Polygonal Refractive Outer Billiards

Jaewoo Park
University Laboratory High School

March 8, 2026

Abstract

Extending recent work on *refractive billiards*, we introduce and study the corresponding *refractive outer billiards* system about a convex polygon. Gutkin and Simányi showed in 1992 that for regular outer billiards, orbits about a certain class of polygons called *quasi-rational* polygons are bounded, and that orbits about *rational* polygons are periodic. Tabachnikov and Culter later proved in 2007 that every outer billiard system about a convex polygon admits a periodic trajectory. We generalize both results to the refractive setting.

1 Introduction

Let Γ be a closed convex curve in the plane. The *outer billiards map* T around Γ is defined as follows. Let E be the region outside of Γ, and let $o \in E$ be a point in the plane. There are two *supporting rays* $L_+(o)$ and $L_-(o)$ from o to Γ, such that the entirety of Γ lies to the right of L_+ and to the left of L_-. If L_+ has a unique intersection with Γ, then T is a reflection about the point of tangency. If the intersection is *not* unique, then T is undefined. See Figure 1.

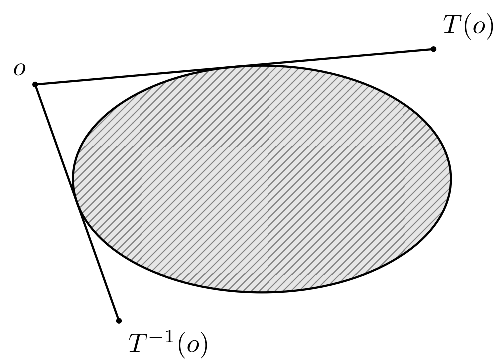

Figure 1: The outer billiards map T.

Moser first popularized outer billiards as a toy model for the solar system in [5]. Problems about the stability and structure of orbits are of great interest: for instance, Moser showed that if Γ is C^6-curved and has positive curvature, then all orbits are bounded [6]. Dolgopyat and Fayad later proved that outer billiards around a semicircle has unbounded orbits [9].

Our interest lies in the case where Γ is a convex polygon; therefore, we will now use P instead of Γ. The study of orbits about polygons is rich and complex. Gutkin-Simányi [1], Kolodziej [4], and Vivaldi-Shaidenko [7] independently proved that all orbits about a class of polygons called *quasi-rational polygons* are bounded, and that all orbits about *rational polygons* are periodic. All rational polygons are quasi-rational; for instance, all regular n-gons are quasi-rational, but only the regular 3, 4 and 6-gons are rational. Tabachnikov and Culter [2] also showed that every polygon has periodic orbits and that the points that induce periodic orbits have a positive measure in the plane. For other notable results, see [8, 10, 12].

Recently, in [3], we generalized regular (Birkhoff) billiards to the *refractive billiards* system. In this paper, we build on our paper, introducing the *refractive outer billiards* system. We generalize the main results of [1] and [2].

The refractive outer billiard map T around a convex n-gon P is defined as follows. Throughout this paper, we follow Gutkin's notation. Let $\lambda_1, \lambda_2, \ldots, \lambda_k$ be a sequence of positive *refractive indices* (viewed modulo k) that satisfy

$$\prod_{i=1}^{k} \lambda_i = 1.$$

Given a point o outside P, the map T is well-defined (with respect to all choices of refractive indices) when o does not lie on a line through a side of P. If this is the case, let A_1 denote the intersection of $L_+(o)$ with P. Then $T(o)$ is the reflection of o about A_1, followed by a scaling of λ_1 about A_1. Generally, $T^m(o)$ is the reflection of $T^{m-1}(o)$ about A_{m-1}, followed by a scaling of λ_m about A_{m-1}. See Figure 2.

Note that the inverse map T^{-1} is also defined in a similar way. Let A_0 denote the intersection of $L_-(o)$ with P. Then $T^{-1}(o)$ is the reflection of o about A_0, followed by a scaling of $1/\lambda_1$ about A_1, and the general case is defined similarly.

Let P be a convex polygon, T the refractive outer billiard map, and $\lambda_1, \ldots, \lambda_k$ refractive indices. Then, our main results are as follows.

Theorem 1. *If P is quasi-rational, then every orbit of T is bounded.*

Theorem 2. *If P is rational and each of $\lambda_1, \ldots, \lambda_k$ is rational, then every orbit of T is periodic.*

Theorem 3. *Given any P and $\lambda_1, \ldots, \lambda_k$, there exists a periodic orbit.*

In Section 2, we review Gutkin's construction for outer billiards. We define the necklace map, the cone and ray construction, the necklace polygon, and define what it means for a polygon to be rational and quasi-rational.

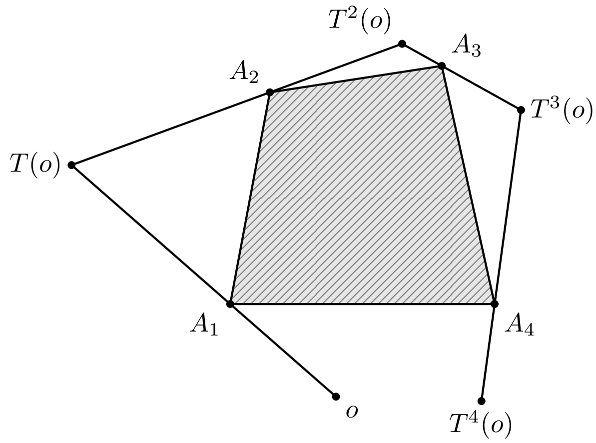

Figure 2: Refractive outer billiards with $\lambda_1 = 3/2,\ \lambda_2 = 2/3,\ \lambda_3 = 2,\ \lambda_4 = 1/2$.

In Section 3, we prove the generalizations of [1] for refractive outer billiards. Using the *refractive necklace map*, we establish that the *first return map* for the system is bounded if P is quasi-rational. We show that if P is rational with each refraction coefficient rational, then the possible configurations of P are finite, which implies periodicity.

In Section 4, we prove the generalization of [2] for refractive outer billiards. We use the fact that the refractive necklace map roughly follows Gutkin's necklace polygon.

In Appendix A, we present the motivation behind refractive outer billiards, highlighting how the projective duality between billiards and outer billiards extends to the duality between refractive billiards and refractive outer billiards.

2 Basic Definitions

In this section, we consider *regular* outer billiards. We summarize and adapt the constructions of Gutkin [1] for completeness. Define σ_ℓ for $\ell \in \mathbb{Z}^+$ to be subsets of the plane such that if $o \in \sigma_\ell$, then either $T^\ell(o)$ or $T^{-\ell}(o)$ is undefined. For instance, the union of all lines through the sides of P is σ_1. Inductively, we see that each σ_ℓ is a finite union of straight lines, so that the set $\sigma = \cup_{n=1}^{\infty} \sigma_n$ is a countable union of straight lines. It follows that σ is a set of zero measure in the plane.

We call the set of points that lie outside of P and are not in σ the *strongly regular points* about P. For a strongly regular point, T^ℓ is well-defined for all integer ℓ. From now on, all points that we consider will be strongly regular. For all arguments, the exclusion of a set of measure zero does not impact the logic.

We now introduce the necklace map, based on a key idea: we reflect the polygons instead of the points. Take $P = P_0$ and a strongly regular point o.

Call the left tangency point A_1 the *head*, and the right tangency point A_0 the *tail*. We then reflect P_0 about A_1 to get P_1. More generally, a polygon P_ℓ has head $A_{\ell+1}$ and tail A_ℓ. Reflecting P_ℓ about its head gives $P_{\ell+1}$, and reflecting about its tail gives $P_{\ell-1}$. Given P_0, we can continue this process infinitely in both directions (since o is strongly regular) to obtain a sequence $\{\dots, P_{-1}, P_0, P_1, \dots\}$ called the *necklace of P about o*. The *necklace map* sends P_ℓ to $P_{\ell+1}$.

We can temporarily forget the outer billiards map and investigate the necklace instead, due to the following theorem. See Figure 3.

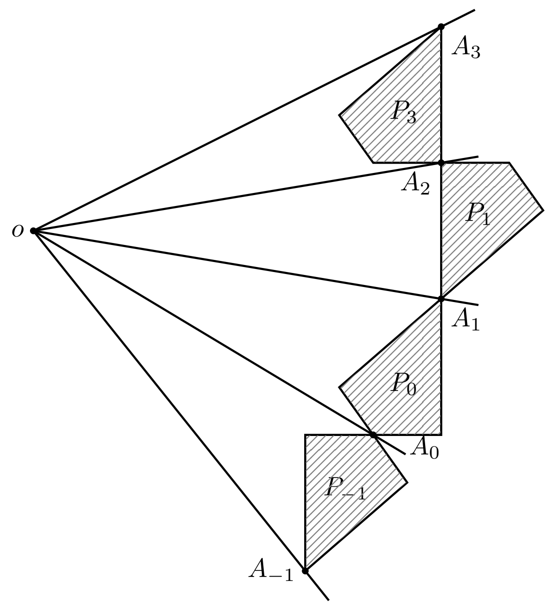

Figure 3: Part of the necklace of $P = P_0$ about o.

Theorem 4 (Gutkin [1, pp. 434–435]). *The necklace of a convex polygon P about o is defined simultaneously with the orbit $\{T^\ell(o) \mid \ell \in \mathbb{Z}\}$. Specifically, we have that for $\ell \in \mathbb{Z}^+$,*

- $T^\ell(o) = (r_1 \circ \cdots \circ r_\ell)(o)$,
- $T^{-\ell}(o) = (r_{-1} \circ \cdots \circ r_{-\ell})(o)$,

in the sense that the relative position of P to $T^\ell(o)$ is the same as the relative position of P_ℓ to o. Here, r_i denotes Euclidean reflection about A_i.

Moreover, the orbit $\{T^\ell(o) \mid \ell \in \mathbb{Z}\}$ is bounded if and only if the necklace is bounded, and periodic if and only if the necklace is periodic.

Now we introduce Gutkin's *cone and ray* construction.

Definition. Fix a convex n-gon P in the plane. Since P is an n-gon, we can draw n straight lines $\ell_1, \dots, \ell_n$ through o that are parallel to each side of P

(note that if P has parallel sides, these lines may overlap). The lines partition the plane into $2n$ cones, possibly degenerate, which we denote $C_1, \ldots, C_{2n}$, in counter-clockwise order.

The lines $\ell_1, \ldots, \ell_n$ create $2n$ rays in the plane, labeled as follows. Let cone C_i be bounded by rays R_i and R_{i+1}, and so forth, with the rays being labeled counter-clockwise. Note that we have the identities $C_{2n+1} = C_1, R_{2n+1} = R_1$ and $C_{n=i} = -C_i, R_{n+i} = -R_i$. See Figure 4.

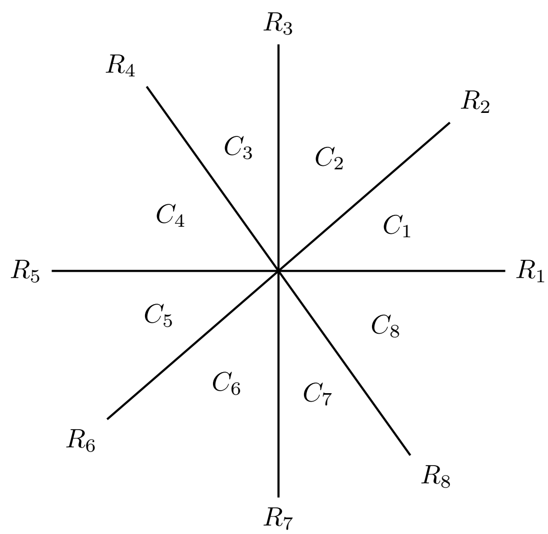

Figure 4: The cone and ray construction for P_0 in Figure 3.

We say that a polygon on the plane is *inside* the cone C_j if it intersects C_j and does not intersect the next cone C_{j+1}. Let G be the group of reflections and translations of the plane. Then, define $\mathscr{P}$ to be the set of polygons *congruent* to P_0 (i.e., polygons Q such that $Q = g(P_0)$ for some $g \in G$) that are also strongly regular about o. Denoting our original polygon by P_0, the following lemma holds.

Lemma 5 (Gutkin [1, pp. 436–437])**.** *Choose a cone C_i. Then, consider the subset $\mathscr{P}_i \subset \mathscr{P}$ of polygons inside the cone C_i. For each $Q \in \mathscr{P}_i$, the head and the tail about o are well-defined. Furthermore, the vector connecting the head and the tail does not depend on the choice of $Q \in \mathscr{P}_i$.*

For a given cone C_j, we call the vector starting from the tail and ending at the head of any $Q \in \mathscr{P}_i$ the *necklace vector in cone C_j*, and denote it as $\vec{a}_j$. Note that $\vec{a}_{j+n} = -\vec{a}_j$.

Now, choose a point $A_1 \in R_1$. Draw the ray emanating from A_1 in the direction of $\vec{a}_1$, until it intersects R_2 at A_2. Then draw the ray in direction $\vec{a}_2$ from A_2, and repeat the process until we return to the ray $R_{2n+1} = R_1$ at A_{2n+1}.

Lemma 6 (Gutkin [1, pp. 438–440])**.** *The polygonal line generated by this process is closed; i.e.,* $A_{2n+1} = A_1$.

This process traces a polygon called the *necklace polygon*, denoted Q.

Remark. Any changes in the definition of Q, i.e., the position of o, the choice of R_1, and the position of A_1 on R_1, changes Q by translations and dilations only. Moreover, the necklace polygon is a convex, centrally symmetric $2n$-gon.

Using this definition of the necklace polygon, we can now define *quasi-rationality.*

Definition. Take $Q = A_1 \dots A_{2n+1}$ to be a necklace polygon of P. Then there exist positive real numbers $r_1, r_2, \dots, r_{2n}$ such that for $1 \le i \le 2n$, we have

$$\overrightarrow{A_i A_{i+1}} = r_i \vec{a}_i.$$

Note that $r_{n+i} = r_i$.

We say that the polygon P is *quasi-rational* if $r_1, \dots, r_n$ are rational up to a common factor, i.e., $(r_1 : r_2 : \cdots : r_n) \in \mathbb{QP}^{n-1}$. We say that P is *rational* if the vertices of P belong to a lattice.

It should again be noted that every rational polygon is quasi-rational—refer to [1] for a proof. We also collect a definition that is useful down the line.

Definition. A *truncated strip* is an infinite strip bordered by two parallel rays and a polygonal line. See Figure 5.

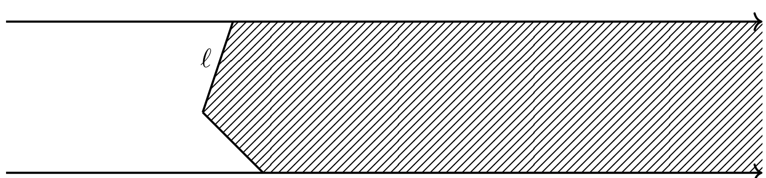

Figure 5: A truncated strip.

3 Orbits in Refractive Billiards

We now extend to refractive billiards, where we now have *refractive coefficients* $\lambda_1, \dots, \lambda_k$ that multiply to 1. Similarly to the regular outer billiard case, the points for which $T^\ell(o)$ is not well-defined for any $\ell \in \mathbb{Z}^+$ is a countable union of finitely many lines, and thus is a set of zero measure. We will only work with *strongly regular* points.

The analog of the necklace construction is as follows. Fix a point o at the origin and choose a polygon $P = P_0$ on the plane. Then, P_1 is the reflection of P_0 about its head A_1, followed by a scaling by $1/\lambda_1$ about A_1. More generally,

P_ℓ is the reflection of $P_{\ell-1}$ about A_ℓ, followed by a scaling by $1/\lambda_\ell$ about A_ℓ. It is also the reflection of $P_{\ell+1}$ about its tail P_ℓ, followed by a scaling by $\lambda_{\ell+1}$.

Thus, given a polygon P_0, we obtain the sequence $\{\dots, P_{-1}, P_0, P_1, \dots\}$, which we will call the *refractive necklace of P about o*. Moreover, we define the *refractive necklace map* as the transformation that sends P_ℓ to $P_{\ell+1}$. With this definition in mind, there is a natural correspondence between refractive billiards and the refractive necklace. See Figure 6.

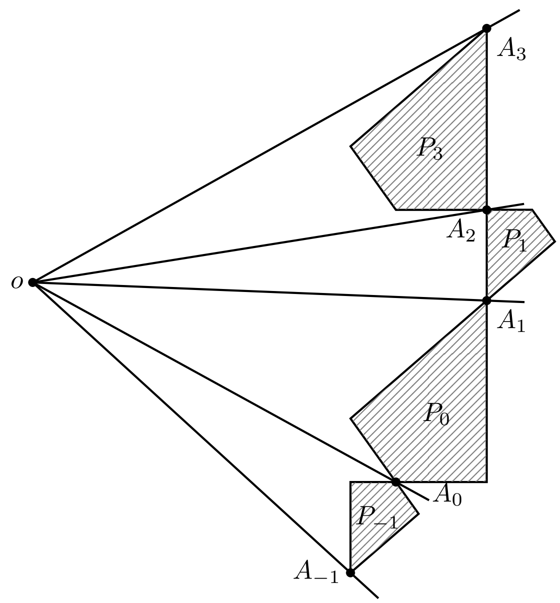

Figure 6: Part of the refractive necklace with $\lambda_1 = 2, \lambda_2 = 1/2$.

Theorem 7. *The refractive necklace of P about o is defined simultaneously with the orbit $\{T^\ell(o) \mid \ell \in \mathbb{Z}\}$. Specifically, we have that for $\ell > 0$,*

- $T^\ell(o) = (r_1 \circ \cdots \circ r_\ell)(o)$,
- $T^{-\ell}(o) = (r_{-1} \circ \cdots \circ r_{-\ell})(o)$,

in the following sense: the relative position of $P = P_0$ to $T^\ell(o)$ is the same as the relative position of P_ℓ to o. Here, r_i denotes the Euclidean reflection transformation about the point A_i, followed by a scaling by λ_i about A_i.

Moreover, the infinite orbit $\{T^\ell(o) \mid \ell \in \mathbb{Z}\}$ is bounded if and only if the necklace is bounded, and periodic if and only if the necklace is periodic.

Proof. Let R_i be the similarity that reflects about A_i and then scales by $1/\lambda_i$ about A_i, so that

$$P_i = R_i(P_{i-1}).$$

Note $r_i := R_i^{-1}$. Let T_i be the refractive outer billiard map with respect to P_i. By construction,

$$T_{i-1} = r_i\, T_i\, R_i = R_i^{-1} T_i R_i,$$

so in particular

$$T = T_0 = r_1 T_1 R_1.$$

Let $o_1 := T(o) = r_1(o)$. We prove by induction on $n \geq 1$ that

$$T^n(o) = (r_1 \circ \cdots \circ r_n)(o).$$

For $n = 1$ this is exactly the definition of T at o, so the base case holds.

Assume the formula holds for necklaces of length $n-1$. Using the conjugacy and the fact that $R_1(o_1) = o$, we have

$$T^{n-1}(o_1) = (r_1 T_1^{n-1} R_1)(o_1) = r_1 T_1^{n-1}(o).$$

Thus

$$T^n(o) = T^{n-1}(T(o)) = T^{n-1}(o_1) = r_1 T_1^{n-1}(o).$$

Now the shorter necklace $\{P_1, \ldots, P_n\}$ with map T_1 is of the same type, so by the induction hypothesis applied to T_1,

$$T_1^{n-1}(o) = (r_2 \circ \cdots \circ r_n)(o),$$

and hence

$$T^n(o) = (r_1 \circ \cdots \circ r_n)(o).$$

The proof for negative iterations is similar.

For the other part of the theorem, each r_i is a similarity with scale factor λ_i, so S_k is a similarity with factor $\Lambda_k := \lambda_1 \cdots \lambda_k$. Since the product of the λ_i over one circuit of the polygon is 1, the factors Λ_k stay uniformly bounded and bounded away from 0. Thus

$$d\big(T^k(o), P_0\big) = d\big(S_k(o), P_0\big) = \lambda_1 \cdots \lambda_k\, d\big(o, S_k^{-1}(P_0)\big) = \lambda_1 \cdots \lambda_k\, d\big(o, P_k\big)$$

shows that the orbit $\{T^k(o)\}$ is bounded if and only if the necklace $\{P_k\}$ is bounded.

If $T^N(o) = o$ for some $N > 0$, then S_N is a similarity fixing o with $|\Lambda_N| = 1$, hence either the identity or a half–turn about o; in both cases the necklace is periodic. Conversely, a periodic necklace forces S_N to be either the identity or a half–turn, so $T^N(o) = S_N(o) = o$ and the orbit is periodic. □

Using this correspondence, we can study the refractive necklace instead of the refractive billiards map. Along the necklace, we have exactly $2k$ different possible configurations of polygons up to translation: a polygon is either oriented the same as P or a reflection of P, and there are k possible sizes, namely P scaled by a factor of $\{1/\Lambda_i \mid 1 \leq i \leq k\}$.

Let $\mathcal{S}$ be the set of possible configurations of polygons on the plane through iterations of the necklace map. Define $\mathcal{S}_i^+$ as the set of polygons that are congruent to P_i and have the same orientation as P, and $\mathcal{S}_i^-$ as the polygons congruent to P_i with opposite orientation. So, the polygons in $\mathcal{S}_i = \mathcal{S}_i^- \sqcup \mathcal{S}_i^+$ are congruent to P scaled by a factor of $1/\Lambda_i$. Immediately, we have

$$\mathcal{S} = \mathcal{S}_0^- \sqcup \mathcal{S}_0^+ \sqcup \mathcal{S}_1^- \sqcup \cdots \sqcup \mathcal{S}_{n-1}^- \sqcup \mathcal{S}_{n-1}^+.$$

Returning to the cone and ray construction from the previous section, pick a ray R_j. We will use $\pm$ to simultaneously define constructions that hold for both $+$ and $-$.

Definition. Let $\mathcal{S}_{R_j}$ denote the set of polygons $Q \in \mathcal{S}$ that intersect R_j but not R_{j+1}. Geometrically, this represents the polygons in $\mathcal{S}$ that have ray R_j as the "furthest" ray it touches.

Define $\mathcal{S}_{i,j}^+$ to be the intersection of $\mathcal{S}_i^+$ and $\mathcal{S}_{R_j}$, and define $\mathcal{S}_{i,j}^-$ similarly. This is the set of polygons in $\mathcal{S}$ that are congruent to P_i, in the same orientation as P (if $+$) or the opposite orientation (if $-$), and intersects ray R_j but not R_{j+1}. Note that $\mathcal{S}_i^+ \sqcup \mathcal{S}_i^-$ contains all polygons congruent to P_i, but $\mathcal{S}_{i,j}^+ \sqcup \mathcal{S}_{i,j}^-$ only contains the polygons congruent to P_i which intersect R_j but not R_{j+1}. Thus, generally,

$$\mathcal{S}_i^\pm \neq \bigsqcup_{j=1}^{2n} \mathcal{S}_{i,j}^\pm.$$

Within each set $\mathcal{S}_{i,j}^\pm$, each polygon is the same orientation and size, so they differ by a translation only. In other words, polygons $Q, Q' \in \mathcal{S}_{i,j}^+$ satisfy $Q = Q' + \vec{u}$ for some vector $\vec{u}$. Then the heads of Q and Q' in R_j, A_+ and A'_+ respectively, satisfy $A_+ = A'_+ + \vec{u}$.

Thus, there exists a bijective correspondence between the set $\mathcal{S}_{i,j}^\pm$ and its *set of heads* $H_{i,j}^\pm$. The *head function* $h_{i,j}^\pm : \mathcal{S}_{i,j}^\pm \to H_{i,j}^\pm$ that sends a polygon to its head satisfies the relation $h(Q + \vec{u}) = h(Q) + \vec{u}$ for all $Q \in \mathcal{S}_{i,j}^\pm$ and $\vec{u}$ such that $Q + \vec{u} \in \mathcal{S}_{i,j}^\pm$.

There exist unique vectors $\vec{d}_j \in R_j$ and $\vec{b}_j \in R_{j+1}$ such that $\vec{b}_j, \vec{d}_j$, and $\vec{a}_j$ for a triangle in the plane.

Lemma 8. *The set $H_{i,j}^\pm$ is a truncated strip for each $1 \leq i \leq n$ and $1 \leq j \leq 2m$. In particular, $H_{i,j}^\pm$ is bounded by R_j, the ray $\tilde{R}_j = R_j + \Lambda_i \vec{b}_j$, and some polygonal line. Moreover, $H_{i,j}^\pm$ is $H_{1,j}^\pm$ scaled by Λ_i. See Figure 7.*

Proof. The first part is a direct generalization of the proof in pp. 441–442 of [1], so it is omitted. For the second part, note that polygons in $\mathcal{S}_{i,j}^\pm$ are just polygons in $\mathcal{S}_{1,j}^\pm$ scaled by Λ_i. By virtue of the construction of the truncated strip in [1], it follows that $H_{i,j}^\pm$ is simply $H_{1,j}^\pm$ scaled by Λ_i. □

Note, of course, that the strips $\mathcal{S}_{i,j}^+$ and $\mathcal{S}_{i,j}^-$ need not be the same, since their orientations are different.

For each cone C_j, we have $2k$ truncated strips. We take the *disjoint union of truncated strips*

$$H_j := \bigsqcup_{i=1}^{k} (H_{i,j}^+ \sqcup H_{i,j}^-).$$

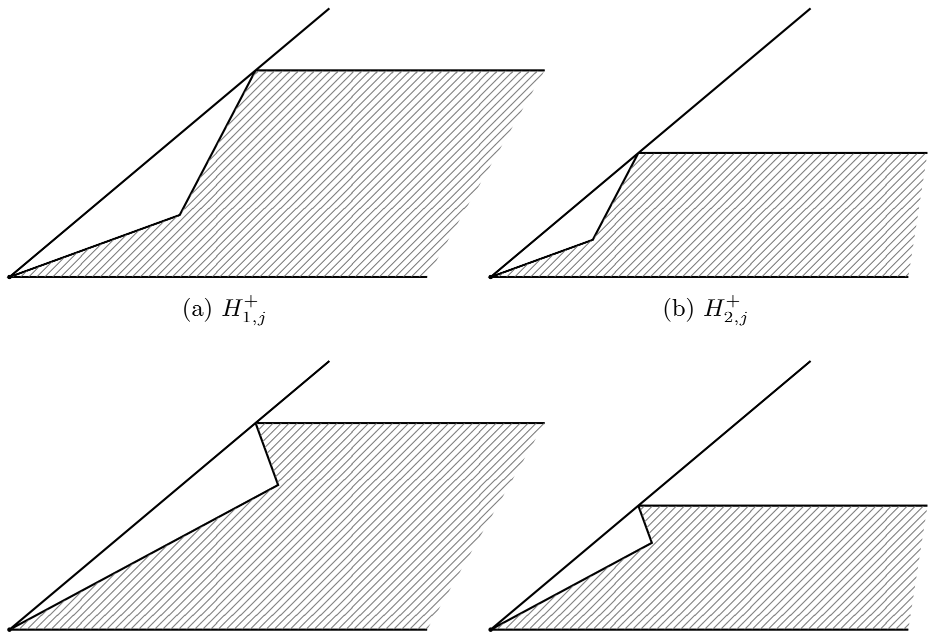

(a) $H^+_{1,j}$ (b) $H^+_{2,j}$

(c) $H^-_{1,j}$ (d) $H^-_{2,j}$

Figure 7: Four example regions with $\lambda_2 = 3/5$.

By the bijective correspondence between $\mathcal{S}^{\pm}_{i,j}$ and $H^{\pm}_{i,j}$, each point in H_j corresponds to a unique polygon.

Using this new language, we can define the *next cone map.*

Definition. Take a polygon $Q \in \mathcal{S}^{\pm}_{i,j}$. Continue applying the refractive necklace map to Q until it intersects ray R_{j+1}. Denote this new polygon as Q'. Note that this process terminates since the vector $\vec{a}_j$ is not parallel to either of the rays R_j or R_{j+1}.

The *next cone map in cone* C_j, denoted f_j, maps heads to heads: $f_j : H_j \to H_{j+1}$. It maps the head of Q in H_j to the head *in cone* C_{j+1} of Q', which lies in H_{j+1}. Importantly, we need to add a *head correction vector* $\vec{h}^{\pm}_i$ that accounts for a change in the head vertex, since the head vertex depends on the orientation of Q' and the ray R_{j+1}.

Remark. Taking H to be $H_1 \sqcup H_2 \sqcup \cdots \sqcup H_{2k}$, the maps f_j, when combined, naturally extend to the *general* next cone map $f : H \to H$. Note that f satisfies $f(H_j) \subseteq H_{j+1}$.

The map f_j can also be viewed as a map of tuples $(A, i, \varepsilon) \mapsto (\tilde{A}, \tilde{i}, \tilde{\varepsilon})$. Here, A represents the location of the polygon's head in C_j; the values i, ε represent the truncated strip $H^{\varepsilon}_{i,j}$ to which the polygon belongs. Similarly, $\tilde{A}$ is the location of the head in cone C_{j+1}, and the polygon lies in $H^{\tilde{\varepsilon}}_{\tilde{i},j+1}$. By our prior discussion, we know that $\tilde{A}$ is $A + C \cdot \vec{a}_j + \vec{h}$, where C is some constant and $\vec{h}$ is some head correction vector.

To find the precise values of C and $\vec{h}$, we partition each truncated strip $H^{\pm}_{i,j}$ into regions $\pi^{\pm}_{i,j,m}$ for $m \in \mathbb{Z}^+$. The region $\pi^{\pm}_{i,j,m}$ is the set of heads in $H^{\pm}_{i,j}$ such that f_j is equal to m iterations of the refractive necklace map. In the following proposition, we describe these regions.

Proposition 9. *On ray R_j, take the infinite sequence of points that have distance $\Lambda_i \vec{d_j}$, $\Lambda_{i+1}\vec{d_j}$, $\Lambda_{i+2}\vec{d_j}$, $\dots$ away from each other, starting at the origin. Since n is finite, this sequence has period n. Now, draw a line parallel to R_{j+1} through each point. Taking intersections of these lines with the ray $R_j + \Lambda_i \vec{b}_j$, we obtain a periodic sequence of parallelograms.*

The region $\pi^{\pm}_{i,j,k}$ is the intersection of the k-th parallelogram with the truncated strip $H^{\pm}_{i,j}$. Note that π_1 may not be a parallelogram since $\mathcal{S}^{\varepsilon}_i$ is a truncated strip.

Proof. This follows from the fact that the first iteration of a head in $H^{\pm}_{i,j}$ adds $\Lambda_i \vec{a_j}$ to P, the second iteration adds $\Lambda_{i+1}\vec{a_j}$, and so on. Partitioning the truncated strip via lines parallel to $\vec{a_j}$ of these lengths, we get the construction above. □

Indeed, using this partition, we can completely describe the next cone map f_j.

Lemma 10. *Let $f_j : H_j \to H_{j+1}$ be the next cone map on cone R_j. Using the tuple notation above, f_j maps (A, i, ε) to $(\tilde{A}, \tilde{i}, \tilde{\varepsilon})$. If A lies in $\pi^{\varepsilon}_{i,j,m}$, we have the following:*

$$\begin{cases} \tilde{A} = A + \big(\Lambda_{i+1} + \Lambda_{i+2} + \cdots + \Lambda_{i+m}\big) \cdot \vec{a} + \vec{h}^{\tilde{\varepsilon}}_i, \\ \tilde{\varepsilon} = (-1)^m \varepsilon, \\ \tilde{i} \equiv i + m \bmod k. \end{cases} \tag{1}$$

Proof. Since our polygon has scale Λ_i, the first iteration of the refractive necklace map gives a polygon of scale Λ_{i+1}. Repeating m times, we get a total of m polygons with scales $\Lambda_{i+1}, \dots, \Lambda_{i+m}$. Multiply this constant by $\vec{a}$, then add the head correction vector $\vec{h}^{\tilde{\varepsilon}}_i$ to get the first part. The other two parts follow immediately from the definition of π_k. □

Corollary. *Recall the definition of $\vec{b}, \vec{d}$ from Proposition 9. Writing $\Lambda = 1 + \Lambda_1 + \cdots + \Lambda_{k-1}$, the following relation holds:*

$$f_j\Big(A + 2\Lambda \vec{d_j}\Big) = f(A) + 2\Lambda \vec{b_j}.$$

Proof. Suppose A lies in $\pi^{\varepsilon}_{i,j,m}$. By construction of π, it follows that $A + 2\Lambda\vec{d_j}$ lies in π_{m+2k}, so by Lemma 10 the head positions agree. Note that the head correction vector is also identical. Moreover, $(-1)^{m+2k}\varepsilon = (-1)^m\varepsilon$, and $i+m \equiv i + m + 2k \pmod{k}$. □

Now, we introduce the *first return map* F.

Definition. Consider a polygon with head in H_1. The *first return map* $F = f_{2n} \circ \cdots \circ f_1$ represents the head of the first polygon that intersects R_1 again after completing a "full loop." In the same sense as the next cone map f, the function $F : H_1 \to H_{2n+1} = H_1$ maps a tuple (A, i, ε) to $(\tilde{A}, \tilde{i}, \tilde{\varepsilon})$.

With the first return map defined, we can finally start analyzing the structure of orbits.

Proposition 11. *Let the polygon P be quasi-rational. Then there exists $n \in \mathbb{Z}^+$ such that*

$$F(x + 2n\Lambda\vec{d_1}) = F(x) + 2n\Lambda\vec{b_1}.$$

Proof. We have $oA_{j+1} = r_j\vec{b_j} = r_{j+1}\vec{d}_{j+1}$. If the polygon P is quasi-rational, then we can assume that

$$r_j = \Lambda n_j, n_j \in \mathbb{N}\ (1 \le j \le 2n).$$

Now, applying the identity

$$f_j\left(A + 2\vec{d_j}\Lambda\right) = f(A) + 2\vec{b_j}\Lambda,$$

we get

$$f_j(x + 2n_j\Lambda\vec{d_i}) = f_j(x) + 2n_j\Lambda\vec{b_j} = f_j(x) + 2n_{j+1}\Lambda\vec{d}_{i+1}.$$

Iterating for $i = 1, \ldots, 2m$, we get

$$(f_{2m} \circ \cdots \circ f_2 \circ f_1)(x + 2n_1\Lambda\vec{d_1}) = (f_{2m} \circ \cdots \circ f_2 \circ f_1)(x) + 2n_1\Lambda\vec{d_1},$$

so our proof is complete: set $n = 2n_1$. □

Corollary. *Let $\vec{p} := 2n_1\Lambda\vec{d_1}$, and choose $k \in \mathbb{Z}$, $x \in H_1$. If the point $x + k\vec{p}$ lies in H_1, then*

$$F(x + k\vec{p}) = F(x) + k\vec{p}.$$

Using the corollary above, we can reduce the entire first return map to a *fundamental domain* Π, which is the disjoint union of all $\pi^{\pm}_{i,1,k}$ with bottom parts inside some choice of $\vec{p}$ in R_1.

Thus, we can redefine F as follows:

Definition. Let $\Phi : \Pi \to \Pi$ and $\tau : \Pi \to \mathbb{Z}$ be functions such that $F(x) = \Phi(x) + \tau(x)\vec{p}$.

We have the following three results of Gutkin that we can apply.

Lemma 12 (Gutkin [1, p. 445]). *The pair (Φ, τ) uniquely determines the first return map F. The mapping Φ is invertible and $\Phi, \Phi^{-1} : \Pi \to \Pi$ are local translations.*

Corollary (Gutkin [1, p. 446]). *The first return map F is uniquely determined by the pair of maps $(\Phi : \Pi \to \Pi, \tau : \Pi \to \mathbb{Z}_{\geq -1})$. The function τ corresponding to an invertible mapping F can take values $-1, 0, 1$ only.*

Theorem 13 (Gutkin [1, p. 446]). *If F is invertible and satisfies the condition in Theorem 11, then the orbits $\left\{F^k(x) : -\infty < k < \infty\right\}$ are bounded. If F is invertible and the translation vectors $\vec{t}(\varepsilon, i, k)$ defining F generate a discrete group, then the orbits of F are periodic.*

Thus we obtain the boundedness theorem for the refractive outer billiards system:

Theorem 14. *Let T be the refractive dual billiard mapping about a polygon P. If P is quasi-rational then the orbits of T are bounded.*

Proof. We know that the orbit is bounded if and only if the necklace is bounded. The infinite necklace is bounded if and only if the first return map is bounded. □

Rationality is slightly different:

Theorem 15. *If P is rational and each of $\lambda_1, \ldots, \lambda_n$ are rational, then the orbits of the refractive dual billiards map T are periodic.*

Proof. First, note that since P is rational, it is also quasi-rational. By Theorem 14, we know that the orbit is bounded. If P is rational and each of $\lambda_1, \ldots, \lambda_n$ are rational, then the actual set of possible heads in H_1 is a finite set. Keep iterating the first return map until we get a periodic orbit. □

4 The Existence of Periodic Orbits

Now, we generalize Culter's theorem: every refractive outer billiard has a periodic orbit. It is worth noting that the proof for this in [2] avoids the usage of the necklace map technique. However, we present an approach using necklaces here.

The main idea of the proof is as follows. Instead of directly searching for a periodic orbit of T, we will look for a periodic orbit of T^{2n}. After $2n$ iterations of the refractive necklace map, the polygon returns to its original shape: in other words, $T^{2n}(P)$ is a translation of P.

We will look for a necklace that satisfies the following conditions:

1. Each ray contains exactly one polygon of the refractive necklace, which are all translations of each other. Denote the polygon on ray R_j as Q_j.

2. The following identity holds:

$$Q_{j+1} = Q_j + 2\Lambda \overrightarrow{a_j} \cdot p_j,$$

   where p_j is some positive integer satisfying $p_{n+i} = p_i$.

3. The *head* of Q_j lies in the parallelogram $\pi^+_{1,j,2p_jn}$.

The last two conditions ensure that $Q_{j+1} = T^{2n \cdot p_j}(Q_j)$. If these conditions are met, we can simply fill in the remaining refractive necklace to obtain a periodic orbit.

First, we introduce a lemma that helps us simplify the head condition.

Lemma 16. *Consider a convex n-gon P and its $2n$ cones. Suppose P lies in a cone bordered by rays R and R'. Construct the parallelogram π with sides parallel to R and R' and $\overrightarrow{A_-A_+}$ as a diagonal. Then $P \subseteq \pi$.*

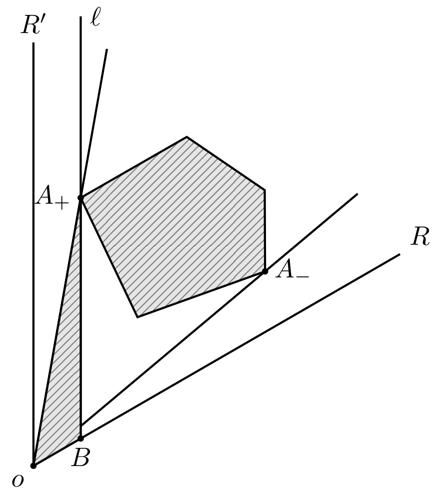

Figure 8: A diagram for the proof of Lemma 16.

Proof. See Figure 8. Draw ℓ, the line that passes through A_+ and is parallel to R'. We claim that every vertex of P lies on the opposite side of ℓ as R'.

Let B be the intersection of ℓ and R. If vertex v of P was in the truncated strip formed by $R', \overline{OA_+}$, and ℓ, then this contradicts the fact that $\overline{OA_+}$ is a support line of P. If instead $v \in \triangle OA_+B$, then by the convexity of P, it follows that the vertex adjacent to A_+, denoted v', lies in the truncated strip formed by R', ℓ, and R. We only need to consider the case where v' is in $\triangle OA_+B$. In this case, the ray from o in the direction of $v'A_+$ lies strictly between R and R', so that C is not a cone: the ray from o in the direction of $v'A_+$ lies between R' and oA_+. Similarly, it follows that every vertex of P lies inside π. By convexity, $P \subseteq \pi$. □

Now, pick an arbitrary point x in the interior of P. Using the lemma above, we have the following fact.

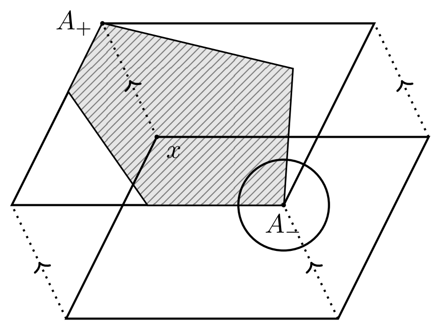

Figure 9: An ε-ball for x centered at A_-. Each dotted vector is $\vec{v}$.

Proposition 17. *For each ray R_j, there exists $\varepsilon_j > 0$ such that if P is translated so that x is in one of the ε_j-balls centered at $2\Lambda r_j p_j$, where $p_j \in \mathbb{Z}^+$, then the head of the translated polygon lies inside the region $\pi^+_{1,j,2p_j}$.*

Proof. First note that the "bottom right" corner of $\pi^+_{1,j,2p_j}$ is placed at $2\Lambda\vec{r}_j p_j$ along R_j from o. Therefore, we can prove this for general translations of π. Let $\vec{v}$ be the vector difference of x and A_+. Since $\vec{v}$ is completely contained in the parallelogram, the points q for which $q + \vec{v}$ is inside the parallelogram is a translation of π, which contains the bottom right corner. It follows that such an ε_j exists. See Figure 9. □

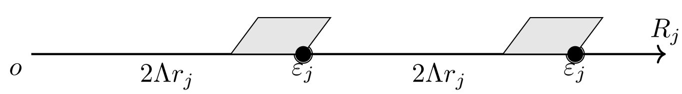

Figure 10: The ε_j-balls along R_j. The first parallelogram is $\pi^+_{1,j,2n}$, and the second is $\pi^+_{1,j,4n}$.

Taking $\min(\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_{2n}) = \varepsilon$, we know that if on the ray R_j the point x lies in a ε-ball centered at distance $2\Lambda r_j p_j$ for any $p_j \in \mathbb{Z}^+$, the refractive necklace map will send the polygon to the next cone in $2p_j$ iterations.

At this point, it suffices to find a polygon with vertices in these ε-balls such that the vector difference between points on adjacent rays R_j and R_{j+1} is an integer multiple of $2\Lambda\vec{a}_j$. Then, we can place x at each vertex of this new polygon to obtain a periodic orbit.

Lemma 18 (Tabachnikov [2, pp. 4–5])**.** *For any $\delta > 0$, there exists arbitrarily large $q \in \mathbb{R}$ and positive integers $p_1, p_2, \ldots, p_{2n}$ such that for each j,*

$$|qr_j - p_j| < \delta.$$

Remark. In other words, we can approximate the ratio $(r_1 : r_2 : \cdots : r_n) \in \mathbb{RP}^{n-1}$ arbitrarily well with $(p_1 : \cdots : p_n) \in \mathbb{ZP}^{n-1}$. Note that since $r_{n+j} = r_j$, we also have $p_{n+j} = p_j$.

Theorem 19. *For any choice of P and $\lambda_1, \dots, \lambda_k$, there exists a periodic orbit. Moreover, the set of points that form a periodic orbit has positive measure in the plane.*

Proof. We seek a more specific polygon with side length vectors

$$2\Lambda\vec{a}_1 p_1, \dots, 2\Lambda\vec{a}_n p_n, -2\Lambda\vec{a}_1 p_1, \dots, -2\Lambda\vec{a}_n p_n.$$

Note that the sum of these is the zero vector, so that we actually return to the original point.

We will start from the point $2\Lambda\vec{r}_1 p_1$ on R_1 and add the vectors above, claiming that each subsequent point stays close to the point $2\Lambda\vec{b}_j p_j$ on R_j. More precisely, we claim that there exist $p_1, p_2, \dots, p_{2n} \in \mathbb{Z}^+$ such that for each j,

$$||2\Lambda\vec{b}_j p_j - (2\Lambda\vec{b}_1 p_1 + \sum_{i=1}^{j-1} 2\Lambda\vec{a}_i p_i)|| < \varepsilon \cdot \frac{j}{2n}.$$

We will prove this using an inductive argument:

$$\begin{aligned}
&||2\Lambda\vec{b}_j p_j - (2\Lambda\vec{b}_1 p_1 + \sum_{i=1}^{j-1} 2\Lambda\vec{a}_i p_i)|| \\
&= ||2\Lambda\vec{b}_{j-1} p_{j-1} - (2\Lambda\vec{b}_1 p_1 + \sum_{i=1}^{j-2} 2\Lambda\vec{a}_i p_i) - 2\Lambda\vec{a}_{j-1} p_{j-1} - 2\Lambda\vec{b}_{j-1} p_{j-1} + 2\Lambda\vec{b}_j p_j|| \\
&< \varepsilon \cdot \frac{j-1}{2n} + 2\Lambda||\vec{b}_j p_j - \vec{a}_{j-1} p_{j-1} - \vec{b}_{j-1} p_{j-1}|| \\
&= \varepsilon \cdot \frac{j-1}{2n} + 2\Lambda||\vec{b}_j p_j - \vec{b}_j p_{j-1} \cdot \frac{r_j}{r_{j-1}}|| \\
&= \varepsilon \cdot \frac{j-1}{2n} + 2\Lambda(p_j - \frac{r_j}{r_{j-1}} p_{j-1})||\vec{b}_j|| \\
&< \varepsilon \cdot \frac{j}{2n},
\end{aligned}$$

where the last step follows since we can make $p_j - \frac{r_j}{r_{j-1}} p_{j-1}$ arbitrarily small. □

5 Acknowledgements

I express my sincere gratitude to my mentor, Prof. Maxim Arnold, for his invaluable patience, dedication, and feedback throughout this project. I am also grateful to him for providing the topic of this paper. I also thank the MIT PRIMES-USA program for making this research possible.

References

[1] E. Gutkin, N. Simanyi, Dual polygonal billiards and necklace dynamics, *Comm. Math. Phys.* **143** (1992), 431–449.

[2] S. Tabachnikov, A proof of Culter's theorem on the existence of periodic orbits in polygonal outer billiards, *Geom. Dedicata* **129** (2007), 83–87.

[3] M. Arnold, J. Park, Snell meets Fagnano. Path optimization through an imperfect mirror, available online at the URL: `https://arxiv.org/abs/2512.02236`.

[4] R. Kolodziej, The antibilliard outside a polygon, *Bull. Polish Acad. Sci. Math.* **37** (1989), 163–168.

[5] J. Moser, Is the solar system stable?, *Math. Intelligencer* **1** (1978), 65–71.

[6] J. Moser, *Stable and random motions in dynamical systems*, Annals of Mathematics Studies, vol. 77, Princeton Univ. Press, 1973.

[7] F. Vivaldi, A. V. Shaidenko, Global stability of a class of discontinuous dual billiards, *Comm. Math. Phys.* **110** (1987), 625–640.

[8] D. Genin, *Regular and chaotic dynamics of outer billiards*, Ph.D. thesis, Pennsylvania State University, 2005.

[9] D. Dolgopyat, B. Fayad, Unbounded orbits for semicircular outer billiard, *Ann. Henri Poincaré* **10** (2009), 357–375.

[10] R. E. Schwartz, Unbounded orbits for outer billiards. I, *J. Mod. Dyn.* **1** (2007), 371–424.

[11] S. Tabachnikov, *Geometry and billiards*, Amer. Math. Soc., 2005.

[12] P. Boyland, Dual Billiards, twist maps, and impact oscillators, *Nonlinearity* **9** (1996), 1411–1438.

A Motivating Refractive Outer Billiards

The outer billiards system gets its other name, *dual billiards*, from the *projective duality* that it shares with classical billiards. To illustrate this duality, define the outer billiards system on S^2, the 2-sphere.

There exists a natural correspondence between the set of oriented great circles and the set of points on the sphere given by a "projective" duality: every oriented great circle corresponds to its *north pole*. Given an oriented curve $\gamma \subset S^2$, we define its *dual curve* γ^* as follows.

The curve γ defines a one-parameter family of oriented tangent lines. Extend each such tangent line to obtain a corresponding family of oriented great circles. Then, the dual curve γ^* is the collection of north poles of these great circles. Note that for great circles a, b and corresponding north poles A, B, the spherical distance AB equals the angle between the lines a and b. See Figure 11.

This operation is indeed a duality: one can show that γ^* is obtained from γ by moving each point a distance of $\pi/2$ in the direction orthogonal to γ, and

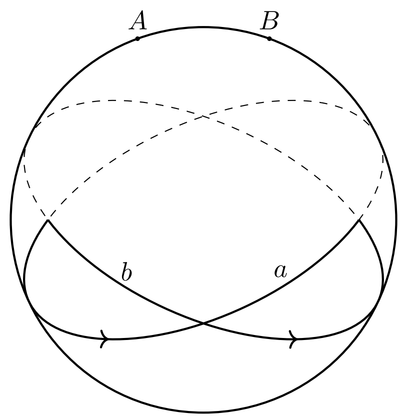

Figure 11: The projective duality.

that $(\gamma^*)^*$ is the antipodal curve of γ on S^2. Note also the following property: if a north pole B is on the line a, the north pole A lies on the line b.

Consider a billiard reflection inside γ, where a ray a reflects off the tangent line ℓ at a point P and travels in a new direction b. The three lines $a, b,$ and ℓ all pass through P, so that all three points $A, B,$ and L lie on the line p. Since L is also on γ^*, we get $L = p \cap \gamma^*$. Moreover, since the angle of incidence is equal to the angle of reflection, we have $AL = LB$ from the length-angle duality. Therefore, A, B, and L all lie on the same line with $AL = LB$.

The Birkhoff billiard map taking a to b after a reflection at $P \in \gamma$ corresponds exactly to the outer billiards map taking A to B after reflection about $P^* \in \gamma^*$. This duality holds only in S^2: in the plane, there is no direct relation between the systems. For more details, see [11].

Extending this duality to *refractive inner billiards* defined in [3] gives refractive *outer* billiards. The name stems from the refraction phenomena in optics: for two materials with refractive indices n and $\tilde{n}$, Snell's law states that

$$\frac{n}{\tilde{n}} = \frac{\sin\tilde{\theta}}{\sin\theta},$$

where θ is the angle of incidence and $\tilde{\theta}$ is the angle of refraction.

Consider a billiard ball moving inside a table Γ. When the ball collides with the boundary $\partial\Gamma$, the refractive billiards system refracts the ball instead of reflecting it. We are given refractive indices $\lambda_1, \ldots, \lambda_k$, with $\lambda_1 \ldots \lambda_k = 1$. For the ith refraction, $\sin\tilde{\theta}_i / \sin\theta_i = \lambda_i$, where θ_i is the i^{th} angle of incidence and $\tilde{\theta}_i$ the i^{th} angle of refraction. Then, we reflect the ray exiting the table back inside.

It can be shown that the same projective duality gives rise to refractive outer billiards.