \overline{U}$,

(1) $\overline{D_\lambda(R)} \subset f_\lambda(D_\lambda(R))$

(2) No critical point of f_λ is contained in $D_\lambda(R)$.

PROOF. The point x_{λ_0} is of course a repelling fixed point, and it is well-known that it persists as a repelling fixed point of f_λ, in some neighborhood of λ_0. The multiplier, $M_\lambda = f'_\lambda(x_\lambda)$, varies analytically with λ. So we can certainly find a bounded open neighborhood, $U \subset \mathbb{C} \times \mathbb{C}$, of λ_0 satisfying the following finite conditions. For any $\lambda \in \overline{U}$,

$$|M_\lambda| \geq K_2 > 1 \qquad \text{where } K > K_2 > 1.$$

For each $\lambda \in \overline{U}$, f_λ looks locally like $M_\lambda(z - x_\lambda) + x_\lambda$. The point x_λ is still repelling, and if we choose a radius, R, small enough, the image of the disk $D_\lambda(R)$ under f_λ is approximately equal to $D_\lambda(|M_\lambda|R)$. Since $|M_\lambda|$ is bounded away from one, we can certainly find an R small enough so that property (1) is satisfied and f_λ is univalent on $D_\lambda(R)$, which implies property (2). So for every $\lambda \in \overline{U}$, there is *some* non-zero R satisfying (1) and (2). But we need to choose R to be universal.

We proceed by assuming such a universal radius does not exist. Then for each integer $i > 1$, we can find a $\lambda_i \in \overline{U}$ so that either f_{λ_i} is not univalent on $D_{\lambda_i}(R)$, for every $R \geq 1/i$, or equation (1) is violated, i.e.

$$\overline{D_{\lambda_i}(R)} \not\subset f_{\lambda_i}(D_{\lambda_i}(R)) \qquad \text{if } R \geq 1/i.$$

Since U is bounded in $\mathbb{C} \times \mathbb{C}$, $\overline{U}$ is a compact set, and we can find an accumulation point, $\mu \in \overline{U}$, of the sequence of points $\{\lambda_i\}$ and the corresponding map f_μ. By continuity, f_μ inherits many properties. The map must still have a fixed point x_μ, and for each $R > 0$, f_μ has a critical point in $D_\mu(R)$ or equation (1) is not satisfied. Since f_μ has only a finite number of critical points, either x_μ is a critical point of f_μ, or no R satisfies equation (1). But again by continuity we must have $|M_\mu| \geq K_2$, so f_μ is repelling at x_μ, so x_μ

is certainly not critical. As noted before there is *some* radius R for which f_μ is univalent on $D_\mu(R)$ and equation (1) is satisfied. This contradiction finishes the proof. Q.E.D.

We prove one more similar result about the existence of certain potential curves. This lemma provides control of an escaping critical point, making sure it is still close to the Julia set. Compare with Lemma [7.10].

Lemma 7.12. Suppose f_{λ_0} is a map in $\mathcal{G}$, where $\mathcal{G}$ is parameterized by $\lambda = (A, B) \in \mathbb{C} \times \mathbb{C}$. Suppose f_{λ_0} has no escaping critical orbits. Then, given $\epsilon > 0$, there exists a neighborhood of λ_0, $U \subset \mathbb{C} \times \mathbb{C}$, such that for $\lambda \in U$, the potential curve $G^{1+\epsilon}$ of f_λ is a simple closed curve in the plane, which contains the critical values, the critical points, and all precritical points of f_λ.

PROOF. This result is trivial if the filled Julia set of f_λ stays connected within a neighborhood of λ_0, but in general neighborhoods of λ_0 may contain maps with escaping critical orbits. If $\lambda_0 = (A_0, B_0)$, since $f_{\lambda_0} \in \mathcal{G}$, we know A_0 is non-zero. We choose a neighborhood of λ_0, V. If V is chosen to be small enough, then there exist constants $K_1, K_2 > 0$ so that for $\lambda = (A, B) \in V$,

$$\frac{1}{K_1} \leq |A| \leq K_1 \quad \text{and} \quad |B| \leq K_2.$$

Hence, we will be able to apply Lemma [7.10] to V.

Let an arbitrary $\epsilon > 0$ be chosen. Recall that each Böttcher coordinate, φ_λ, can be extended onto some maximal ray convex set which we label M_λ. The potential curve $G^{1+\epsilon}$, if it is defined at all, intersects rays of every angle, which then must be defined for every radius greater than $1+\epsilon$, by the definition of M_λ. In other words, if $G^{1+\epsilon}$ is in M_λ, then all points outside this potential are also in M_λ. There are no critical or precritical points of f in M_λ, so this potential, $G^{1+\epsilon}$, must enclose all the critical and precritical points. The proof then is reduced to showing that $G^{1+\epsilon}$ exists for f_λ near f_{λ_0}. Define the set W to be all $\lambda \in V$, where $G^{1+\epsilon}$ is not defined for f_λ. The point λ_0 is *not* in this set because f_{λ_0} has a connected Julia set and, by a corollary of Proposition [7.4], all potential curves are defined for f_{λ_0}. Therefore, if we can show W is closed, then $V - W$ is the desired neighborhood of λ_0, and we are done.

We prove W is closed by showing it contains its limit points. Suppose $\{\lambda_i\}_{i=1}^{\infty}$ is a sequence of points in W converging to λ_∞. By Lemma [7.10], there is a radius R and radius r so that $\{z \in \mathbb{C} \mid |z| > R\}$ is contained in the basin at infinity of f_λ, and the potential curve G^r is contained in this set, for all $\lambda \in V$. Recall that if $|\varphi_\lambda(x)| = s$, then $|\varphi_\lambda(f_\lambda(x))| = s^3$, and $f_\lambda(x)$ is outside of any potential curve of radius smaller than s^3. Therefore there exists a constant k, independent of λ, so that if $|\varphi_\lambda(x)| > 1+\epsilon$, then $|\varphi_\lambda(f_\lambda^{\circ k}(x))| > r$ and $f_\lambda^{\circ k}(x)$ is outside of G^r.

We claim that for any i, there exists a critical point of f_{λ_i}, which we label c_i, so that $f_{\lambda_i}^{\circ k}(c_{\lambda_i})$ is outside G^r. Since the Böttcher coordinate of each f_{λ_i} does not extend to include the potential curve of radius $1+\epsilon$, there must be a ray $\mathcal{R}_\theta^t$ that hits a critical or precritical point of f_{λ_i} before it extends to a radius of $1+\epsilon$. Let a be this critical or precritical point, and let x be a point on this ray which we can choose to be arbitrarily close to a. Let $j \geq 0$

be the integer which makes $f_{\lambda_i}^{\circ j}(a)$ critical. Then $f_{\lambda_i}^{\circ j}(a)$ is our c_i. Since x is close to a, $f_{\lambda_i}^{\circ j}(x)$ is close to c_i. Certainly since $j + k \geq k$, $f_{\lambda_i}^{\circ(j+k)}(x)$ is outside of G^r, therefore $f_{\lambda_i}^{\circ k}(c_i)$ must be also be outside as claimed.

Let λ_∞ be the limit of the $\{\lambda_i\}$. Since it is true for each i, by continuity and Lemma [2.3], f_{λ_∞} must have a critical point c_∞, with $f_{\lambda_\infty}^{\circ k}(c_\infty)$ outside of the potential curve G^r. The Böttcher coordinate of f_{λ_∞} cannot extend down to a radius of $1 + \epsilon$. There must be points, x, arbitrarily close to c_∞, with $|\varphi_{\lambda_\infty}(x)| \geq 1 + \epsilon$, and so there is an obstruction to extending φ_{λ_∞} down to this radius. Therefore λ_∞ is in W, and W is closed. Q.E.D.

Holomorphic motions of rays. We are now ready to construct a holomorphic motion on rays.

Lemma 7.13. Suppose f_{λ_0} is a map in $\mathcal{G}$, parameterized by $\lambda = (A, B) \in \mathbb{C} \times \mathbb{C}$, and f_{λ_0} has no escaping critical orbits. Suppose $\mathcal{R}_\theta$ is an external ray of f_{λ_0} landing at a repelling preperiodic point x which is not a critical or precritical point. Then there exists a neighborhood U of λ_0, open in $\mathbb{C} \times \mathbb{C}$, such that for $\lambda \in U$, the entire external ray $\mathcal{R}_\theta$ is defined for f_λ, i.e. the radial line $\hat{\mathcal{R}}_\theta$ in $\hat{\mathbb{C}} - \bar{D}$ is in the range of the Böttcher coordinate φ_λ. Furthermore, this ray lands at a repelling preperiodic point x_λ that defines a function, with $x_{\lambda_0} = x$, that varies analytically with $\lambda = (A, B)$.

Proof. We put off dealing with the preperiodic case until later and assume that x and $\mathcal{R}_\theta$ are periodic. We can normalize this situation even further. By choosing a high enough iterate of f_{λ_0}, we may assume that x is a repelling fixed point and $\mathcal{R}_\theta$ is a fixed ray, i.e. forward invariant under f_{λ_0}.

This normalization satisfies the hypotheses of both Lemma [7.11] and Lemma [7.12], so we can apply both these results. From Lemma [7.11] we can find a universal radius R and open neighborhood of λ_0, $U_1 \subset \mathbb{C} \times \mathbb{C}$, so that the disk $D_\lambda(R) = \{z \in \mathbb{C} \mid |z - x_\lambda| < R\}$ satisfies

$$\overline{D_\lambda(R)} \subset f_\lambda(D_\lambda(R)) \qquad \text{for every } \lambda \in U_1,$$

and $D_\lambda(R)$ contains no critical point of f_λ. Note that for every $\lambda \in U_1$, the inverse branch, f_λ^{-1}, from $f_\lambda(D_\lambda(R))$ onto $D_\lambda(R)$ is univalent and well-defined, forming a contraction towards the fixed point.

We look specifically at the original map, f_{λ_0}, again. The ray $\mathcal{R}_\theta$ lands at x_{λ_0}, so as we trace the ray in from infinity, it at some point must enter $D_{\lambda_0}(R)$ and not leave. Since $\mathcal{R}_\theta$ is fixed, if we pick a point on the ray, w_{λ_0}, close enough to x_{λ_0}, then $f_{\lambda_0}(w_{\lambda_0})$ is also on the ray and the segment of the ray between these two points is still contained in $D_{\lambda_0}(R)$. In the image space under φ_{λ_0}, the Böttcher coordinate, we get two corresponding points, $\hat{w}$ and $\hat{w}^3$, which must be on the same radial line of $\hat{\mathbb{C}} - \bar{D}$, and we get the corresponding radial segment in between. We parameterize this segment as follows.

$$\hat{w}(t) = t\hat{w} + (1-t)\hat{w}^3 \qquad 0 \le t \le 1.$$

For any λ where the image of φ_λ extends to this radial segment, we can parameterize the actual ray segment as follows.

$$w_\lambda(t) = \varphi_\lambda^{-1}(\hat{w}(t)) \qquad 0 \le t \le 1.$$

We claim that this ray segment is defined for λ in a neighborhood of λ_0.

Define $1+\epsilon$ to be the radius of a circle in $\hat{\mathbb{C}} - \bar{D}$ that does not contain this radial segment. For instance, with $\epsilon = \frac{1}{2}(|\hat{w}| - 1)$, we have $|\hat{w}| > 1$ and $|\hat{w}^3| > |\hat{w}| > 1 + \epsilon$. Now we are ready to apply Lemma [7.12] for f_{λ_0} and x_{λ_0}. Therefore, given this ϵ, we can find another open neighborhood, U_2 in $\mathbb{C} \times \mathbb{C}$, of λ_0 so that all precritical points and critical values are contained in the potential curve $G^{1+\epsilon}$, for $\lambda \in U_2$. Under the Böttcher coordinate, this means that either the image of φ_λ extends to all of $\hat{\mathbb{C}} - \bar{D}$ or at least to the complement of the closed disk of radius $1+\epsilon$, for every $\lambda \in U_2$. Therefore the radial segment between $\hat{w}$ and $\hat{w}^3$ are always in the image of φ_λ and $w_\lambda(t)$ is defined for all $\lambda \in U_2$ and $t \in [0, 1]$.

Since φ_λ depends analytically on λ, the points $w_\lambda(t)$ are analytical with respect to λ. The open disk $D_\lambda(R)$ also varies continuously with respect to λ. Since the boundary of $D_\lambda(R)$ and the ray segment $\{w_\lambda(t)\}_{0 \le t \le 1}$, are disjoint compact sets, we can find an open neighborhood of λ_0, $U_3 \subset U_2$, so that for every $\lambda \in U_3$, the whole ray segment $\{w_\lambda(t)\}_{0 \le t \le 1}$ is still contained in $D_\lambda(R)$. The open neighborhood, $U = U_1 \cap U_2 \cap U_3$, is our desired set.

For each $\lambda \in U$, we need to check two things. First, we check the entire ray $\mathcal{R}_\theta$ is well defined for f_λ. Because λ is in U_2, we have already shown that the ray of angle θ down to a radius of $1+\epsilon$, $\mathcal{R}_\theta^{(1+\epsilon)}$, is defined. This includes the segment of points $\hat{w}(t)$ for $0 \le t \le 1$. We will divide the rest of the ray into similar segments and inductively verify that we can extend φ_λ to each segment. Second, we will check that the sequence of ray segments converge to the fixed point x_λ, and therefore the ray $\mathcal{R}_\theta$ lands at x_λ.

Suppose we take the inverse image of the segment $\{\hat{w}(t)\}$, $0 \leq t \leq 1$, under the map $z \mapsto z^3$. We get three new radial segments, with three distinct angles. One of these angles must in fact be θ because θ is fixed under the tripling action of z^3. Since $[\hat{w}(1)]^3 = \hat{w}(0)$, the inner boundary of the old segment corresponds to the outer boundary of the new segment with angle θ. We shall check that the range φ_λ can be extended to this particular segment, extending the image of $\mathcal{R}_\theta$ under φ_λ from the radial line $\hat{\mathcal{R}}_\theta^{|\hat{w}(1)|}$ to the longer radial line $\hat{\mathcal{R}}_\theta^{|\hat{w}(1)|^{1/3}}$.

Switching back to the dynamical plane of f_λ, the corresponding operation is to take the preimage of the ray segment between $w_\lambda(0)$ and $w_\lambda(1)$ via the inverse branch of f_λ mapping $D_\lambda(R)$ into itself. This new segment is obviously still in $D_\lambda(R)$ and has as endpoints $w_\lambda(1)$ and the image of $w_\lambda(1)$ under the inverse branch. Assuming this new segment does not hit any critical or precritical points, because of the functional equation, this must be an extension of the ray $\mathcal{R}_\theta$, and this extension stays in the disk $D_\lambda(R)$. We have shown even more.

Clearly the whole construction can be repeated to generate the entire radial line $\hat{\mathcal{R}}_\theta$ and ray $\mathcal{R}_\theta$. So long as critical and precritical points are avoided, we get a sequence of new segments of $\mathcal{R}_\theta$, all parameterized as follows. With $0 \leq t \leq 1$, the first new segment looks like $\{f_\lambda^{-1}(w_\lambda(t))\}$, the next new segment looks like $\{f_\lambda^{-1} \circ f_\lambda^{-1}(w_\lambda(t))\}$, and so on. Each point $w_\lambda(t)$ is iterated by the branch f_λ^{-1} which contracts to the fixed point x_λ, proving that $\mathcal{R}_\theta$ lands at x_λ, if it is defined.

It is simple to verify that new segments avoid critical and precritical points of f_λ. Suppose some segment contains a precritical point. All the points on the forward orbit of this precritical point, including the critical point, are either on previously defined segments or on the original unextended ray $\mathcal{R}_\theta^{(1+\epsilon)}$. Therefore the critical point is either on $\mathcal{R}_\theta^{(1+\epsilon)}$, which is not true by the definition of rays, or the critical point is in $D_\lambda(R)$, which is not true by Lemma [7.11]. Therefore the entire external ray $\mathcal{R}_\theta$ is defined for f_λ.

The proof for preperiodic points is just about the same, but we have to keep even tighter control over the critical orbit. Recall that a preperiodic ray is the preimage under some iterate of the dynamical map f_λ of some periodic ray. Suppose the ray $\mathcal{R}_\alpha$ of f_{λ_0} is mapped by $f_{\lambda_0}^{\circ j}$ on the periodic ray $\mathcal{R}_\beta$. We know $\mathcal{R}_\beta$ persists for λ near λ_0 by what we have just proved. This will assist us in checking that f_λ does not contain any critical or precritical points on $\mathcal{R}_\alpha$.

Lemma [7.12] will work in exactly the same way to keep precritical points off of $\mathcal{R}_\alpha$ down to a radius of $1+\epsilon$ for λ in some neighborhood of λ_0. By hypothesis the landing point of $\mathcal{R}_\alpha$ is not critical or even precritical, so there must exist some disk V centered on the landing point x which is mapped by $f_{\lambda_0}^{\circ j}$ to a neighborhood of the landing point of $\mathcal{R}_\beta$. We choose V small enough so that it does not contain any of the finite number of critical or precritical points of f_{λ_0} that have a critical point within the first j points of their forward orbits. Since x and these finite points all vary continuously with λ we can construct V to be universal for λ in some neighborhood of λ_0. This construction is similar in spirit to Lemma [7.11]. We choose a radius $1+\epsilon$, as before, so that $\mathcal{R}_\alpha$ for f_{λ_0} is entirely inside V below that radius. The rest of $\mathcal{R}_\alpha$ below this radius can be shown to avoid critical and precritical points by pulling back $\mathcal{R}_\beta$ through $f_\lambda^{\circ j}$, which is univalent on V. The precritical points that are critical after j iterations are avoided by construction of V. Any other precritical point must also be avoided or else its forward image under $f_\lambda^{\circ j}$ would be a precritical point on $\mathcal{R}_\beta$. The entire $\mathcal{R}_\alpha$ is now constructed. Q.E.D.

Quasiconformal extensions of holomorphic motions on rays. We apply the λ-Lemma to holomorphic motions constructed with rays. This produces quasiconformal homeomorphisms which map rays of one map f onto the corresponding rays of another map $\hat{f}$. In Chapter VIII part 3, we show that if the rays happen to form the boundary of the domains of box mappings, the homeomorphism can form a branchwise equivalence.

Theorem 7.14. Suppose $f_{\lambda_0} \in \mathcal{F}$ is in $\mathcal{L} = \{f_\lambda\}_{\lambda \in L}$, where L is bone-loop in parameter space. Let $\mathcal{A}$ be some subset of $\mathbb{C}$, symmetric with respect to $\mathbb{R}$, containing a finite number of external rays of f_{λ_0}, which land on real repelling preperiodic points of f_{λ_0}, and also containing all points outside of some potential curve of f_{λ_0}, G^r. Let $f_{\lambda_1} \in \mathcal{F}$ be any other map in $\mathcal{L}$. Then there exists a quasiconformal homeomorphism H of the entire sphere such that

$$H|_{\mathcal{A}} = \varphi_{\lambda_1}^{-1} \circ \varphi_{\lambda_0},$$

where φ_{λ_0} and φ_{λ_1} are the Böttcher coordinates for f_{λ_0} and f_{λ_1} respectively.

Proof. The definition of $H|_{\mathcal{A}}$ is representative of a whole family of maps $\varphi_\mu^{-1} \circ \varphi_\lambda$, where μ and λ are allowed to vary within our parameter space $\mathcal{G}$.

We will show that these maps form injective holomorphic motions on certain subsets of the plane, like $\mathcal{A}$. The heart of the proof is the λ-Lemma, which will extend these maps to full quasiconformal homeomorphism of $\mathbb{C}$, and H will be easily obtained by composition of these maps.

Recall that any map in $\mathcal{F}$ has no escaping critical orbits. We start by choosing an $\epsilon > 0$ so that $1 + \epsilon < r$ and then apply Lemma [7.12] to every λ for f_λ in $\mathcal{L}$. The union of all the neighborhoods obtained is an open set U_0 in $\mathbb{C} \times \mathbb{C}$ containing all the parameters in L,and the potential curve $G^{1+\epsilon}$ of f_λ always contains the critical values, critical points, and precritical points of f_λ, for $\lambda \in U_0$. Another way of viewing this is that the circle centered at the origin of radius $1 + \epsilon$ is in the range of φ_λ, for all $\lambda \in U_0$. Similarly, for each ray of f_{λ_0} in $\mathcal{A}$, the corresponding ray for any $f_\lambda \in \mathcal{L}$ is preperiodic and must land at a repelling preperiodic point (See Theorem [7.7]). Thus we can apply Lemma [7.13] to each map f_λ in $\mathcal{L}$ and take the union of the resulting neighborhoods. We obtain a finite collection of open sets containing L, $\{U_i\}$, one for each ray in $\mathcal{A}$, so that every f_λ for $\lambda \in U_i$ has a well-defined ray of the same angle. In other words, if $\mathcal{R}_{\theta_i}$ is in $\mathcal{A}$ for f_{λ_0}, then the radial line $\hat{\mathcal{R}}_{\theta_i}$ in $\hat{\mathbb{C}} - \bar{D}$ is in the range of φ_λ, for every $\lambda \in U_i$.

The sets U_0 and $\{U_i\}$ form a finite collection of open neighborhoods of the parameters in $\mathbb{C} \times \mathbb{C}$ forming the bone-loop L. So the intersection of all these sets forms an open set U, also containing L. The image of $\mathcal{A}$ under φ_{λ_0} is *straightened*, formed of radial lines and concentric circles, and for every $\lambda \in U$, φ_λ^{-1} has this, $\varphi_{\lambda_0}(\mathcal{A})$ in its domain. Therefore $H = \varphi_\lambda^{-1} \circ \varphi_{\lambda_0}$ is well-defined on $\mathcal{A}$.

The set of parameters L form a compact subspace of $\mathbb{C}\times\mathbb{C}$ or of $\mathbb{R}\times\mathbb{R}$. To prove this theorem, it is enough to initially construct conjugacies between f_λ and f_{λ_0} only for λ varying through some open neighborhood of λ_0 in L. Since λ_0 is arbitrary, we could then find such open neighborhoods around every $\lambda_0 \in L$. Then by using a finite sub-cover of L, we can construct the desired extensions between any two maps in $\mathcal{L}$, by composing a finite number of these initial extensions. We fix an arbitrary λ_0 and construct this open set of parameters.

Since λ_0 is contained in U, an open set in $\mathbb{C}\times\mathbb{C}$, we can find an open ball, V, around λ_0 in U. If $d(\lambda_1, \lambda_2)$ is the standard metric on $\mathbb{C}\times\mathbb{C}$, then for some radius $R > 0$, the set $V = \{\lambda \in \mathbb{C}\times\mathbb{C} \mid d(\lambda, \lambda_0) < R\}$ is completely contained in U. The real axes $\mathbb{R}\times\mathbb{R}$, which contain the maps in $\mathcal{F}$, are naturally embedded in $\mathbb{C}\times\mathbb{C}$. Since V is centered at $\lambda_0 \in \mathbb{R}\times\mathbb{R}$, it intersects $\mathbb{R}\times\mathbb{R}$ in a two real-dimensional disk, also of radius R. We label a disk slightly smaller. Let $W = \{\lambda \in \mathbb{R}\times\mathbb{R} \mid d(\lambda, \lambda_0) < \frac{1}{4}R\}$. We attempt to show, for every $\lambda \in W$, that a quasiconformal extension of $\varphi_\lambda^{-1} \circ \varphi_{\lambda_0}$ exists.

Fix an arbitrary point $(A_1, B_1) = \lambda_1$ in W, and label the points $\lambda_0 = (A_0, B_0)$ and $\nu = (A_1, B_0)$ in $\mathbb{C}\times\mathbb{C}$. To invoke the λ-Lemma, we need to construct a parameter space which is a disk of exactly one complex dimension. To this end we define

$$D_1 = \{(w, B_0) \in \mathbb{C}\times\mathbb{C} \mid |w - A_0| < \frac{1}{3}R\} \qquad \text{and}$$

$$D_2 = \{(A_1, w) \in \mathbb{C}\times\mathbb{C} \mid |w - B_0| < \frac{1}{3}R\}.$$

Notice that if (w, B_0) is in D_1, then

$$d((w, B_0), (A_0, B_0)) = |w - A_0| < \frac{1}{3}R.$$

If (A_1, w) is in D_2, then by the triangle inequality

$$\begin{aligned} d((A_1, w), (A_0, B_0)) &\leq d((A_1, w), (A_1, B_0)) + d((A_1, B_0), (A_0, B_0)) \\ &= |w - B_0| + |A_1 - A_0| < \frac{1}{3}R + \frac{1}{4}R. \end{aligned}$$

Therefore both D_1 and D_2 are contained in $V \subset U$, and note that ν is in D_1.

We are now prepared to define a holomorphic motion suitable for extension by the λ-Lemma. On $\mathcal{A}$ define the map

$$i_w = \varphi_\lambda^{-1} \circ \varphi_{\lambda_0} \qquad \text{with } |w| < 1 \text{ and } \lambda = (A_0 + \tfrac{R}{3}w, B_0) \in D_1.$$

Since $i_w(z)$ is constructed out of Böttcher coordinates, for each fixed z, $i_w(z)$ is a holomorphic function in w. Furthermore, i_w is injective, and for $w \in \mathbb{R}$, i_w is symmetric and preserves the orientation of the $\mathbb{R}$ because each of the coordinates have the same properties. The map i_0 is clearly the identity on $\mathcal{A}$. A second motion can be defined on $\varphi_\nu^{-1} \circ \varphi_{\lambda_0}(\mathcal{A})$.

$$j_w = \varphi_\lambda^{-1} \circ \varphi_\nu \qquad \text{with } |w| < 1 \text{ and } \lambda = (A_1, B_0 + \tfrac{R}{3}w) \in D_2.$$

The motion $j_w(z)$ is also injective, and for $w \in \mathbb{R}$, is also symmetric and preserves the orientation of $\mathbb{R}$.

For $w = \frac{3}{R}(A_1 - A_0)$, i_w is extended by the λ-Lemma, Theorem [7.8], to a quasiconformal map H_1 with

$$H_1(z) = i_w(z) = \varphi_\nu^{-1} \circ \varphi_{\lambda_0}(z) \qquad \text{for } z \in \mathcal{A}.$$

For $w = \frac{3}{R}(B_1 - B_0)$, j_w is extended by the λ-Lemma to the quasiconformal map H_2 with

$$H_2(z) = j_w(z) = \varphi_{\lambda_1}^{-1} \circ \varphi_\nu(z) \qquad \text{for } z \in \varphi_\nu^{-1} \circ \varphi_{\lambda_0}(\mathcal{A}).$$

Composing these maps produces a quasiconformal map $H = H_2 \circ H_1$, and restricted to $\mathcal{A}$, we have

$$H|_\mathcal{A} = \varphi_{\lambda_1}^{-1} \circ \varphi_\nu \circ \varphi_\nu^{-1} \circ \varphi_{\lambda_0} = \varphi_{\lambda_1}^{-1} \circ \varphi_{\lambda_0}.$$

Q.E.D.

Corollary. The quasiconformal map H can be chosen to be symmetric and to preserve the orientation of the real line.

PROOF. We have constructed a quasiconformal homeomorphism of $\mathbb{C}$ that is equal to $\varphi_{\lambda_1}^{-1} \circ \varphi_{\lambda_0}$ on the set $\mathcal{A}$. Since $\mathcal{A}$ is symmetric about $\mathbb{R}$ and the Böttcher coordinates are derived from real polynomials, the map H is symmetric and preserves the orientation of the real line when restricted to $\mathcal{A}$. Unfortunately, H is not necessarily symmetric on the complement.

The complement of $\mathcal{A}$ is a finite collection of topological disks with boundaries equal to pieces of rays and potentials. If some component of the complement, intersects $\mathbb{R}$, it must be symmetric with respect to $\mathbb{R}$. Let Δ be such a component, and let $\hat{\Delta} = H(\Delta)$ be the corresponding component in the complement of $\hat{\mathcal{A}} = H(\mathcal{A})$. We show that on Δ, H can be redefined so that H remains unchanged on the boundary and is still quasiconformal but maps $\Delta \cap \mathbb{R}$ onto $\hat{\Delta} \cap \mathbb{R}$. Performing this operation on all components of the complement of $\mathcal{A}$ intersecting $\mathbb{R}$ will produce a homeomorphism that preserves the orientation of $\mathbb{R}$ and is glued together from a finite number of quasiconformal pieces. The resulting homeomorphism is quasiconformal. To guarantee that H is symmetric, we reflect the restriction of H to the upper half plane $\mathbb{H}^+$. In other words we replace $H(z)$, for $z \in \mathbb{H}^-$, with $\bar{H}(\bar{z})$. This will not change H on $\mathcal{A}$ or $\mathbb{R}$ because H is already symmetric here. The final result is a quasiconformal map equal to $\varphi_{\lambda_1}^{-1} \circ \varphi_{\lambda_0}$ on $\mathcal{A}$, which is symmetric everywhere and preserves the orientation of $\mathbb{R}$.

We begin the construction of an H on Δ which preserves $\mathbb{R}$. Define a Riemann map g which takes the upper half plane $\mathbb{H}^+$ onto Δ. This map will extend continuously to the real line. If we require that 0 and ∞ map to the boundary points of $\Delta \cap \mathbb{R}$, then the Riemann map will be unique, and by symmetry, g must map the positive imaginary axis onto $\Delta \cap \mathbb{R}$. Similarly, define the Riemann map $\hat{g}$, mapping $\mathbb{H}^+$ onto $\hat{\Delta}$. We know

$$g(-r) = \bar{g}(r) \in \partial\Delta \quad \text{and} \quad \hat{g}(-r) = \bar{\hat{g}}(r) \in \partial\Delta \qquad \text{for } r \in \mathbb{R}.$$

Since H is symmetric on the boundary of Δ, we also know

$$\hat{g}^{-1} \circ H \circ g(-r) = \hat{g}^{-1} \circ H(\bar{g}(r)) = \hat{g}^{-1}(\bar{H} \circ g(r)) = -\hat{g}^{-1} \circ H \circ g(r). \quad (1)$$

Let Q_1 and Q_2 be two injective maps, one defined on the first quarter of $\mathbb{C}$ and one on the second quarter, with range equal to $\mathbb{H}^+$. In polar coordinates, they are defined as

$$Q_1(r,\theta) = (r, 2\theta) \qquad \text{and} \qquad Q_2(r,\theta) = (r, 2\theta - \pi).$$

This definition gives

$$Q_1(ri) = -r \quad \text{and} \quad Q_2(ri) = r \qquad \text{for } r > 0. \tag{2}$$

$$Q_1(r) = r \quad \text{and} \quad Q_2(-r) = -r \qquad \text{for } r > 0. \tag{3}$$

It is easy to check that Q_1 and Q_2 are both quasiconformal. The two maps $Q_1^{-1} \circ \hat{g}^{-1} \circ H \circ g \circ Q_1$ and $Q_2^{-1} \circ \hat{g}^{-1} \circ H \circ g \circ Q_2$ are then quasiconformal homeomorphisms of the first and second quarters of $\mathbb{C}$ respectively onto themselves. By equation (2) and equation (1), both these maps agree on the positive imaginary axis, and we can glue them together to form a quasiconformal homeomorphism Q_3 taking the upper half plane $\mathbb{H}^+$ onto itself.

Now replace H on Δ with $H_2 = \hat{g} \circ Q_3 \circ g^{-1}$, which is quasiconformal and maps Δ to $\hat{\Delta}$. If $z \in \partial\Delta \cap \mathbb{H}^+$, then $g^{-1}(z) \in \mathbb{R}^-$, and

$$Q_3(g^{-1}(z)) = Q_2^{-1} \circ \hat{g}^{-1} \circ H \circ g \circ Q_2(g^{-1}(z)) = Q_2^{-1} \circ \hat{g}^{-1} \circ H(z),$$

by equation (3). Since $H(z) \in \partial\hat{\Delta} \cap \mathbb{H}^+$, we know $\hat{g}^{-1}(H(z)) \in \mathbb{R}^-$, and by equation (3) again

$$Q_3(g^{-1}(z)) = Q_2^{-1}(\hat{g}^{-1}(H(z))) = \hat{g}^{-1}(H(z)).$$

Therefore $H_2(z) = H(z)$, for $z \in \partial\Delta \cap \mathbb{H}^+$. A similar proof works for $z \in \partial\Delta \cap \mathbb{H}^-$, so H_2 agrees with H on the boundary of Δ.

Now if $z \in \Delta \cap \mathbb{R}$, then $g^{-1}(z)$ is on the imaginary axis, and if z is on the imaginary axis, then $\hat{g}(z)$ is in $\hat{\Delta} \cap \mathbb{R}$. Therefore $H_2(z)$ is in $\hat{\Delta} \cap \mathbb{R}$. Therefore H_2 is quasiconformal and preserves the real line. Q.E.D.

VIII. Applications to cubic maps

§8.1. Real box mappings induced from cubic polynomials.

We want to construct box mappings induced from maps $f \in \mathcal{F}$. We must keep in mind these issues involved in construction.

(1) **Markov partition** Suppose U is the image, under the box mapping, of Δ, a connected component of the domain, and Δ' is any other such domain component. Then either $U \cap \Delta' = \Delta'$ or $U \cap \Delta' = \emptyset$. This property is required for any box mapping.

(2) **Central domain** There is at most one connected component of the domain on which the box mapping is folding, the central domain.

(3) **Critical recurrence** In order for the inducing algorithm to proceed to a limit, the forward orbit of the critical point in the central domain must be contained in the domain of the box mapping. The only exception to this is if the critical point is *non-recurrent* under f, meaning the forward orbit does not contain the critical point itself as a limit point. As will be seen, such polynomials can produce box mappings with no central branch, where the inducing algorithm is not needed.

There is one other minor property we will need which is not strictly required for box mappings but is needed to apply Theorems [6.2] and [6.1].

(4) **Type II** The image of any component of the domain is equal to B', the image of the central branch.

A box mapping is a generalization of a *first return* map and often has an inducing interval I similar to that used to construct first return maps. The inducing interval I can be considered as the smallest interval containing the range of the box mapping. The image of points under our initial box mappings will in fact be defined as the first return of that point under the iteration of the dynamical map f to an inducing interval I. Careful choice of the inducing interval causes the first return map to satisfy all four properties. In particular we cut the immediate basin of the periodic point out of the inducing interval for maps on a bone-loop, so that only one critical point is in the domain of the resulting map. As is the case of any first return map, selection of the inducing interval I and the dynamical function f almost completely defines the final box mapping. In view of the four properties listed above, the construction of I is natural.

For the rest of this chapter we will fix a real cubic polynomial $f \in \mathcal{F}$ that has, like maps on a bone-loop, exactly one periodic critical point and one chaotic critical point in the Julia set of f.

Inducing intervals. A box mapping induced from a cubic map is a first return map to a special *inducing interval*, constructed to eliminate the one of the two critical points.

Definition. (F_i) Since f has an attracting periodic orbit, its Fatou set has a bounded component that contains the periodic critical point. We label this set F_1. We know F_1 is a topological disk, symmetric about the real line, and it intersects $\mathbb{R}$ in a single open interval. If the critical point has period k, the entire forward orbit of F_1 must consist of k components, F_1 through F_k, each symmetric and intersecting $\mathbb{R}$ in a single open interval. Together these sets from the immediate basin of the periodic critical orbit.

Under f, the sets map onto each other in a cycle, in a one-to-one manner, except for F_1 which is mapped in a two-to-one manner. Note that the boundary points of the open interval intersections of these sets must all be periodic or preperiodic points.

Definition. (the inducing interval) The interval (a, b) is an **inducing interval** for f if the following are true.

(1) (a, b) contains the chaotic critical point.

(2) (a, b) is in the complement of $\mathbb{C} - \bigcup_{i=1}^{k} \overline{F_i}$

(3) Both a and b are repelling preperiodic points whose orbits do not intersect (a, b).

Look at the complement of the $\overline{F_k}$, $\mathbb{C} - \bigcup_{i=1}^{k} \overline{F_i}$. We know the chaotic critical point is not in the $\{F_i\}$, so it must be in the complement. The connected component containing the chaotic critical point is a natural candidate for an inducing interval.

Lemma 8.1. Suppose $f \in \mathcal{F}$ is on a bone-loop. Let I be the interval in $\mathbb{C} - \bigcup_{i=1}^{k} \bar{F}_i$ containing the chaotic critical point of f. Then I is an inducing interval for f.

PROOF. Property (1) and (2) of the inducing interval are satisfied by construction. Let x be either boundary point of I. By construction, x is also a boundary point of some connected component of the immediate attracting basin of f, and by continuity, so is every point in the forward orbit of x. The set of all boundary points of the immediate attracting basin of f forms a finite set which does not intersect the interior of I, as I itself is in the complement of the closure of the immediate basin. Q.E.D.

Cubic first return maps. We fix an inducing interval I for f. We now check to see if the first return map to I is a box mapping.

Definition. (Cubic first return map, Φ) The map Φ, with respect to a map $f \in \mathcal{F}$ on a bone-loop and an inducing interval I, is defined for points in $(0, 1)$ to be the first return map of f to I.

The first return map Φ will be undefined on entire intervals in the Fatou set whose iterates fall into the immediate basin and thus never hit the inducing interval. There may also be points whose orbit never hit the closure of the immediate basin but never enter I either. However these points must always be isolated because an entire interval in the Julia set that avoided the critical point would be a wandering interval. If the critical orbit is recurrent, then Φ will be defined on every point of that orbit. Each point always avoids the immediate basin and must eventually return close to the critical point in I

The chaotic critical orbit of f will never hit the components of the attracting basin, F_i, but it may not return to I itself. The question of whether it does or not splits our construction of box mappings into two distinct cases. If the critical point does return, there will be some entire interval around the critical point that returns to I with it. The map Φ will be **folding** or two-to-one here. Because we have eliminated the other critical point from consideration and since the chaotic critical point is in I, no other interval on which the cubic first return map is continuous will be folding. This single folding interval will lead to the central branch of our box mapping.

If the critical orbit does not return, things will be much the same. However there will be no interval at all on which the first return map is folding. This will lead to a box mapping with no central domain. This case is very similar to the behavior of Misiurewicz polynomials in the quadratic case, which have a non-recurrent critical orbit. The most obvious situation where this might happen is when the critical orbit hits the boundary of I before it hits the interior. The proof that Φ is a box mapping goes in exactly same way, whether or not the critical orbit returns to I. The properties of box mappings concerning the central domain are all optional. We only need to take into account both possibilities.

Suppose the critical orbit does return to I, then we will have a *central domain.*

Definition. (the central domain of Φ) By assumption, the function Φ is defined at c, the chaotic critical point. There must be some maximal open interval, containing c, where points near c also return to I on the same iteration as c. Clearly Φ is continuous on this interval. We label the interval B, the central domain of Φ.

Such open intervals in the domain of Φ can be constructed around *any* point where Φ is defined. We will call these intervals *components of the domain of* Φ, and later, when we extend these definitions to the complex plane, this name will also apply to open disks in the plane, symmetric about the real line. For now, we continue to consider only real intervals. We make some preliminary observations about these intervals and what Φ looks like restricted to them.

Lemma 8.2. Let (a, b) be any maximal open interval on which Φ is defined as a single iteration of f, $f^{\circ i}$, for some integer i. Then the extensions of Φ to a and b, $\lim_{x \to a^+} \Phi(x)$ and $\lim_{x \to b^-} \Phi(x)$, are boundary points of I. Thus a and b are preperiodic, and Φ is undefined at a and b.

PROOF. The map $\Phi|_{(a,b)} = f^{\circ i}$ is continuous, so we get immediately

$$\lim_{x \to a^+} \Phi(x) = f^{\circ i}(a) \quad \text{and} \quad \lim_{x \to b^-} \Phi(x) = f^{\circ i}(b).$$

The point $f^{\circ i}(a)$ must be part of the closure of the inducing interval I by continuity, but it cannot be in the interior. If it were, each of the first i points in the orbit of a would be in the interior of the complement of the F_k. With a slight perturbation of a, these orbit points would still be in the interior, and Φ would still be defined on a larger interval. Thus the set (a, b) would not be maximal. Therefore $f^{\circ i}(a)$ must be a boundary point of I. We also claim, for any $j < i$, $f^{\circ j}(a)$ is not in the interior of I. Otherwise, points in (a, b) arbitrarily close to a would also return to the interior of I after j iterations, rather than i iterations, and $\Phi|_{(a,b)}$ would not equal $f^{\circ i}$. Similar arguments also hold for the point b.

We have shown that the orbits of a and b under f first intersect I on its boundary and the conclusion follows from the definition of an inducing interval. Q.E.D.

This lemma proves that maximal domains, on which Φ is some constant iterate of the map f can also be characterized as connected components of the domain of Φ. We can also characterize what Φ looks like on most of these intervals.

Lemma 8.3. Suppose Δ is any connected component of the domain of Φ that does not contain a critical point of f, in other words, any component which is not B. Then Φ restricted to Δ is a homeomorphism onto the interior of I.

PROOF. Recall Φ is always undefined on the immediate basin of the periodic critical point of f. Now suppose x is a critical point of Φ in Δ. By definition x is a critical point of some iterate of f, $f^{\circ k} = \Phi|_\Delta$, for some integer k. This can happen only if one of the first k points in the orbit of x under f, $\{x, f(x), f^{\circ 2}(x), \dots, f^{\circ(k-1)}(x)\}$, is a critical point of f, chaotic or periodic. None of these points can be the periodic critical point, or else $\Phi(x)$ would be undefined. The point x itself cannot be the chaotic critical point, which is not in Δ. The next $k-1$ points in the orbit cannot be the chaotic critical point because this point is in the interior of I, and the orbit of x does not intersect I until the k-th iterate. We have reached a contradiction, so x cannot be a critical point of Φ. This shows that $\Phi|_\Delta$ is a homeomorphism, and by Lemma [8.2], $\Phi|_\Delta$ must extend to the boundary of I. Q.E.D.

This shows that Φ has the same qualitative behavior on all the components in its domain, except possibly the central domain B. We establish a notational convention here to use the symbol Δ to represent arbitrary components of the domain of Φ on which Φ is a injective. We have the symbol B to represent the central domain of Φ, and in general, Δ will not be used to refer to this interval.

The Lemma [8.2] has consequences for the central domain B, if it exists, as well.

Lemma 8.4. Suppose B is equal to the open interval (a, b). Then a and b are preperiodic, $f(a) = f(b)$, and Φ restricted to $B - \{c\}$ is a two to one cover onto its image.

PROOF. The map $\Phi|_B$ is folding on B because B contains the critical point c, and this must be the only fold. By Lemma [8.2], a and b are preperiodic, and the range of $\Phi|_B$ extends to the boundary points of I. Because Φ folds only once on B, the boundary points of B are then not only preperiodic, but their orbits must hit the same boundary point of I on the same iteration. In other words, if $\Phi|_B = f^{\circ i}$, then $f^{\circ i}(a) = f^{\circ i}(b) \in \partial I$.

Consider the orbit of the set B as it eventually returns to I. For each $j < i$, the set $f^{\circ j}(B)$ is connected and must be contained in some connected component of the domain of Φ. If $j \neq 0$, then the component containing $f^{\circ j}(B)$ must be one of the monotone components Δ outside of I. In any case, all the folding that Φ does to B must occur in the first iterate of f. We must have $f(a) = f(b)$.

It now follows quickly that Φ is a two to one cover on $B - \{c\}$. Since

c is the only critical point of $\Phi|_B$, we must have that Φ restricted to either (a, c) or (c, b) is an injective map each with image $(f(c), f(a)) = (f(c), f(b))$. Q.E.D.

§8.2. Extending real box mappings.

We have demonstrated most of the key properties of Φ, and in fact Φ satisfies the definition of a real Type II Box mapping. However we will leave the remaining details of the proof until later. Now we concentrate on extending Φ to a holomorphic box mapping. For each interval in the domain of Φ, we will construct an extension of that interval to the complex plane and define Φ to be the analytic continuation of Φ. The extension of the interval is constructed so that its boundaries are pieces of potential curves and rays landing at the endpoints of the interval. The resulting topological disk is symmetric about the real line and its intersection with $\mathbb{R}$ is the original interval. This is an example of a *Yoccoz puzzle piece.* We have already shown how such an object can be used to build a quasiconformal map using the λ-Lemma.

Puzzle pieces and the resulting *Yoccoz partition* of the plane form a link between the classification questions we are studying here, monotonicity and conjugacy classes for instance, and geometric questions about Julia sets and parameter space. It is therefore natural that puzzle pieces form the frame upon which quasiconformal conjugacies are built, the geometric objects we are most interested in. For now we show how puzzle pieces can also provide a complex region which mirrors the dynamics on the real line.

Extended interval regions. For real maps, every ray $\mathcal{R}_\theta$ has a natural reflection $\mathcal{R}_{-\theta}$, which is literally the reflection of $\mathcal{R}_\theta$ about $\mathbb{R}$. In particular, if $\mathcal{R}_\theta$ lands at a real point x, then so does $\mathcal{R}_{-\theta}$. In such a situation, the set of points forming the two rays and the landing point x will be called a **ray pair**. The union of the ray pair with infinity forms a Jordan curve in $\hat{\mathbb{C}}$, dividing the sphere into two pieces. As a convenience, we define a notation for these pieces.

Definition. (ray pair sets) Let $\mathcal{R}_\theta$, $\mathcal{R}_{-\theta}$ and the common landing point be a ray pair of f. Also let a be some point in the complement of the ray pair. As noted above, the ray pair divides the sphere into two open pieces. We define $O_\theta(a)$ to be the set containing a.

Recall that by Theorem [7.6], if x is a repelling periodic point, then there are always ray pairs landing at it. Now, if x is preperiodic and its orbit hits a periodic point which is repelling, there must be inverse images of the rays landing at the periodic point that form rays and ray pairs landing at the preperiodic x. As is proved in Lemma [8.2], the boundary points of domain components of the first return map Φ are all preperiodic and always have at least one ray pair landing at them.

For domain components Δ of Φ, or any interval with repelling preperiodic boundaries, we can construct an extended domain, also labeled Δ, whose intersection with $\mathbb{R}$ is the original interval.

Definition. (extended interval regions) Let $\Delta = (a, b)$ be an open interval of the real line with a boundary formed by repelling preperiodic points of a real polynomial f. We know at each boundary point there is at least one ray pair of f which lands there. Let $\mathcal{R}_\alpha$ and $\mathcal{R}_{-\alpha}$ form the *rightmost* pair landing at the left boundary point a, and let $\mathcal{R}_\beta$ and $\mathcal{R}_{-\beta}$ form the *leftmost* pair landing at the right boundary point, b. We then define the *extended interval region* to be the set $O_\alpha(b) \cap O_\beta(a)$. This is clearly bounded by the α and β ray pairs and is the smallest of all regions bounded by ray pairs landing at a and b, which contain the interval (a, b).

Definition. (bounded interval regions) A bounded interval region, Δ^r, is obtained from an extended interval region Δ by truncating it with the potential curve G^r. Thus Δ^r is the intersection of Δ with the disk bounded by G^r.

The careful choice of boundary ray pairs, in situations where multiple pairs land at a point, is necessary to handle issues of containment. The Markov properties all stem from the construction of sets that nest properly. We need to make sure that if one interval contains another interval, then the same containment holds for the corresponding extended interval regions. One major problem has already been circumvented by using rays as boundaries because distinct rays of f do not intersect each other. This combined with the careful choice of rays among multiple ray pairs landing at a common point, yields this easy but crucial lemma.

Lemma 8.5. Let Δ_1 and Δ_2 be real intervals with repelling preperiodic boundary points under $f \in \mathcal{F}$. Suppose Δ_1 and Δ_2 satisfy one of the three containment properties

(1) $\Delta_1 \subset \Delta_2$

(2) $\Delta_2 \subsetneq \Delta_1$

(3) $\Delta_1 \cap \Delta_2 = \emptyset$

Then Δ_1 and Δ_2 considered as extended interval regions or bounded interval regions, Δ_1^r and Δ_2^r, satisfy the same property.

PROOF. Suppose $\Delta_1 = (a_1, b_1)$ and $\Delta_2 = (a_2, b_2)$. Corresponding to the points a_1, b_1, a_2, and b_2, suppose α_1, β_1, α_2, and β_2 are the angles used to construct the ray pair boundaries of the extended interval regions of Δ_1 and Δ_2.

We claim that if $a_1 \leq a_2$, then $O_{\alpha_2}(b_2) \subset O_{\alpha_1}(b_1)$. If a_1 and a_2 are equal, then we must also have $\alpha_1 = \alpha_2$, both angles representing the rightmost ray pair landing at $a_1 = a_2$. Since both b_1 and b_2 fall to the right of $a_1 = a_2$, we have $O_{\alpha_1}(b_1) = O_{\alpha_2}(b_2)$ and the claim is true. If we have the strict inequality $a_1 < a_2$, then the ray pairs represented by α_1 and α_2 are distinct sets. The sets $O_{\alpha_1}(b_1)$ and $O_{\alpha_2}(b_2)$ are simply connected sets that share no boundary in $\mathbb{C}$. The real line between a_2 and b_1 does not contain a_1 by the inequality, i.e. there is a path between $b_1 \in O_{\alpha_1}(b_1)$ and the boundary of $O_{\alpha_2}(b_2)$ which does not cross the boundary of $O_{\alpha_1}(b_1)$. We must conclude that the claim is again true. A nearly identical proof shows that if $b_1 \geq b_2$, then we have $O_{\beta_2}(a_2) \subset O_{\beta_1}(a_1)$.

Now suppose $\Delta_2 \subset \Delta_1$, then we must have $a_1 \leq a_2$ and $b_1 \geq b_2$. But then by the above claims, we get

$$\Delta_2 = O_{\alpha_2}(b_2) \cap O_{\beta_2}(a_2) \quad \subset \quad O_{\alpha_1}(b_1) \cap O_{\beta_1}(a_1) = \Delta_1.$$

Similar proofs hold for containment properties (1) and (3). Q.E.D.

There is also the issue of whether extended interval regions map to other extended interval regions under iteration by f. This next lemma together with the last lemma constitute most of the work involved in showing our extended first return map has the necessary Markov properties to be a holomorphic box mapping.

Lemma 8.6. Let Δ_1 be a real interval with repelling preperiodic boundary points under $f \in \mathcal{F}$. Fix an integer k, and suppose $\overline{\Delta_1}$ contains no critical point of $f^{\circ k}$. Let $\Delta_2 = f^{\circ k}(\Delta_1)$. Then, if Δ_1 and Δ_2 are considered as extended interval regions, it is still true that $\Delta_2 = f^{\circ k}(\Delta_1)$.

PROOF. Let us label the boundary points as $\Delta_1 = (a_1, b_1)$ and $\Delta_2 = (a_2, b_2)$. The boundary of Δ_1 must be mapped to the boundary of Δ_2, but this can be done in an orientation preserving or orientation reversing manner. We assume that

$$f^{\circ k}(a_1) = a_2 \qquad \text{and} \qquad f^{\circ k}(b_1) = b_2,$$

and a similar proof will work for the orientation reversing case. Let us label the left boundary of the extended region Δ_1, the rightmost ray pair landing at a_1, as $\mathcal{R}_{\alpha_1}$ and $\mathcal{R}_{-\alpha_1}$. Similarly, label the left boundary ray pair of Δ_2 as $\mathcal{R}_{\alpha_2}$ and $\mathcal{R}_{-\alpha_2}$. We must show only that $f^{\circ k}(\mathcal{R}_{\alpha_1} \cup \mathcal{R}_{-\alpha_1}) = \mathcal{R}_{\alpha_2} \cup \mathcal{R}_{-\alpha_2}$.

An identical argument shows that $f^{\circ k}$ maps the right boundary of Δ_1 to the right boundary of Δ_2. Then since the boundaries of Δ_1 and Δ_2 correspond under $f^{\circ k}$, and since the interior of Δ_1 must map into the interior of Δ_2 because they do so as real intervals, we must have the desired equation on extended interval regions, $\Delta_2 = f^{\circ k}(\Delta_1)$.

Notice that, since the first k points in the orbit of a_1 do not hit a critical point of f, locally $f^{\circ k}$ looks like $a_2 + C(x - a_1)$, a rotation and dilation followed by a translation. Furthermore, the constant C must be both real and positive, since f is real and we are assuming $f^{\circ k}$ is orientation preserving on Δ_1. Thus, locally $f^{\circ k}$ preserves the upper and lower half planes and the *left to right order* of any ray pairs landing at a_1. Now suppose that $f^{\circ k}(\mathcal{R}_{\alpha_1} \cup \mathcal{R}_{-\alpha_1}) \neq \mathcal{R}_{\alpha_2} \cup \mathcal{R}_{-\alpha_2}$. Since the image of any ray under an iterate of f is another ray, this equation implies there are multiple ray pairs landing at a_2. Moreover, there must be corresponding ray pairs landing at a_1. In other words, there must be a different ray pair, $\mathcal{R}_{\beta_1} \cup \mathcal{R}_{-\beta_1}$, which by definition of the extended region Δ_1 is not the rightmost pair landing at a_1, that is mapped by $f^{\circ k}$ onto $\mathcal{R}_{\alpha_2} \cup \mathcal{R}_{-\alpha_2}$, which is the rightmost pair landing at a_2. This violates the local orientation preserving properties of $f^{\circ k}$ and is a contradiction. Q.E.D.

Constructing holomorphic box mappings. The two boundary points of the inducing interval I are not in the domain of Φ, considered as a real map by definition. This fact has two consequences. First, any interval in the domain of Φ must be either inside or outside of I, and second, the image of each interval under Φ is inside of I by Lemma [8.3]. So given any component of the domain of Φ and the injective image of another component, either one set must contain the other, or they are totally disjoint. This is the desired Markov property of box mappings. By combining this real line property with the properties of extended interval regions, Lemmas [8.6] and [8.5], we construct a complex version of Φ, which satisfies the conditions necessary to be a holomorphic box mapping. Lemmas [8.6] and [8.5] can be considered as justification for the extended definition of Φ.

The map Φ restricted to each component of its domain on the real line will be extended as a holomorphic map to the bounded interval region determined by the domain component. Let r be the universal radius, determined by the bone-loop containing f, of the potential curve that is bounded away from 0 and thus well inside the basin at infinity of f. We start by taking the extended interval region of the inducing interval I and cutting it off at the potential curve G^r. This set is a bounded interval region which we label I^r.

Definition. (Φ extended from injective components) Suppose Δ is a connected component of the domain of the real Φ on which Φ is injective. We know that Φ restricted to Δ is equal to $f^{\circ k}$ for some k. Define Φ on the extended domain Δ^s, where $s = r^{\frac{1}{3}^k}$, as $f^{\circ k}$ considered as a holomorphic map. Note that Φ is univalent on Δ^s, and the image of Δ^s is contained in I^r.

If a central domain exists, we do not merely extend the central interval B. We need to make special arrangements so that the image of the extended central domain under Φ is I^r. Suppose the central branch of the real box mapping is $\Phi|_B = f^{\circ k}$. Since the critical point c returns to I, the critical value $f(c)$ must be in the domain of Φ. In particular it must lie in some monotone domain Δ, which has an extension Δ^s, $s = r^{\frac{1}{3}^{(k-1)}}$, on which the holomorphic Φ is already defined. Furthermore, since the image of B under f is connected and contained in the domain of the real Φ, it must be contained in the single connected component Δ, and the image of the boundary points of B under f must be a single boundary point of Δ. In terms of the ray pairs bounding the extended interval regions, both ray pairs bounding B are mapped by f to one of the ray pairs bounding Δ. Therefore, the image of the extended interval region of B is the open set bounded by this ray pair and containing Δ, and Δ^s.

Definition. (Φ extended from the central component) The domain of the central branch is defined to be the open set $f^{-1}(\Delta^s)$, which is contained in, but is not all of, the bounded interval region B^r. This new set still contains the entire original interval B however. The set $f^{-1}(\Delta^s)$ will still be called the extended central component B, although this does not match the definition of an extended interval region. On the extended B, Φ is defined to be $f^{\circ k}$ considered as a holomorphic map.

This completes the definition of Φ. The image of any point under Φ not contained in the extended central components or one of the extended univalent components is undefined.

Proposition 8.7. The extended map Φ is a holomorphic box mapping.

PROOF. The domain of Φ is the union of bounded interval regions derived from the connected components of the domain of the original real version of Φ. These extended regions are open disks, thus the entire domain of definition U is an open set. By Lemma [8.3], all branches of the real map $\Phi|_{\mathbb{R}}$ are homeomorphic onto their images, except possibly for the central branch mapping the domain containing the chaotic critical point c. Suppose Δ^s is the bounded interval region extended from an interval on which $\Phi|_{\mathbb{R}}$ is monotone. Then Δ^s can contain no critical point of f and is mapped univalently onto another bounded interval region by Lemma [8.6]. Further iterations under f of this bounded interval region, until I^r is reached, are all mapped univalently by f because the underlying real interval contains no critical point of f and no critical point of f exists off of the real line. Therefore, all the branches of the extended Φ, except the central branch, are univalent onto their images.

By the extension construction and Lemma [8.3], the image of any univalent branch $f^{\circ k_1}$ is the bounded interval region of I, I^r. The image of the central branch B, if it exists, is also I^r by construction. We must check, for the Markov properties required of holomorphic box mappings, that any other domain component Δ^s is either contained in I^r or totally disjoint from it. The interval $\Delta^s \cap \mathbb{R}$ is either contained in or disjoint from the interval I because $\Delta^s \cap \mathbb{R}$ is entirely contained in the domain of $\Phi|_{\mathbb{R}}$ and the boundary points of I are not in this domain. By Lemma [8.5], the extended interval regions Δ and I share the same containment property. The same property can then by shown to hold for the bounded interval regions Δ^s and I^r by noting that the *cut off* for Δ^s is G^s which is contained in the cut off for I^r,

G^r. The Markov property holds for the central domain B also because, by definition, B is contained in the bounded interval region B^r, and the same reasoning shows that B^r must be contained in I^r.

We must also check that $\Phi|_{B-\{c\}}$ is a two to one cover onto its image. It suffices to show only that f is two to one on $B^r - \{c\}$, since $B = f|_{B^r}^{-1}(\Delta^s)$ for some region Δ^s on which $\Phi = f^{\circ(k-1)}$ is univalent. To show that f is two to one on $B^r - \{c\}$, we first show that this is true when we restrict our attention to the real line. Suppose $B^r \cap \mathbb{R} = (a, b)$. By Lemma[8.4], $\Phi|_{\mathbb{R}}$ is two to one on $(a, b) - \{c\}$, $f(a) = f(b)$, and hence the map f must also be two to one on $(a, b) - \{c\}$.

Now let $\mathcal{R}_\alpha$ and $\mathcal{R}_{-\alpha}$ be the rightmost ray pair landing at a, and let $\mathcal{R}_\beta$ and $\mathcal{R}_{-\beta}$ be the leftmost ray pair landing at b. We claim that these two ray pairs have the same image under f. This follows from the same argument used in Lemma [8.6]. Since a and b are on opposite sides of the critical point c, if f reverses the orientation of $\mathbb{R}$ locally near a then it must preserve the orientation locally near b. In this case the rightmost ray pair landing at a must be mapped by f to the leftmost pair landing at $f(a)$, and the leftmost ray pair landing at b must be mapped by f to the same leftmost ray pair landing at $f(b) = f(a)$. An identical argument can be made if f is locally orientation preserving near a and orientation reversing near b. In either case, $f(\mathcal{R}_\alpha \cup \mathcal{R}_{-\alpha}) = f(\mathcal{R}_\beta \cup \mathcal{R}_{-\beta})$.

It now follows immediately that $f(B^r \cap \mathbb{H}^+)$ and $f(B^r \cap \mathbb{H}^-)$ have the same boundary, and thus represent the same set, which we label V. This set is bounded by the ray pair $f(\mathcal{R}_\alpha \cup \mathcal{R}_{-\alpha}) = f(\mathcal{R}_\beta \cup \mathcal{R}_{-\beta})$ and the potential curve G^{r^3}. We obtain V from the entire set bounded by these curves by removing the part of the real line between the critical value $f(c)$ and the landing point of the ray pair. Inspection reveals that V is simply connected, and this shows that $B^r \cap \mathbb{H}^+$ and $B^r \cap \mathbb{H}^-$, which contain no critical point of f, are both mapped univalently by f onto V. This finishes the proof that f is two to one because any point in V has exactly one preimage in $B^r \cap \mathbb{H}^+$ and one in $B^r \cap \mathbb{H}^-$. Q.E.D.

Corollary. The restriction of Φ to the real line is a real box mapping.

PROOF. The restriction of the extended Φ to the real line is of course just the original first return map. Almost all the properties of real box mappings are inherited from Φ being a holomorphic box mapping. The monotone and containment properties of intervals in the domain of Φ are inherited from the corresponding interval regions. Actually the proof that Φ is a holomorphic box mapping shows that extended interval regions have the necessary properties by first showing that property for the underlying interval.

Thus all we have left to prove is that branches of $\Phi|_{\mathbb{R}}$ have non-positive Schwarzian derivative. The polynomial f is cubic and bimodal, so the first derivative of f will be quadratic with two distinct real roots. Polynomials with such property always have a negative Schwarzian derivative. Compositions of functions with a negative Schwarzian also have a negative Schwarzian. All branches of a box mapping induced from f are compositions of f with itself, so the branches must have a negative Schwarzian derivative. Q.E.D.

§8.3. Constructing branchwise equivalences.

Suppose f and $\hat{f} \in \mathcal{F}$ are two maps on the same bone-loop. We know there is a real conjugacy between them. Suppose these maps each have a corresponding inducing interval, I^r and $\hat{I}^r$ respectively, which are mapped one onto the other by the conjugacy. In the same way, the construction of the first return maps to $I|_{\mathbb{R}}$ and $\hat{I}|_{\mathbb{R}}$, will parallel each other. In other words, any maximal interval on which Φ is defined as $f^{\circ i}$, will be mapped to the corresponding domain of $\hat{\Phi}$, where $\hat{\Phi}$ is defined as $\hat{f}^{\circ i}$. The conjugacy between f and $\hat{f}$ is also a conjugacy between Φ and $\hat{\Phi}$, since conjugacies are preserved under iteration, and clearly this conjugacy is a branchwise equivalence between Φ and $\hat{\Phi}$ as well. Establishing the correspondence between the connected domains of Φ and $\hat{\Phi}$, considered as box mappings is thus straightforward and follows naturally from the fact that the conjugacy between f and $\hat{f}$ is always a branchwise equivalence.

Of course the ultimate goal is to show the conjugacy can be extended to a quasiconformal conjugacy, and to do this, we need to induce from an initial branchwise equivalence which is quasiconformal. The conjugacy itself is thus not a suitable starting point. We need to find a branchwise equivalence in the same similarity class, which is also quasiconformal. From this starting point, the inducing algorithm constructs a quasiconformal conjugacy.

To find this initial quasiconformal equivalence, we will use the inductive construction in Lemma [7.1] to reproduce the first return maps, Φ and $\hat{\Phi}$. The lemma will then produce the desired branchwise equivalence between our initial box maps.

Inducing branchwise equivalences. The constant r and $s = r^{\frac{1}{3}}$ will be present throughout the construction. This is the same r used in the construction of the extended Φ and $\hat{\Phi}$. It is a universal radius used to specify the potential curve G^r, which exists for any map f in the neighborhood of the bone-loop. The existence of such a radius is proved by Lemma [7.10] and Lemma [7.12].

Our goal here is to construct pairs of holomorphic box mappings with no central branch and a branchwise equivalence that will satisfy the conditions of Lemma [7.1].

Recall the notation for components of the attracting basin of f on a bone-loop, $\{F_i\}$. Suppose f and $\hat{f}$ are two maps on the same bone-loop, then they are conjugate maps. Therefore both maps have corresponding basins for a single super attracting periodic orbit. We define $\hat{F}_i$ to be the interval corresponding to F_i under the conjugacy between f and $\hat{f}$. The end-points of all these intervals are preperiodic points, so we can extend these intervals into the plane as bounded interval regions. We cut off these interval regions at the potential curves G^s and $\hat{G}^s$ respectively as boundaries, producing the family of bounded interval regions, $\{F_i^s\}$ and $\{\hat{F}_i^s\}$.

We prepare to construct the map ϕ, and $\hat{\phi}$ will be constructed similarly. There are a number of *marked* points on the real line, depending on the map f, which we will use to construct the rays constituting the initial domain boundaries of ϕ. The first return map Φ is constructed from f, and the boundary points of the central domain B of this map, if it exists, are the first marked points. We consider the entire forward orbit of each of these points and the boundary points of each F_i and then I itself as marked. These boundary points are all preperiodic, so the set of marked points is finite. Finally we mark the repelling periodic points 0 and 1, and then collect the finite number of ray-pairs landing at the marked points into a collection $\mathcal{R}_*$. These rays, their landing points, and the potential curve G^s partition the plane into a finite number of pieces. Those intersecting the real line look like bounded interval regions Δ^s, including the regions $\{F_i^s\}$. We are ready to define ϕ.

Definition. (ϕ) The domain of ϕ is the union of all domains, Δ^s, that intersects $\mathbb{R}$ between 0 and 1, unless that domain is one of the $\{F_i^s\}$ or the bounded interval region B^s containing the extended central domain B. Then $\phi(z)$, for z in its domain, is defined to be $f(z)$.

Definition. (ϕ_0) The definition of ϕ_0 is identical to that of ϕ, except we define ϕ_0 to be the identity on any component of the domain inside I^s.

Definition. ($\mathcal{A}$) The set $\mathcal{A}$ is the union of all the rays in $\mathcal{R}_*$, their landing points, and all points outside, and including, the potential curve G^s.

There are parallel definitions for $\hat{\phi}$, $\hat{\phi}_0$, and $\hat{\mathcal{A}}$.

Definition. (h) On the set $\mathcal{A}$ we define $h = \hat{\varphi}^{-1} \circ \varphi$. Where φ and $\hat{\varphi}$ are the Böttcher coordinates of f and $\hat{f}$ respectively. By Theorem [7.14] and its corollary, since $\mathcal{A}$ includes only a finite number of rays and a collection of potential curves with radius bounded from below by s, this map can be extended to a symmetric quasiconformal homeomorphism of the plane preserving the orientation of the real line.

The dilatation of the induction process, K, is now defined to be the bound on the complex dilatation of h.

Lemma 8.8. The maps ϕ, ϕ_0, $\hat{\phi}$, $\hat{\phi}_0$ are all symmetric holomorphic box mappings with no central branch. Furthermore h is a branchwise equivalence between ϕ and $\hat{\phi}$ and between ϕ_0 and $\hat{\phi}_0$.

PROOF. The domain of ϕ is an open set consisting of connected components which are topological disks. The map ϕ is equal to f where it is defined, but the regions of f containing critical points have been removed from the domain of ϕ. Restricting ϕ to the remaining components of the domain must therefore yields univalent maps.

To prove that ϕ is a box mapping with no central branch, we must show that images of domain components either contain or are disjoint from other domain components. Let V be the image of some component and let Δ^s be some other component. First examine the restriction of these sets to the real line. The interval region Δ^s intersects $\mathbb{R}$ in an interval whose boundaries are two of the finite preperiodic ray landing points found in $\mathcal{A}$. There can be no other real point of $\mathcal{A}$ inside this interval because all such points are not in the domain of ϕ by definition. On the other hand, the boundary points of $V \cap \mathbb{R}$ must also be in $\mathcal{A}$ because by definition they are in $f(\mathcal{A})$ and $\mathcal{A}$ is forward invariant by construction. But it is not necessarily true that this interval contains no other real point of $\mathcal{A}$. There can be no partial overlap of $V \cap \mathbb{R}$ and $\Delta^s \cap \mathbb{R}$ or else Δ^s would contain a point of $\mathcal{A}$. Therefore as intervals V must either contain or be disjoint from Δ^s. The set Δ^s is bounded by the potential with radius s by definition, while V is bounded by radius r. Therefore, by Lemma [8.5], V either contains or is disjoint from Δ^s considered as bounded interval regions. This proves ϕ is a box mapping with no central branch.

Since ϕ_0 has the same domain components as ϕ, the proof for ϕ_0 is exactly the same. The image V may be the image of some Δ^s under the identity map rather than f, but this does not add any complications. Similar proofs also show $\hat{\phi}$ and $\hat{\phi}_0$ are box mappings with no central branch. The branches from all these maps are iterates of f or $\hat{f}$, or they are identity maps. In any case, all four maps are symmetric.

The boundaries of domain components of ϕ and ϕ_0 make up $\mathcal{A}$ and are mapped by h onto $\hat{\mathcal{A}}$, the boundaries of components of $\hat{\phi}_0$. Therefore the components of the complement of $\mathcal{A}$ are mapped onto the components of the complement of $\hat{\mathcal{A}}$. Since h respects the conjugacy between f and $\hat{f}$ on $\mathcal{A}$, we must have F_i^s mapping to $\hat{F}_i^s$, for each i, and the extended central domain B^s mapping to $\hat{B}^s$. By the pigeon hole principal, we must have domain components of ϕ mapping to domain components of $\hat{\phi}$. Notice that any branch of ϕ is equal to f, and any branch of $\hat{\phi}$ is equal to $\hat{f}$. Let $z \in \mathcal{A}$ be an arbitrary boundary point of the domain of ϕ. By applying the basic functional equation for Böttcher coordinates, we obtain

$$h \circ f(z) = \hat{\varphi}^{-1} \circ \varphi \circ f(z) = \hat{\varphi}^{-1}([\varphi(z)]^3) = \hat{f} \circ \hat{\varphi}^{-1} \circ \varphi(z) = \hat{f} \circ h(z).$$

This immediately implies the desired equations for branchwise equivalences. The same proof also works for ϕ_0 and $\hat{\phi}_0$, but in some cases, branches equal to f and $\hat{f}$ may be replaced by the identity function. The same equation is obviously still true here. This shows h is a branchwise equivalence between both pairs of maps. Q.E.D.

We can build two sequences of functions, $\{\phi_i\}$ and $\{\hat{\phi}_i\}$. Where $\phi_1 = \phi$, $\hat{\phi}_1 = \hat{\phi}$,

$$\phi_{i+1}(z) = \phi_0 \circ \phi_i(z) \qquad \text{and,}$$

$$\hat{\phi}_{i+1}(z) = \hat{\phi}_0 \circ \hat{\phi}_i(z).$$

Notice that at each step of the induction, the domain of ϕ_{i+1} or $\hat{\phi}_{i+1}$ may shrink because the composition may not be defined. In view of Lemma [8.8], by Lemma [7.1], these are all symmetric holomorphic box mappings with no central branch, and there exists a sequence, $\{h_i\}$, of K-quasiconformal branchwise equivalences, which are symmetric and preserve the orientation of $\mathbb{R}$.

The limiting branchwise equivalence. We check that the sequence $\{\phi_i\}$ converges to a branchwise equivalence.

Lemma 8.9. For z in the domain of the first return map Φ of f, but not in the central domain, $\phi_i(z)$ converges to $\Phi(z)$. Otherwise, the $\{\phi_i(z)\}$ have no limit or become undefined.

PROOF. First notice that $\Phi(z)$ and $\lim_{i\to\infty} \phi_i(z)$ agree for $z \in \mathbb{R}$. Let z be an arbitrary point in the interval $[0, 1]$, and let us follow the sequence

$\{\phi_i(z)\}_{i=0}^{\infty}$. In general, $\phi_0(z) = \phi_1(z) = f(z)$, and starting from here, the sequence can be given by,

$$\phi_{i+1}(z) = \phi_0 \circ \phi_i(z) = f(\phi_i(z)) = f^{\circ i}(z),$$

until a terminating condition is reached, either when ϕ_0 becomes undefined or becomes the identity map. This sequence is thus equal to the orbit of z under f until that orbit hits the interior of the real interval I, in which case the sequence becomes fixed because ϕ_0 is the identity, or until the orbit hits $\mathcal{A}$ or some F_i, in which case the sequence becomes undefined. The value of $\Phi(z)$ is also determined by following the orbit of z under f. $\Phi(z)$ is defined if the orbit of z hits the interior of I before any F_i, the boundary of any F_i, or the boundary of I. This is almost exactly the same condition, since $\mathcal{A} \cap \mathbb{R}$ contains the boundary points of these intervals. Unfortunately $\mathcal{A}$ contains additional points from the forward orbits of the boundary of B, the central domain.

Suppose the orbit of z hits the interior of I before any point in $\mathcal{A} \cap \mathbb{R}$, then clearly we have $\Phi(z) = \lim_{i\to\infty} \phi_i(z)$. Suppose the orbit of z hits a point in $\mathcal{A}$ before returning to the interior of I. Then, either the orbit hits a boundary point of some F_i or I, or it hits some point in the orbits of the boundary points of B. In either case, the sequence $\phi_i(z)$ becomes undefined by definition because ϕ_0 is undefined here. Since the orbits of the boundary of B hit the boundary points of the interval I or some F_i by definition of B, $\Phi(z)$ is still undefined for both cases. We have shown

$$\Phi|_{\mathbb{R}} = \lim_{i\to\infty} \phi_i|_{\mathbb{R}}.$$

Let Δ be some interval in the common domain. Suppose we have $\Phi|_\Delta = \lim_{i\to\infty} \phi_i|_\Delta = f^{\circ j}$. Then by definition of the extended map, Φ is defined on Δ^t, $t = r^{\frac{1}{3}^j}$. At each stage of the induction, the branches of ϕ_i also look like $f^{\circ j}$, and domain components are bounded interval regions equal to the pull-back of some Δ^s through $f^{\circ(j-1)}$. These domains have radius $t = s^{\frac{1}{3}^{(j-1)}} = r^{\frac{1}{3}^j}$. So the limit of the sequence $\{\phi_i\}$ must also be defined on Δ^t. Since every domain component of Φ or the $\{\phi_i\}$ fall into this case, except the central domain B, we are done. Q.E.D.

We have shown that the only place where $\lim_{i\to\infty} \phi_i(z)$ does not equal $\Phi(z)$ is for points in the extended central domain, and the same holds for $\lim_{i\to\infty} \hat{\phi}_i(z)$ and $\hat{\Phi}(z)$. Is is suitable to define the maps

$$\phi_\infty(z) = \lim_{i\to\infty} \phi_i(z) \qquad \text{and} \qquad \hat{\phi}_\infty(z) = \lim_{i\to\infty} \hat{\phi}_i(z),$$

defined where the limits exist. Because of Lemma [8.9], both ϕ_∞ and $\hat{\phi}_\infty$ are holomorphic box mappings with no central branch. In particular, they are equal to Φ and $\hat{\Phi}$ with the central branches removed.

Lemma 8.10. There exists a branchwise equivalence between ϕ_∞ and $\hat{\phi}_\infty$, which is K-quasiconformal, symmetric, and preserves the orientation of the real line.

Proof. In view of Lemma [8.8], we can apply the results of Lemma [7.1] to the $\{\phi_i\}$ and $\{\hat{\phi}_i\}$. There must exist branchwise equivalences $\{h_i\}$ between each pair of maps $(\phi_i, \hat{\phi}_i)$, which are K-quasiconformal, symmetric, and preserve the orientation of $\mathbb{R}$. It natural to use the limit of the $\{h_i\}$ as a candidate for a branchwise equivalence between ϕ_∞ and $\hat{\phi}_\infty$. We do not however need to expend the effort to show the limit exists. The sequence of K-quasiconformal homeomorphisms forms a normal family, so we can pick an arbitrary convergent subsequence of $\{h_i\}$. We label this arbitrary limit, also a K-quasiconformal homeomorphism, $\tilde{H}$. The map $\tilde{H}$ is also symmetric and preserves the orientation of $\mathbb{R}$ because it is the limit of maps with the same properties. Furthermore, we show $\tilde{H}$ is the desired branchwise equivalence between ϕ_∞ and $\hat{\phi}_\infty$.

Let Δ^s be any domain of ϕ_∞. For large enough i, Δ^s is also a component of the domain of ϕ_i. By Lemma [8.9], for $z \in \Delta^s$, $\lim_{i\to\infty} \phi_i(z) = \Phi(z)$. The sequence $\{\phi_i(z)\}_{i=0}^{\infty}$ must become fixed at the stage in the orbit of z when it returns to I. But by definition, the step where $\{\phi_i(z)\}$ becomes fixed, is exactly where $\{h_i(z)\}$ becomes fixed. The sequence h_i must therefore have a limit at this point and $\tilde{H}(z)$ is equal to this limit. Therefore $\tilde{H}$ is equal to a branchwise equivalence on Δ^s and, by continuity, on the boundary of Δ^s as well. Thus $\tilde{H}$ maps the domain component Δ^s of ϕ_∞ onto the corresponding component of $\hat{\phi}_\infty$ and satisfies the functional equation of the boundary. Since all components of the domain of ϕ_∞, and $\hat{\phi}_\infty$, arise in this manner, $\tilde{H}$ must be a full branchwise equivalence itself Q.E.D.

Adding the central branch. The map $\tilde{H}$ constructed in Lemma [8.10] is almost the desired branchwise equivalence between Φ and $\hat{\Phi}$. We have only the issue of the central branch remaining. Recall that intervals containing the critical points of f and $\hat{f}$ were removed from consideration in the induction. However, the branchwise equivalence must also take the central domain of Φ

onto the central domain of $\hat{\Phi}$. It is possible to redefine $\tilde{H}(z)$ on the bounded interval region B^r so that it maps B, the central domain of Φ, onto the central domain of $\hat{\Phi}$, $\hat{B}$.

Theorem 8.11. There exists a branchwise equivalence between Φ and $\hat{\Phi}$, which is quasiconformal, symmetric, and preserves the orientation of the real line.

PROOF. The branchwise equivalence, $\tilde{H}$, between ϕ_∞ and $\hat{\phi}_\infty$ needs to be redefined on B^r. Recall that $\tilde{H}$ is the limit of a subsequence of the branchwise equivalences $\{h_i\}$ between the $\{\phi_i\}$ and the $\{\hat{\phi}_i\}$. We know that

$$\tilde{H}(z) = h_0(z) = h_i(z) \qquad \text{for } z \in B^r \cup \mathcal{A} \text{ and for } i \geq 0$$

because in this case $\phi_i(z)$ is undefined for all $i \geq 1$ and therefore $h_i(z)$ stays fixed. In particular, $\tilde{H}$ maps B^r onto $\hat{B}^r$, since h_0 does by definition.

Recall that f and $\hat{f}$ map their chaotic critical points into monotone domains of Φ and $\hat{\Phi}$, which we label Δ^s and $\hat{\Delta}^s$ respectively, and recall that the central domains, B and $\hat{B}$ are defined as $f^{-1}(\Delta^s)$ and $\hat{f}^{-1}(\hat{\Delta}^s)$ respectively. We label the ray pairs bounding B as $\mathcal{R}_a \cup \mathcal{R}_{-a}$ and $\mathcal{R}_b \cup \mathcal{R}_{-b}$, and the corresponding pairs bounding $\hat{B}$ as $\hat{\mathcal{R}}_a \cup \hat{\mathcal{R}}_{-a}$ and $\hat{\mathcal{R}}_b \cup \hat{\mathcal{R}}_{-b}$. $\mathcal{R}_a \cup \mathcal{R}_{-a}$ and $\mathcal{R}_b \cup \mathcal{R}_{-b}$ map under f to a single ray pair, $\mathcal{R}_c \cup \mathcal{R}_{-c}$, part of the boundary of Δ^s. Similarly $\hat{\mathcal{R}}_a \cup \hat{\mathcal{R}}_{-a}$ and $\hat{\mathcal{R}}_b \cup \hat{\mathcal{R}}_{-b}$ map under $\hat{f}$ to $\hat{\mathcal{R}}_c \cup \hat{\mathcal{R}}_{-c}$ bounding $\hat{\Delta}^s$. All these ray pairs are in $\mathcal{A}$ or $\hat{\mathcal{A}}$, so $\tilde{H}$ inherits from h_0 the property that $\mathcal{R}_a$ maps to $\hat{\mathcal{R}}_a$, $\mathcal{R}_b$ to $\hat{\mathcal{R}}_b$, $\mathcal{R}_c$ to $\hat{\mathcal{R}}_c$, etc.

The map $\tilde{H}$ must map Δ^s onto some domain component of $\hat{\Phi}$. This component must be $\hat{\Delta}^s$ because $\tilde{H}$ maps $\mathcal{R}_c$ bounding Δ^s to $\hat{\mathcal{R}}_c$ bounding

$\hat{\Delta}^s$. We also know $\tilde{H}$ maps the potential curve G^{r^3} to the corresponding curve $\hat{G}^{r^3}$. As $\mathcal{R}_c \cup \mathcal{R}_{-c} \cup G^{r^3}$ and $\hat{\mathcal{R}}_c \cup \hat{\mathcal{R}}_{-c} \cup \hat{G}^{r^3}$ form the boundary of $f(B^r)$ and $\hat{f}(\hat{B}^r)$ respectively, $\tilde{H}$ must map these sets onto each other, so we are in a position to pull $\tilde{H}$, defined on $f(B^r)$, back to a map defined on B^r.

This is a simple lifting situation, giving us a symmetric homeomorphism H between B^r and $\hat{B}^r$ satisfying

$$\tilde{H} \circ f(z) = \hat{f} \circ H(z) \qquad \text{for } z \in \bar{B^r}. \tag{1}$$

We know H is quasiconformal by the same argument used in lifting through univalent maps because f and $\hat{f}$ are conformal and $\tilde{H}$ is K-quasiconformal. We claim $\tilde{H}$ agrees with H on the boundary of B^r. Inheriting this from h_0, $\tilde{H}$ also satisfies

$$\tilde{H} \circ f(z) = \hat{f} \circ \tilde{H}(z) \qquad \text{for } z \in \partial B^r \subset \mathcal{A}.$$

Thus $\tilde{H}$ is also a lift through f and $\hat{f}$ of $\tilde{H}$ restricted to the boundary of $f(B^r)$. Since the lifting is unique, given that we require the lift to be orientation preserving on $\mathbb{R}$, we must have

$$H|_{\partial B^r} = \tilde{H}|_{\partial B^r}.$$

Thus if we extend H be defining

$$H(z) = \tilde{H}(z) \qquad \text{for } z \notin B^r,$$

we get a homeomorphism. Moreover the extended H is quasiconformal. We also claim that H maps the central domain $f^{-1}(\Delta^s)$ onto $\hat{f}^{-1}(\hat{\Delta}^s)$. From equation (1), we get $\hat{\Delta}^s = \tilde{H}(\Delta^s) = \hat{f} \circ H(B)$. So $H(B) = \hat{f}^{-1}(\hat{\Delta}^s) = \hat{B}$.

Thus H maps the domain of Φ, including the central domain, onto the domain of $\hat{\Phi}$. The map H is quasiconformal, symmetric, and preserves the orientation of $\mathbb{R}$. Furthermore H, by Lemma [8.10], is a branchwise equivalence outside of the central domain. All we have left to check is that H has the proper definition for a branchwise equivalence for points on the boundary of B.

By definition, $\Phi|_B = f^{\circ j}$ and $\hat{\Phi}|_{\hat{B}} = \hat{f}^{\circ j}$, for some integer j. Recall that Δ^s and $\hat{\Delta}^s$ contain the respective chaotic critical values of f and $\hat{f}$. We must have $\Phi|_{\Delta^s} = f^{\circ(j-1)}$ and $\hat{\Phi}|_{\hat{\Delta}^s} = \hat{f}^{\circ(j-1)}$. By definition H is equal to $\tilde{H}$ on Δ^s, a branchwise equivalence by Lemma [8.10], so we have

$$\hat{f}^{\circ(j-1)} \circ H(z) = H \circ f^{\circ(j-1)}(z) \qquad \text{for } z \in \partial\Delta^s. \tag{2}$$

For $z \in \partial B$, by definition $f(z)$ is in $\partial\Delta^s$, so combining equation (1), equation (2), and the definition of H, we obtain

$$\hat{f}^{\circ j} \circ H(z) = \hat{f}^{\circ(j-1)} \circ H(f(z)) = H \circ f^{\circ j}(z) \qquad \text{for } z \in \partial B.$$

This verifies that H satisfies the functional equation for branchwise equivalences on the boundary of B. Q.E.D.

IX. Quasiconformal conjugacies on bone-loops

In this chapter we finish the proof of Theorem [5.4]. This provides the final piece in the proof of the Connected Bone Conjecture and monotonicity for the real cubic family.

§9.1. Final filling.

The next proposition is the final step in the proof of Theorem [5.4] in all non-renormalizable cases. Box mappings with no central branch can be obtained in these cases, and pull-back immediately produces a quasiconformal map matching the conjugacy on the forward orbit of the chaotic critical point.

Proposition 9.1. Let ϕ and $\hat{\phi}$ be holomorphic box mappings with no central branch, induced from real polynomials f and $\hat{f}$ that are topologically conjugate on the real line. Let h be a branchwise equivalence, in the same similarity class as the conjugacy, that is K-quasiconformal, symmetric, and preserves the orientation of the real line.

Let J be the intersection of the Julia set of f with $\mathbb{R}$. Suppose that the only open intervals not intersecting the domain of ϕ are in the Fatou set of f. Finally suppose that ϕ is not equal to the identity on any component of its domain. Then there exists a quasiconformal homeomorphism H that is symmetric and preserves the orientation of the real line, and which satisfies

$$H \circ f(z) = \hat{f} \circ H(z),$$

for any z in J.

PROOF. With $\phi_0 = \phi_1 = \phi$ and $\hat{\phi}_0 = \hat{\phi}_1 = \hat{\phi}$, define a new sequence of maps

$$\phi_{i+1} = \phi_0 \circ \phi_i \qquad \text{and} \qquad \hat{\phi}_{i+1} = \hat{\phi}_0 \circ \hat{\phi}_i.$$

With initial branchwise equivalence h, Lemma [7.1] applies. All the $\{\phi_i\}$ and $\{\hat{\phi}_i\}$ are box mappings with no central branch, and between each pair $(\phi_i, \hat{\phi}_i)$ there exists a branchwise equivalence h_i, which is K-quasiconformal, symmetric, and preserves the orientation of $\mathbb{R}$. Let H be the limit of a convergent subsequence of $\{h_i\}$, which must exist as the K-quasiconformal maps form a normal family. The map H itself is K-quasiconformal, and it inherits symmetry and orientation from the the $\{h_i\}$. We claim H is equal to the conjugacy between f and $\hat{f}$ for every real point in the Julia set of f. The conclusion follows.

Since h is in the similarity class of the conjugacy, each of the h_i is also in the similarity class of the conjugacy. Therefore each h_i, on the boundary of its domain, is equal to the conjugacy. By the pull-back construction the only open intervals on which ϕ_i is undefined are inverse images of open intervals on which ϕ is undefined. These intervals are in the Fatou set of f, so intervals not intersecting the domain of ϕ_i are also in the Fatou set. So, if $\phi_i(z)$ becomes undefined for $z \in J$, z must still be in the closure of the domain of ϕ_i. Therefore, either z is the boundary point of some domain component of ϕ_i or there is a sequence of such boundary points converging to z. Either way since h_i is equal to the conjugacy on boundary points of components, h_i is equal to the conjugacy on z by continuity.

The final observation is that, in the limit, the ϕ_i become undefined everywhere. If there is some interval on which the box mappings stay defined, since the iterates of f defining the ϕ_i go to infinity, this interval would be contained in the Julia set and under iteration would never hit the critical point, a contradiction. Thus every point is the limit of points on which the h_i become fixed and equal to the conjugacy. Therefore the limit of the h_i H must agree with the conjugacy between f and $\hat{f}$ for every real point in the Julia set of f. Q.E.D.

§9.2. Non-renormalizable case.

Recall that the inducing algorithm for box mappings may produce a *terminal* box mapping if there is a restrictive interval I. In this case, for some box mapping, the central branch will map I into itself and the boundary of I into itself as well. If the box mapping is induced from f, the central branch is an iterate of f, and in the special case of a box mapping induced from a map on a bone-loop, we also know the orbit of I under f never intersects the periodic critical point because this point has been removed from the domain of the box mapping. This inspires the next definition.

Definition. (maximal restrictive intervals) Suppose $f \in mapfamily$ is on a bone-loop. A maximal restrictive interval I of f must be mapped by $f^{\circ j}$, for some j, into itself, and $f^{\circ j}(\partial I)$ must be contained in ∂I as well. Furthermore the interval I must contain the chaotic critical point of f, but neither I nor the forward orbit of I under f can contain the periodic critical point.

Cubic maps with this property are analogous to *renormalizable* quadratic maps, which also exhibit restrictive intervals. We will call a map on a bone-loop **renormalizable** if it has a maximal restrictive interval. We now finish the proof of Theorem [5.4] in the non-renormalizable case.

Theorem 9.2. Let f and $\hat{f} \in \mathcal{F}$ be two maps on the same bone-loop. Suppose f is not renormalizable. Then there exists a special quasiconformal conjugacy between f and $\hat{f}$.

Proof. We know there exists a topological conjugacy between f and $\hat{f}$ on the real line, by the corollary to Proposition [5.3]. We split the proof into two cases based on the recurrence of the chaotic critical point. If the chaotic critical orbit of f is non-recurrent and there exists an interval I_0 containing the critical point which the orbit otherwise never intersects, we attempt to turn this into an inducing interval. Since the chaotic critical point c is in the Julia set of f, we can certainly find real repelling preperiodic orbits close to the critical point. So given a finite set of repelling preperiodic orbits that intersect both pieces of $I_0 - \{c\}$. We take the two repelling preperiodic points, one on either side of c, closest to c as the boundary points of the inducing interval I. If the chaotic critical orbit is recurrent, the inducing interval is defined to be the largest open interval not intersecting the immediate basin of the periodic critical point. In the first case, the orbits of the boundary points do not intersect I by construction, and otherwise Lemma [8.1] proves the orbits do not intersect I. In either case we have an inducing interval I for f, and a corresponding interval $\hat{I}$ for $\hat{f}$, equal to the image of I under the conjugacy.

We can now construct the cubic first return maps for each polynomial and extend them to holomorphic box mappings, Φ and $\hat{\Phi}$, by invoking Proposition [8.7]. Then by Lemma [8.11], there exists a quasiconformal branchwise equivalence, H, between Φ and $\hat{\Phi}$, which is symmetric and preserves the orientation of $\mathbb{R}$. Notice that H is obtained by pull-back from a branchwise equivalence that is in the same similarity class as the conjugacy. So H and any branchwise equivalence obtained by pull-back from H is also the in the similarity class of the conjugacy and critically consistent as well. If the chaotic critical points are non-recurrent, the resulting box mappings will be without a central branch, and otherwise they certainly will have a central branch.

If there is no central branch, we can apply Proposition [9.1] at once to produce a quasiconformal homeomorphism that matches the conjugacy between f and $\hat{f}$ on the entire forward orbit of the chaotic critical point. In view of Theorem [4.7], a special qc-conjugacy between f and $\hat{f}$ can be constructed.

If there is a central branch, we must use the inducing algorithm. By Lemma [8.3], Φ and $\hat{\Phi}$ are *type II* box mappings, as all domain components have image equal to the image of the central domain. By the corollary to Proposition [8.7], the restrictions of these maps to the real line are real box mappings themselves. By restricting Φ to the bounded interval region I^r, where f is a two to one map, we can conjugate Φ, via some conformal map G, to a standard holomorphic box mapping. Similarly there exists a $\hat{G}$ conjugating $\hat{\Phi}$ with another standard box mapping. Conjugating Φ and

$\hat{\Phi}$ with conformal maps in this way preserves all the relevant properties of the box mappings. The new box mappings are still type II, and their restrictions to the real line are real box mappings. Furthermore, the map $\hat{G} \circ H \circ G^{-1}$ is a branchwise equivalence between the new box mappings, which is quasiconformal and critically consistent.

We can perform the inducing algorithm on these new mappings, and the obstructions to inducing, renormalization and non-recurrence, have both been ruled out for this particular case. The algorithm can therefore proceed to a limit, producing two infinite sequences of box mappings $\{\phi_i\}$ and $\{\hat{\phi}_i\}$. Theorem [6.2], and then Proposition [6.1], apply to these sequences. We conclude first that the modulus of $B_i' - B_i$, the annulus constructed by removing the central domain of ϕ_i from its image, grows linearly in i. After a finite number of steps, the modulus will be larger than the key value $4 \log 8$. From this step on, we will be able to find quasiconformal branchwise equivalences between ϕ_i and $\hat{\phi}_i$, whose dilatation grows at the rate $\exp(Q' \exp(-\frac{\epsilon}{4}))$, where Q' is a constant and ϵ grows linearly with the modulus. A brief analysis shows that the infinite product of these factors is still finite, so in the limit, we obtain a branchwise equivalence that is still quasiconformal.

The limit of the inducing algorithm on Φ cannot be a *terminal* box mapping. Such a situation can only happen if a restrictive interval around the chaotic critical point exists, and we have explicitly prevented this by requiring f to be non-renormalizable. The limit of inducing must produce two box mappings with no central branch and a quasiconformal branchwise equivalence between them. This translates back through G and $\hat{G}$ to box mappings again induced from f and $\hat{f}$. We apply Proposition [9.1] and then Theorem [4.7] again to produce the desired quasiconformal conjugacy. Q.E.D.

§9.3. Renormalizable case.

In the event that there does exist a maximal restrictive interval around the chaotic critical point, box mapping are not needed at all. In fact, the renormalization on the restrictive interval is polynomial-like (quadratic-like) in the sense of Douady and Hubbard, and Theorem [5.4] follows from the quadratic version, Theorem [1.1].

Definition. (polynomial-like mappings) A **polynomial-like** map of degree d is a triple (U, U', f) where U and U' are open subsets of $\mathbb{C}$ isomorphic to discs, with U' relatively compact in U, and $f : U' \to U$ a $\mathbb{C}$-analytic mapping, proper of degree d.

Also define K_f to be the set of points in U' whose orbits under f are completely contained in U'. This is a set analogous to the filled-in Julia set of a polynomial. The fundamental result of Douady and Hubbard, which incidently also uses quasiconformal pull-back techniques, is the Straightening Theorem (see [DH]).

Theorem 9.3. Every polynomial like mapping of degree d is conjugate, by a quasiconformal map defined on a neighborhood of K_f, to a polynomial P of degree d. Moreover, if K_f is connected, P is unique up to conjugation by an affine map.

Lemma 9.4. Let $f \in \mathcal{F}$ be on a bone-loop and have a maximal restrictive interval I. Then there exists an open disk U' in $\mathbb{C}$ so that $(U', f(U'), f)$ is a polynomial-like mapping of degree 2.

PROOF. Let $f^{\circ j}$ be the first return map of f to I. The boundary points must be preperiodic under f because $f^{\circ j}$ maps the boundary into itself. The boundary is also repelling, in the Julia set of f. So we can construct an extended interval region Δ whose intersection with the real line is I. We cut off Δ at any existing potential curve G^r, producing a bounded interval region Δ^r.

We claim that $f^{\circ j}(\Delta^r)$ contains Δ^r. There is only one fold on I of $f^{\circ j}$, so both boundary points must map to the same boundary point under $f^{\circ j}$. Therefore both ray pairs bounding Δ^r map to ray pairs landing at the same point, and in fact, they must map to the same ray pair. This shows that $f^{\circ j}$ is properly two to one on Δ^r. The rest of the boundary of $f^{\circ j}(\Delta^r)$ is the potential curve G^{r^3}, which does not intersect Δ^r. Since I maps to itself under $f^{\circ j}$, Δ^r must be contained in $f^{\circ j}(\Delta^r)$.

The set Δ^r is the natural candidate for the set U', but it is not quite compactly contained in $f^{\circ j}(\Delta^r)$. Instead we choose U' to be a set slightly larger. If the distance between the boundary of U' and Δ^r is chosen to be small enough everywhere, then U' will be compactly contained in $f^{\circ j}(U')$ because $f^{\circ j}$ is repelling near the boundary of Δ^r. Q.E.D.

We now conclude the proof of Theorem [5.4] for renormalizable maps.

Theorem 9.5. Let f and $\hat{f}$, both in $\mathcal{F}$, be on a bone-loop and have corresponding maximal restrictive intervals I and $\hat{I}$. Then there exists a special qc-conjugacy between f and $\hat{f}$.

PROOF. Let h_0 be a topological conjugacy of the real line between f and $\hat{f}$, which must exist because f and $\hat{f}$ are on the same bone-loop. Let $f^{\circ j}$ be the iterate of f which maps I into itself. Because the interval maps into itself but folds only once, there must be a finite sequence of intervals containing the forward orbit of I and that are all mapped onto I injectively by iterates of f. We label these $I_1 = I$ through I_n, and there must be corresponding intervals $\hat{I}_1$ through $\hat{I}_n$ where $\hat{f}^{\circ j}$ has the same behavior. The conjugacy h_0 restricted to each of the I_n also conjugates $f^{\circ j}$ with $\hat{f}^{\circ j}$ on $\hat{I}_n$. Both $f^{\circ j}$ and $\hat{f}^{\circ j}$ are bimodal maps on each of these intervals, and they inherit the chaotic critical point from f and $\hat{f}$ respectively.

By Lemma [9.4], we can construct open disks U_n' and $\hat{U}_n'$, containing I_n and $\hat{I}_n$ respectively, that make $f^{\circ j}$ and $\hat{f}^{\circ j}$ into quadratic-like maps. To each set we can apply Theorem [9.3] and construct a quasiconformal conjugacy H between $f^{\circ j}$ and some real quadratic map Q and another conjugacy $\hat{H}$ between $\hat{f}^{\circ j}$ and $\hat{Q}$.

The map Q has a unique fixed point and first preimage of the fixed point that must correspond under H with the boundary of I_n. Thus H maps I_n to the entire intersection of the Julia set of Q with the real line. The same situation is true for $\hat{H}$ and $\hat{Q}$ as well. The map $\hat{H} \circ h_0 \circ H^{-1}$ forms a real conjugacy between Q and $\hat{Q}$ defined on the interval between the fixed point and its first preimage. Therefore Q and $\hat{Q}$ are fully conjugate on the real line.

Q and $\hat{Q}$ inherit from $f^{\circ j}$ and $\hat{f}^{\circ j}$ a bounded critical orbit and no attracting cycles. Thus Theorem [1.1] applies, and there is an extension of $\hat{H} \circ h_0 \circ H^{-1}$ to a full quasiconformal conjugacy h between Q and $\hat{Q}$.

The map $\hat{H}^{-1} \circ h \circ H$ is quasiconformal, and it extends the conjugacy h_0 between $f^{\circ j}$ and $\hat{f}^{\circ j}$ on some neighborhood containing I_n. This same construction can be used to extend h_0 to a quasiconformal map on a neighborhood of each the I_n. By piecing these maps together we can construct a quasiconformal map which preserves the real line and its orientation. Furthermore this map matches the conjugacy h_0 on the entire chaotic critical orbit. Invoking Theorem [4.7] finishes the proof. Q.E.D.

References

[Ah] L. Ahlfors, *Lectures on Quasiconformal Mappings*, Wadsworth & Brooks/Cole, Belmont, 1987.

[AB] L. Ahlfors and L. Bers, *Riemann's mapping theorem for variable metrics*, Ann. of Math. **72** (1960), 385-404.

[BR] L. Bers and H. Royden, *Holomorphic families of injections*, Acta Math. **157** (1986), 259-286.

[Bl] P. Blanchard, *Complex analytic dynamics on the Riemann sphere*, Bull A.M.S. **11** (1984), 85-141.

[DGMT] S. Dawson, R. Galeeva, J. Milnor, and C. Tresser, *A monotonicity conjecture for real cubic maps*, preprint, IMS, Stony Brook.

[DH] A. Douady and J. Hubbard, *On the dynamics of polynomial-like mappings*, Ann. Sci. Ec. Norm. Sup. **18** (1985), 287-343.

[GS1] J. Graczyk and G. Swiatek, *Hyperbolicity in the real quadratic family*, preprint (1995).

[GS2] ———, *Induced expansion for quadratic polynomials*, Annales Scientifiques d'ENS (to appear).

[GJ] J. Guckenheimer and S. Johnson, *Distortion of S-unimodal maps*, Annals of Math. **132** (1990), 71-130.

[Ja1] M. Jakobson, *Absolutely continuous invariant measures for one parameter families of one-dimensional maps.*, Commun. Math. Phys. **81** (1981), 39-88.

[Ja2] ———, *Quasisymmetric conjugacies for some one-dimensional maps inducing expansion*, Contemporary Mathematics **135** (1992), 203-211.

[JS1] M. Jakobson and G. Swiatek, *Metric properties of non-renormalizable S-unimodal maps. Part I: Induced expansion and invariant measures*, Ergod. Th. & Dynam. Sys. **14** (1994), 721-755.

[JS2] ———, *Metric properties of non-renormalizable S-unimodal maps: II. Quasisymmetric conjugacy classes*, Ergod. Th. & Dynam. Sys. **15** (1995), 871-938.

[MSS] R. Mãñe, P. Sad, and D. Sullivan, *On the dynamics of rational maps*, Ann. Sci. Ec. Norm. Sup **16** (1983), 193-217.

[dMvS] W. de Melo and S. van Strien, *One-Dimensional Dynamics*, Springer, New York, 1993.

[Mi1] J. Milnor, *Remarks on iterated cubic maps*, Experimental Math., vol. 1, 1992, pp. 5-24.

[Mi2] ———, *Dynamics in one complex variable*, preprint (1990), IMS, Stony Brook.

[MT] J. Milnor and W. Thurston, *On iterated maps of the interval* I, II, Lecture Notes in Mathematics, Vol. 1342, Springer, Berlin, 1988, p. 465.

[Su1] D. Sullivan, *On the structure of infinitely many dynamical systems nested inside or outside a given one*, preprint IHES/M/90/75.

[Su2] ———, *Quasiconformal homeomorphisms and dynamics* I. *Solution of Fatou-Julia problem on wandering domains*, Annals Math. **122** (1985), 401-418.

[ST] D. Sullivan and W. Thurston, *Extending holomorphic motions*, Acta Math. **157** (1986), 243-257.

[Sw] G. Swiatek, *Hyperbolicity is dense in the real quadratic family*, preprint (1992), IMS, Stony Brook.

§10.1. Dilatation Lemma.

We prove a property of quasiconformal maps, crucial to pull-back constructions, that composition with a conformal map does not change the global bound on the complex dilatation.

Lemma 10.1. Let f and g be maps from $\hat{\mathbb{C}}$ to itself. If g is conformal then

$$|\chi_{f\circ g}(z)| \;=\; |\chi_f(g(z))|\,.$$

And if f is conformal then

$$|\chi_{f\circ g}(z)| \;=\; |\chi_g(z)|\,.$$

PROOF. The partial derivatives of a composition satisfy

$$\frac{\partial(f\circ g)}{\partial z} \;=\; \left.\frac{\partial f}{\partial z}\right|_{g(z)}\frac{\partial g}{\partial z} \;+\; \left.\frac{\partial f}{\partial \bar z}\right|_{g(z)}\frac{\partial \bar g}{\partial z}, \qquad \text{and}$$

$$\frac{\partial(f\circ g)}{\partial \bar z} \;=\; \left.\frac{\partial f}{\partial z}\right|_{g(z)}\frac{\partial g}{\partial \bar z} \;+\; \left.\frac{\partial f}{\partial \bar z}\right|_{g(z)}\frac{\partial \bar g}{\partial \bar z}.$$

If g is conformal then $\partial g/\partial \bar z$ and $\partial \bar g/\partial z$ are zero and $|\partial g/\partial z| = |\partial \bar g/\partial \bar z|$. Thus by the definition of $\chi_{f\circ g}$ and the above identities

$$|\chi_{f\circ g}(z)| \;=\; \left|\frac{\left.\frac{\partial f}{\partial \bar z}\right|_{g(z)}\frac{\partial \bar g}{\partial \bar z}}{\left.\frac{\partial f}{\partial z}\right|_{g(z)}\frac{\partial g}{\partial z}}\right| \;=\; \left|\frac{\left.\frac{\partial f}{\partial \bar z}\right|_{g(z)}}{\left.\frac{\partial f}{\partial z}\right|_{g(z)}}\right| \;=\; |\chi_f(g(z))|\,.$$

If f is conformal then $\partial f/\partial \bar{z}$ is zero, yielding

$$|\chi_{f\circ g}(z)| = \left| \frac{\left.\frac{\partial f}{\partial z}\right|_{g(z)} \frac{\partial g}{\partial \bar{z}}}{\left.\frac{\partial f}{\partial z}\right|_{g(z)} \frac{\partial g}{\partial z}} \right| = \left| \frac{\frac{\partial g}{\partial \bar{z}}}{\frac{\partial g}{\partial z}} \right| = |\chi_g(z)|.$$

Q.E.D.

It easy to see that any bound on the dilatation of a quasiconformal map is preserved by composition with a conformal map.

§10.2. Analytic dependence on parameters.

We provide a proof of Theorem [2.3]. First a preliminary lemma.

Lemma 10.2. Let A be an open set in $\hat{\mathbb{C}}$ and D the open unit disk. Suppose $f_i(z, \lambda) : A \times D \to \hat{\mathbb{C}}$ are a sequence of functions, analytic in either variable, z or λ, which omit at least three points of $\hat{\mathbb{C}}$ in their range, i.e. $\{0, 1, \infty\}$. Further suppose that for each $\lambda \in D$, the sequence of functions, considered as functions of z, converge uniformly to a limiting function $\varphi_\lambda(z)$. Then for fixed $z \in A$, $\varphi_\lambda(z)$ is analytic in λ.

PROOF. We use the notation $f_i^z(\lambda) = f_i(z, \lambda)$ and $\varphi_z(\lambda) = \varphi_\lambda(z)$ when we want to stress that z is being held fixed.

First notice that the family $f_i^z : D \to \hat{\mathbb{C}}$, $z \in A$, $1 \leq i < \infty$, satisfies the conditions of Montel's Theorem, and so the $\{f_i^z\}$, and $\{\varphi_z\}$, form a normal family. By Arzela's Theorem, this family is equicontinuous as well. Suppose we are given $\epsilon > 0$, then there exists a $\delta > 0$ so that if $|\lambda_1 - \lambda_2| < \delta$, we have

$$|f_i^z(\lambda_1) - f_i^z(\lambda_2)| < \epsilon/3 \qquad \text{for } z \in A \text{ and } 1 \leq i < \infty, \text{ and} \qquad (1)$$

$$|\varphi_z(\lambda_1) - \varphi_z(\lambda_2)| < \epsilon/3 \qquad \text{for } z \in A. \tag{2}$$

Now let $\{\lambda_j\}_{j=1}^{n}$ be a finite set of points in D, chosen so that the δ-balls centered at the λ_j form a finite cover of D. By hypothesis, for any j between 1 and n, there exists an integer N_j so that

$$|f_i^z(\lambda_j) - \varphi_z(\lambda_j)| < \epsilon/3 \qquad \text{for } i > N_j. \tag{3}$$

Define M to be the maximum of the integers $\{N_j\}_{j=1}^{n}$. Now suppose z is fixed in A, and let λ be any point in D. The point λ must be within δ of some λ_j, therefore equations (1) and (2) apply. If $i > M$, then equation (3) applies as well for this particular λ_j, and we get

$$\begin{aligned}|f_i^z(\lambda) - \varphi_z(\lambda)| &= |f_i^z(\lambda) - f_i^z(\lambda_j) + f_i^z(\lambda_j) - \varphi_z(\lambda_j) + \varphi_z(\lambda_j) - \varphi_z(\lambda)| \\ &\le |f_i^z(\lambda) - f_i^z(\lambda_j)| + |f_i^z(\lambda_j) - \varphi_z(\lambda_j)| + |\varphi_z(\lambda_j) - \varphi_z(\lambda)| \\ &< \epsilon/3 + \epsilon/3 + \epsilon/3 = \epsilon.\end{aligned}$$

Since M was choosen independent of $\lambda \in D$, the conclusion is that $\{f_i^z\}_{i=1}^{\infty}$ converges uniformly, as functions of λ, to $\varphi_z(\lambda)$. By Weierstrass's Theorem, φ_z must be an analytic function of λ. Q.E.D.

Proof of Theorem [2.3]. Define $F_\lambda(z)$ to be the lift of $f_\lambda(z)$ through the covering map e^z onto U. Also define

$$\Phi_\lambda(z) = \lim_{k\to\infty} \frac{F_\lambda^{\circ k}}{n^k}.$$

It is the contention of the Böttcher Theorem itself (See [Mi2]) that these lifts exist and the sequence converges uniformly to their limit. Then, with $w = e^z$,

$$\varphi_\lambda(w) = e^{\Phi_\lambda(z)}$$

is a well-defined analytic function on U, and $\varphi_\lambda \circ f_\lambda(z) = [\varphi_\lambda(z)]^n$. The definition of φ_λ depends on the particular lifting, F_λ, of f_λ we choose, but otherwise the Böttcher Theorem is independent of this choice.

Since f_λ is analytic in λ, by carefully choosing the branch of the natural logarithm we use in defining $F_\lambda(z) = \log(f_\lambda(e^z))$, F_λ can be made to depend analytically on λ. Therefore the $\{F_\lambda^{\circ k}/n^k\}_{k=1}^{\infty}$ vary analytically with λ.

Notice that $\Phi_\lambda(z)$ is defined on the domain $V = \exp^{-1} U$. The set $\varphi_\lambda(U) = e^{\Phi_\lambda(V)}$ is contained in $\hat{\mathbb{C}} - \bar{D}$, by Böttcher's Theorem again. Thus $\Phi_\lambda(V)$ is contained in the half-plane $\{z \in \mathbb{C} \mid \operatorname{Re}(z) > 0\}$ and clearly omits three values. Since the $\{F_\lambda^{\circ k}/n^k\}_{k=1}^{\infty}$ converge uniformly to Φ_λ, there is some k beyond which all $F_\lambda^{\circ k}/n^k$ maps V into $\{z \in \mathbb{C} \mid \operatorname{Re}(z) > -1\}$ and these functions also omit three values. So we may assume without loss of generality that the the sequence of functions $\{F_\lambda^{\circ k}/n^k\}_{k=1}^{\infty}$ satisfy the conditions of Lemma [10.2]. The conclusion then is that $\Phi_\lambda(z)$, for fixed $z \in V$, is an analytic function of λ. Since Φ_λ is the lift of φ_λ, we immediately obtain the desired conclusion, $\varphi_\lambda(z)$, for $z \in U$, is an analytic function of λ. Q.E.D.

Note that if the $\varphi_\lambda(z)$ are defined, for some subfamily, for z in a domain larger than that of U, then the dependence on λ still holds, as the functions at this z will be determined by the analytically varying germ in U. The most restrictive condition is finding the universal neighborhood U,which is in the domain of all the φ_λ. See Lemma [7.9] and Lemma [7.10], for finding this neighborhood for the specific polynomial family, $\mathcal{G}$.