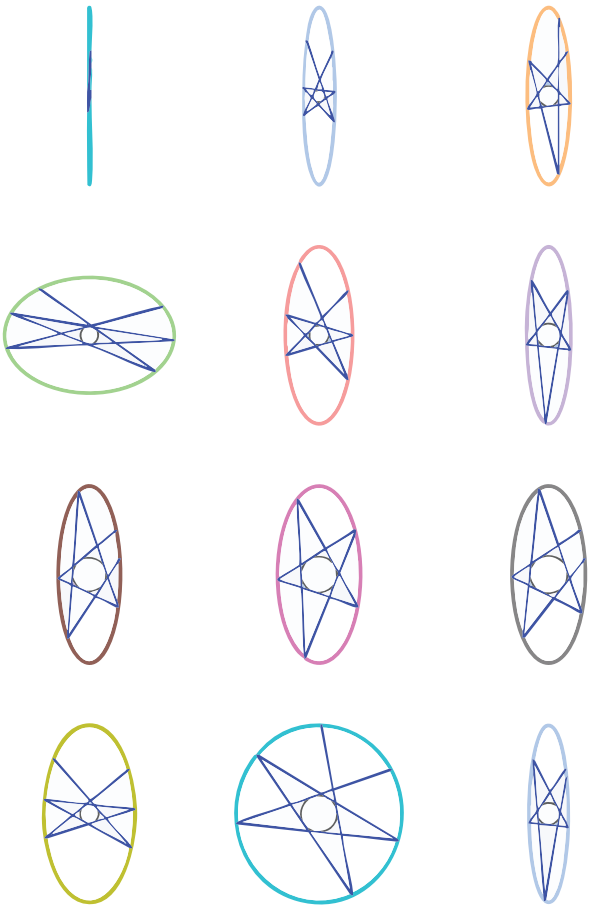
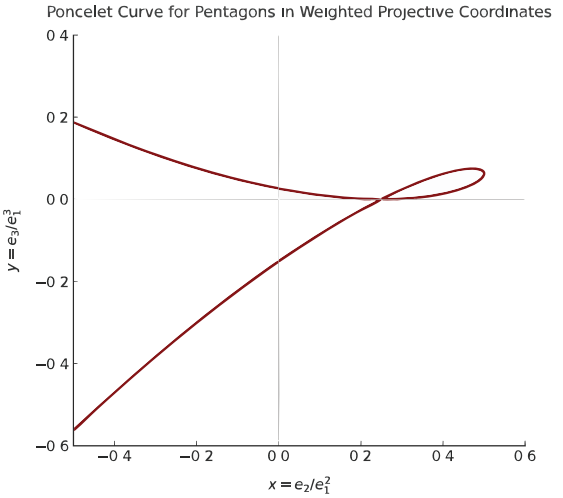
A Complex (Projective) Billiard Game

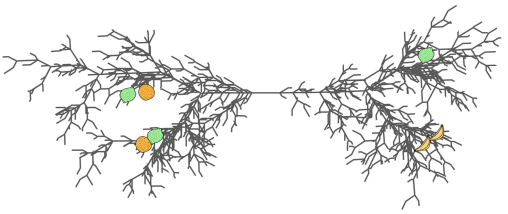
Top (Poncelet Trajectories): Each subplot in the top section displays a pair of conics: the outer circle $x^2 + y^2 = 1$ and an inner ellipse of the form $\left(\frac{x}{a}\right)^2 + \left(\frac{y}{b}\right)^2 = 1$, where (a, b) are derived from sampled Cayley invariants (x, y). A five-step Poncelet trajectory is drawn in blue, starting from a fixed angle and iteratively constructing tangents from the circle to the ellipse. If the polygon closes after five steps, the pair lies on the Poncelet curve.
Bottom (Poncelet Curve): The curve shown is the Poncelet curve associated with pentagonal (5-periodic) Poncelet polygons inscribed in one conic and circumscribed around another. Each point on this curve corresponds to a set of conic pairs for which a closed 5-gon exists that satisfies Poncelet's closure condition. The coordinates (x, y) represent algebraic invariants of the conic pair, specifically $x = e_2/e_1^2$ and $y = e_3/e_1^3$, where e_k are the elementary symmetric functions of the characteristic multipliers derived from the conic configuration. A point on the curve means that a 5-periodic polygon can be inscribed and circumscribed for that particular combination of invariants. For more, see the excellent blog of Oliver Nash at http://oliver-nash.org/2018/07/08/poring-over-poncelet/index.html.





A Complex (Projective) Billiard Game

Poncelet's Porism describes an unexpected property of billiard trajectories between two nested ellipses: if one path returns to its starting point after a finite number of bounces, then all starting points generate periodic trajectories with the same number of bounces. This geometric result connects to elliptic curves and measure-preserving dynamical systems, exemplifying how problems in distinct fields reduce to the same equations through appropriate frameworks.



Elliptical Billiards ∘ Poncelet's Porism ∘ Nested Ellipses ∘ Poncelet Map ∘ Invariant Weighted Measure ∘ Rotation Number ∘ Rational vs Irrational ∘ Benford's Law ∘ Leading Digit Distribution ∘ Gelfand Framework ∘ Dynamical Proof

« La mathématique est l'art de donner le même nom à des choses différentes. »

("Mathematics is the art of giving the same name to different things.")

— Henri Poincaré, 1908

"Algebra is the offer made by the devil to the mathematician.
The devil says: I will give you this powerful machine,
it will answer any question you like.
All you need to do is give me your soul:
give up geometry and you will have this marvelous machine."

— Michael Atiyah, 2001

A Complex (Projective) Billiard Game

> In the early nineteenth century, Jean-Victor Poncelet (1788–1867) pioneered projective geometry by examining how shapes transform under projection. Around 1822, he introduced the concept now known as Poncelet's Porism, demonstrating that a closed polygon can be inscribed in one conic and circumscribed about another conic, provided it exists once for a given number of sides. This insight spurred an intense study of conic sections, with mathematicians such as Carl Gustav Jacob Jacobi and Arthur Cayley extending Poncelet's results to explore more algebraic properties of these curves.
>
> Over the mid to late nineteenth century, researchers recognized a link between geometric theorems such as Poncelet's Porism and physical billiard trajectories. Elliptical billiards, in particular, drew interest when it was observed that the classical reflection law led to periodic paths that echoed Poncelet's closure conditions. By the turn of the twentieth century, these geometric investigations began intertwining with the nascent field of algebraic geometry, revealing that repeated reflections could be described by equations resembling elliptic or hyperelliptic curves.

The rules for the motion of billiards are straightforward. A ball moves in a straight line until it meets a boundary — where it reflects according to the law of geometric optics: the angle of incidence equals the angle of reflection. On a rectangular table, this produces familiar trajectories. Some repeat periodically, some trace diagonals, and others fill up regions with a regular structure. The behavior is fully determined by the shape of the table and the reflection rule.

When the boundary is not polygonal but curved, the situation changes. An ellipse introduces a new constraint. By definition, the sum of distances from any point on the ellipse to two fixed points, the foci, remains constant. Combined with the reflection law, this implies an optical property: any ray emanating from one focus reflects off the boundary and passes through the other. This follows from the fact that, at the point of reflection, the tangent line to the ellipse bisects the angle formed between the incoming and outgoing paths. That is, the segment from the first focus to the boundary, and then from the boundary to the second focus, meets the boundary at equal angles on either side. The configuration is symmetric, and the total path length is stationary with respect to small variations.

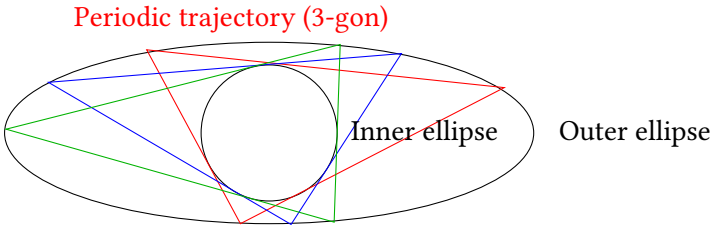
Most trajectories, however, do not begin at a focus. A generic path reflects from arbitrary points on the boundary. Some paths return to their origin after a finite number of reflections; others do not. Some fill annular regions densely, never repeating. The classification of such trajectories depends on both the initial direction and the shape of the boundary.

Now consider two nested, smooth, closed curves, an outer ellipse and an inner ellipse lying entirely within it. Imagine a polygon whose vertices lie on the outer ellipse and whose sides are tangent to the inner ellipse. When the two ellipses are confocal, such a polygon is exactly a closed billiard path — the reflection law at the outer boundary is equivalent to tangency with the confocal caustic. But the construction makes sense for any nested pair: each segment connects two points on the outer ellipse while remaining tangent to the inner one. If such a polygon exists and closes after n steps, it defines a Poncelet polygon.

©2026 David H. Silver, CC BY-NC 4.0

The question is whether such a configuration is rare. If one closed polygon exists, does that imply anything about other starting points? Does the system admit only a single orbit, or does the existence of one periodic path imply a rule?

Poncelet's Porism answers affirmatively. If there exists a single n-gon inscribed in the outer ellipse and tangent to the inner one, then for every point on the outer ellipse there exists such an n-gon. The ellipse is foliated by periodic trajectories of the same type. The existence of one closed polygon implies the existence of an infinite family, each differing only by a rotation of the starting point.



Poncelet's Porism: If one closed trajectory exists, then infinitely many exist.

The geometric construction of Poncelet polygons can be reinterpreted as a dynamical process. Fix two nested ellipses. Choose a point on the outer ellipse and draw a line tangent to the inner ellipse, continuing the segment until it intersects the outer ellipse again. Repeat this procedure: from each new intersection, draw the unique line tangent to the inner ellipse and find its next point of contact with the outer ellipse. The result is a discrete sequence of points on the outer ellipse, each determined from its predecessor. This iteration defines a map from the outer ellipse to itself.

This map — commonly referred to as the Poncelet map — sends a point on the outer ellipse to the next vertex of the corresponding Poncelet polygon.

(Side note: A *porism* is a byproduct of a theorem, usually a corollary which implies infinitely many solutions. Which reminds me of the joke by mathematician Jerry Bona, in which he described the equivalence of the Axiom of Choice (Zermelo, 1904), Zorn's Lemma, and the Well-Ordering Theorem as follows: *"The Axiom of Choice is obviously true, the Well-Ordering Theorem is obviously false, and Zorn's Lemma — who can tell?"* These statements are mathematically equivalent, yet their perceived plausibility varies widely.)

Poncelet originally proved his porism in 1822 using projective geometry and the theory of conics. Later approaches employed algebraic geometry, treating the problem as a question about curves in complex projective space. The configuration of two nested conics defines an elliptic curve, and the closure condition corresponds to torsion points on this curve. These methods, while powerful, require sophisticated machinery from algebraic geometry and complex analysis.

The dynamical systems approach sidesteps this complexity entirely. The question becomes: what kind of transformation is the Poncelet map?

It turns out that this map preserves a specific measure on the ellipse: a notion of "size" for subsets that remains unchanged under iteration. This measure is not uniform arc length but rather arc length weighted by the distance from each point to its tangency point on the

inner ellipse. To see why, first apply an affine transformation to make the inner ellipse into a circle (this preserves tangency and incidence). Now consider nearby points p and p' on the outer ellipse, and their images $T(p)$ and $T(p')$. The tangent chords from these points form two similar triangles; taking the limit as $p' \to p$ and using $\sin\theta \sim \theta$ for small arcs gives $ds_1/ds = |p_1m|/|pm|$, where m is the tangency point on the chord from p. Since the inner ellipse is now a circle, the two tangent segments from p_1 to it have equal length: $|p_1m| = \rho(T(p))$. Thus $ds/\rho(p) = ds_1/\rho(T(p))$, and the weighted measure $d\mu = ds/\rho$ is invariant.

This measure is absolutely continuous with respect to arc length (it has a smooth, positive density) and is non-atomic (individual points have zero measure). The existence of such an invariant measure is a special feature of conic geometry and is not generic for arbitrary smooth curves.

A foundational result in ergodic theory — related to the classification of circle homeomorphisms — states that any orientation-preserving homeomorphism of a circle that preserves a finite, non-atomic, absolutely continuous measure must be topologically conjugate to a rigid rotation (Poincaré, 1885). "Conjugate" here means there exists a homeomorphism φ of the circle such that $\varphi \circ T \circ \varphi^{-1}$ is a pure rotation $R_\alpha : \theta \mapsto \theta + \alpha \mod 2\pi$. The map and the rotation have identical dynamics, just expressed in different coordinates.

The conjugacy parameter α, called the rotation number, is an invariant of the map. It measures the average rate of angular advance per iteration and determines the orbit structure completely. If α is a rational multiple of 2π, say $\alpha = 2\pi p/q$ in lowest terms, then every orbit is periodic with period q. The circle decomposes into congruent periodic orbits, each forming a q-gon in the original ellipse geometry. If α is an irrational multiple of 2π, then no orbit closes. Instead, every orbit is dense: it passes arbitrarily close to every point on the ellipse, filling the curve uniformly in the limit. This is equidistribution, a consequence of Weyl's equidistribution theorem for irrational rotations (Weyl, 1916).

Poncelet's Porism now follows as a corollary. The existence or non-existence of a closed n-gon depends only on whether the rotation number is rational or irrational — a property intrinsic to the pair of ellipses. If one closed polygon exists, the rotation number is rational, and therefore all starting points produce closed polygons with the same number of sides. If no closed polygon exists, the rotation number is irrational, and no starting point produces one.

This same approach appears in problems involving the distribution of leading digits in exponential sequences. The following analysis concerns the frequency of initial digits in powers of integers, governed by irrational rotations on the unit interval.

Consider the powers of 2: $2^1 = 2$, $2^5 = 32$, $2^{10} = 1024$, $2^{53} = 9\,007\,199\,254\,740\,992$. Some begin with 1, others with 3, 9, etc. The digits appear irregularly, but over time, a pattern emerges. The frequency of each digit stabilizes and matches a specific distribution.

The explanation is logarithmic. Write $2^n = 10^{n \log_{10} 2}$. The number of digits in 2^n grows roughly linearly in n, and its leading digit depends only on the decimal part (denoted by $\{x\}$) of $n \log_{10} 2$. As n increases, the sequence $\{n \log_{10} 2\}$ fills the interval $[0, 1)$ evenly, because

$\log_{10} 2$ is irrational. This means that 2^n is equally likely to appear in any logarithmic subinterval of a given order of magnitude.

A number begins with digit d if its logarithm lies between $\log_{10} d$ and $\log_{10}(d+1)$. So the proportion of terms 2^n that begin with digit d approaches

$$\log_{10}\left(1+\frac{1}{d}\right).$$

This is Benford's Law (Benford, 1938). It predicts that 1 appears as the leading digit about 30% of the time, while 9 appears less than 5% of the time. The same reasoning applies to any base-b sequence where $\log_{10} b$ is irrational. The decimal parts of $n \log_{10} b$ become evenly spread, and the digit frequencies converge to the same logarithmic formula.

Benford's Law extends far beyond powers of integers. Real-world datasets that span multiple orders of magnitude — river lengths, city populations, physical constants, file sizes, mountain heights — often follow the same logarithmic distribution. The key requirement is that the data range widely without artificial constraints or preferred scales.

When values are distributed across many orders of magnitude, they tend to be spread evenly on a logarithmic scale rather than a linear one. This occurs naturally when data arise from multiplicative processes or when no particular scale is privileged (so it should be invariant to multiplication by a constant). In such cases, the probability of a value falling into the interval that starts with digit d is proportional to the logarithmic length of that interval: $\log_{10}(1+1/d)$. This geometric fact, combined with scale invariance, produces Benford's distribution without requiring randomness or special assumptions. This property makes Benford's Law useful for fraud detection, where artificial datasets often deviate from the expected pattern.

Benford's Law does not apply universally. It fails when numbers are tightly clustered, rounded, or constrained by conventions, such as ID numbers, product prices, or phone records. But when applicable, it can serve as a diagnostic tool. Large deviations from the expected digit frequencies may indicate fabricated or manipulated data.

Forensic accountants have used Benford's Law to uncover anomalies in financial statements, including during the Enron scandal. Tax authorities apply it to detect suspicious filings. In the 2009 Iranian presidential election, some analysts claimed that statistical tests based on Benford's Law identified irregularities in reported vote counts. In each case, observed digit distributions diverged significantly from the logarithmic baseline, suggesting artificial data generation.

This formulation, noted by Gelfand, is identical to the classification of the Poncelet map. In both cases, a rotation acts on a compact one-dimensional space, and the long-term behavior of orbits (whether periodic or equidistributed) depends only on the rationality of a single parameter. For powers of b, this parameter is $\log_{10} b$; for the Poncelet construction, it is the angular step induced by tangency.

Good Teachers and Invariant Measures

I fell in love with mathematics during my undergraduate studies, particularly through a set theory course taught by Professor Amos Nevo. The material was foundational, focusing on logic, sets, and proofs. Nevo consistently highlighted how abstract structures recur across fields, enhancing its significance. Even when teaching introductory content, he pointed toward broader connections: between algebraic symmetries, analysis, and geometry.

Later, in a seminar on dynamical systems (we were a group of students who tried to enroll in any course Amos offered), he asked me to present a proof of Poncelet's Porism. I began working through the classical approach: projective geometry, dual conics, and elliptic curves. After several weeks of effort, I came to him with the outline. He laughed and said that wasn't the proof he had in mind.

Instead, he pointed me to a short argument we had already studied in class, one based on the existence of an invariant measure. From that, topological conjugacy implies that the Poncelet map behaves like a rigid rotation. Rational rotation number implies periodicity. The porism drops out almost immediately.

The Poncelet map is a beautiful object and exemplifies how abstract tools from dynamical systems, such as invariant measures and topological conjugacy, can be seen doing the heavy lifting on a classical geometry problem. It was also one of the best seminars I ever took for accelerating mathematical maturity.



Mike's understanding of Snell's law made him unbeatable,
but considerably slowed down the game

Poncelet for Two Ellipses

Let C and D be two smooth, strictly convex, nested ellipses in the plane, with $D \subset \text{int}(C)$. The *Poncelet map* $T : C \to C$ is defined as follows: for a point $p \in C$, draw the line through p tangent to the inner ellipse D (choosing one of the two tangent directions consistently and smoothly along C); let $T(p)$ be the second point of intersection of this line with C. This consistent choice defines an orientation-preserving homeomorphism of C.

We show that T preserves a natural measure and is topologically conjugate to a circle rotation. This yields a complete classification of the dynamics of T, and with it, a proof of Poncelet's closure theorem.

Affine Reduction and the Invariant Measure To analyze T, we apply two affine reductions. First, by dilating along one axis of D, we may assume the inner ellipse D is a circle; this preserves incidence and tangency, so the Poncelet map is unchanged. Now suppose C is given by

$$x^2/a^2 + y^2/b^2 = 1, \quad \text{with } a > b > 0.$$

Let $U(x, y) = (x/a, y/b)$, sending C to the unit circle $\widetilde{C}$ and D to an ellipse $\widetilde{D}$. Let $\tilde{s}$ denote arc-length on $\widetilde{C}$. For $p \in C$, let m be the tangency point of the chord $pT(p)$ with D, and let m_1 be the tangency point of the next chord $T(p)T^2(p)$. Set $\rho(p) := |p - m|$. The invariant measure is $d\mu := d\tilde{s} \,/\, \rho$, mixing arc-length on $\widetilde{C}$ with the tangent-segment distance in the original (circular-D) coordinates.

Invariance of the Measure. For nearby $\tilde{p}, \tilde{p}' \in \widetilde{C}$, the chords $\tilde{p}\tilde{p}_1$ and $\tilde{p}'\tilde{p}'_1$ meet at a point $\tilde{n}$, forming similar triangles. As $\tilde{p}' \to \tilde{p}$, $\tilde{n} \to \widetilde{m}$ and chord ratios on $\widetilde{C}$ converge to arc-length ratios ($\sin\theta/\theta \to 1$), so

$$\frac{d\tilde{s}_1}{d\tilde{s}} = \frac{|\tilde{p}_1\widetilde{m}|}{|\tilde{p}\widetilde{m}|} = \frac{|p_1 m|}{|pm|},$$

where the second equality holds because U is linear and preserves distance ratios. Since D is a circle, the two tangent segments from p_1 to D are equal: $|p_1 m| = |p_1 m_1| = \rho(T(p))$.

Hence

$$\frac{d\tilde{s}_1}{d\tilde{s}} = \frac{\rho(T(p))}{\rho(p)}, \quad \text{i.e.,} \quad \frac{d\tilde{s}}{\rho(p)} = \frac{d\tilde{s}_1}{\rho(T(p))},$$

and $d\mu = d\tilde{s} \,/\, \rho$ is T-invariant.

Topological Conjugacy to a Circle Rotation An orientation-preserving circle homeomorphism with a finite non-atomic invariant measure positive on every non-empty open arc has no wandering intervals (a result due to Poincaré); therefore its Poincaré semiconjugacy to a rotation is a conjugacy.

Our measure μ satisfies: **Finite:** Since C and D are smooth and strictly convex with $D \subset \text{int}(C)$, $\tilde{\rho}$ is continuous on the compact set C' and strictly positive, so it has a positive infimum. Thus $\int_{C'} d\tilde{s}/\tilde{\rho} < \infty$. **Non-atomic:** points have zero measure. **Positive on arcs:** $\tilde{\rho}$ is bounded above on any arc, so $1/\tilde{\rho}$ is bounded below by a positive constant, ensuring open arcs have positive measure.

Thus, there exists a homeomorphism $\varphi : C \to S^1$ and $\alpha \in [0, 1)$ such that: $\varphi \circ T \circ \varphi^{-1} = R_\alpha$, where $R_\alpha(\theta) = \theta + \alpha \mod 1$

The number α, called the *rotation number* of T, quantifies the average angular displacement per iteration. Lift the circle map to $F : \mathbb{R} \to \mathbb{R}$ in an angular coordinate; then $\alpha := \lim_{n\to\infty} (F^n(\theta) - \theta)/n$ exists and is independent of θ. The rotation number of T is $\alpha \mod 1$.

The rotation number classifies the dynamics:

- If $\alpha \in \mathbb{Q}$, write $\alpha = m/n$ in lowest terms. Then all orbits are periodic with period n. Every point on C traces a closed n-gon tangent to D.
- If $\alpha \notin \mathbb{Q}$, then no orbit is periodic. The sequence $p, T(p), T^2(p), \ldots$ becomes dense in C, and no Poncelet polygon closes.

References:

Leopold Flatto, *Poncelet's Theorem.* Mathematical Surveys and Monographs, Vol. 56. American Mathematical Society, 2009 (beautiful book!)

For visualization, see bit.ly/poporism.