

Contents

3 Chaotic Dynamics	1
3.1 Symbolic Dynamics	1
3.1.1 Orbits of $q_c(x) = x^2 + c$ for $c < -2$	2
3.1.2 Nested Intervals, Itineraries, and Cantor Sets	3
3.1.3 Metric Spaces and the Sequence Space	6
3.1.4 Equivalent Dynamical Systems: Homeomorphisms and Conjugation	9
3.1.5 Equivalence of $q_c(x) = x^2 + c$ and the Shift Map for $c < -2$	12
3.2 Chaotic Dynamical Systems	15
3.2.1 Motivation for Chaos: Properties of the Shift Map on the Sequence Space	15
3.2.2 Devaney’s Definition of Chaos	18
3.2.3 Equivalence and Chaotic Systems, Examples	20
3.3 Sarkovskii’s Theorem and Applications	22
3.3.1 The Period-3 Theorem	22
3.3.2 The Sarkovskii Ordering and Sarkovskii’s Theorem	25

3 Chaotic Dynamics

In this chapter, our goal is to study chaotic dynamical systems, with an ultimate goal of trying to understand the chaotic behaviors that we saw in our plots of orbit diagrams. In order to do this, we will introduce symbolic dynamics, a powerful tool that will help us understand a number of chaotic systems. Some more technical facts from topology and analysis will be necessary (primarily, some facts about metric spaces, the topology of the real numbers, and homeomorphisms), so we will develop this material to the extent necessary to apply it.

We will then provide a definition of a chaotic dynamical system that is amenable to practical use, and prove that a number of simple systems are chaotic. Finally, we finish with a discussion of Sarkovskii’s Theorem, a striking and unexpected result that (among other things) implies that any continuous function that possesses a 3-cycle exhibits chaotic behavior.


3.1 Symbolic Dynamics

- As motivation, let us summarize what we have learned about the quadratic maps $q_c(x)$ so far:
 - When $c > \frac{1}{4}$, every orbit of the quadratic map $q_c(x)$ tends to ∞ , while when $c = \frac{1}{4}$ the unique fixed point is weakly attracting on the left and weakly repelling on the right.
 - When $c < \frac{1}{4}$, then if p_+ is the larger fixed point, all orbits outside the interval $(-p_+, p_+)$ tend to ∞ .
 - When $-2 < c < \frac{1}{4}$, any point lying in the interval $(-p_+, p_+)$ will have its orbit completely confined to this interval. We can glean some insight about the behavior of q_c from the orbit diagram: some values appear to have an attracting cycle, while other values do not.


- When $c < -2$, the critical orbit diverges to $+\infty$, so there are no attracting cycles.
- We would now like to understand the dynamics of the maps q_c when $c < -2$ more precisely.
 - To do this, we will ultimately show that the dynamics of q_c are modeled by the dynamics of a simple map on a sequence space. We will then switch to analyzing the sequence space, which (it turns out) we can understand completely.
 - Understanding and formalizing the connection between the dynamics on these two spaces will involve some technical results from analysis and topology, which we will develop as needed.

3.1.1 Orbits of $q_c(x) = x^2 + c$ for $c < -2$


- It might appear that the dynamics of $q_c(x) = x^2 + c$ are uninteresting when $c < -2$, given that the critical orbit diverges to $+\infty$, but this is not at all the case.
 - Explicitly, suppose $c < -2$ is fixed, and let $I = [-p_+, p_+]$, where $p_+ = \frac{1 + \sqrt{1 - 4c}}{2}$ is the larger fixed point of $q_c(x)$.
 - Note: All graphics are produced for the case $c = -91/36$, but the analysis holds in general.
- Notice that q_c maps I onto an interval strictly containing I (since $q_c(0) = c$ does not lie in I), and that the set of points in I whose image also lies in I forms a pair of intervals I_0 and I_1 . Here is a picture:




- If a point $x \in I$ has $q_c(x) \notin I$, then the orbit of x necessarily diverges to $+\infty$ by our earlier analysis.
 - So, if the orbit of x does *not* diverge to $+\infty$, it is necessarily the case that $q_c^n(x) \in I$ for every integer $n \geq 1$.
 - On the other hand, clearly if $q_c^n(x) \in I$ for every $n \geq 1$ then the orbit does not go to $+\infty$, since it stays confined to I .
 - Thus, we see that the orbit of x will stay bounded if and only if $q_c^n(x) \in I$ for all integers $n \geq 1$, which is equivalent to saying $x \in q_c^{-n}(I)$ for every integer $n \geq 1$.
 - Since $I \subseteq q_c^{-1}(I) \subseteq q_c^{-2}(I) \subseteq \dots$ that means the set of points x whose orbit stays bounded is given by the infinite intersection $\Lambda = \bigcap_{n=1}^{\infty} q_c^{-n}(I)$.
 - So now: what does this infinite intersection Λ look like? To understand this, let us examine each of the terms; we already saw that $q_c^{-1}(I)$ is the union of the two intervals I_0 and I_1 .
 - Then $q_c^{-2}(I) = q_c^{-1}(I_0 \cup I_1) = q_c^{-1}(I_0) \cup q_c^{-1}(I_1)$ by basic properties of the inverse image. Each of $q_c^{-1}(I_0)$ and $q_c^{-1}(I_1)$ consists of a pair of closed intervals, one in I_0 and the other in I_1 :



- Notice that $q_c^{-2}(I)$ consists simply of the two intervals I_0 and I_1 , but with a middle portion removed from each: this follows because q_c maps I_0 bijectively onto the interval I , so the set of points in I_0 that q_c^2 sends into I will be all of I_0 except the points in the middle of the interval satisfying $q_c^{-1}(x) < -p_+$. (The same argument holds for I_1 .)
- In a similar way we see that $q_c^{-3}(I)$ will be a collection of eight closed intervals obtained by removing a piece from the middle of each of the four intervals in $q_c^{-2}(I)$:



- We can see (by an easy induction) that $q_c^{-n}(I)$ will consist of 2^n closed intervals, half of which are in I_0 and the other half of which are in I_1 . Here is a typical picture of the successive inverse images $q_c^{-n}(I)$ for $0 \leq n \leq 6$:




- We will show later that whenever $c < -2$, the sizes of the intervals will shrink to zero as we take $n \rightarrow \infty$. Assuming for now that all these intervals have size shrinking to 0, a very reasonable question is: how do we know that the set $\Lambda = \bigcap_{n=1}^{\infty} q_c^{-n}(I)$ is even nonempty? (After all, the intersection cannot contain any intervals at all.)
 - One way is simply to exhibit some points that lie in the intersection: namely, the set $q_c^{-k}(p_+)$ for any $k \geq 1$: after k iterations, every point in this set lands at p_+ , so the orbit certainly always stays in I .
 - Notice that $q_c^{-k}(p_+)$ is a set containing 2^k points, since each point in I has two preimages under q_c , so in fact we have exhibited infinitely many points in the intersection. Indeed, it is not hard to see that $q_c^{-k}(p_+)$ is the set of endpoints of the intervals in $q_c^{-k}(I)$.
- It might seem that, as we iteratively remove the middle portion of each interval, as we take the limit we will only be left with the endpoints of the intervals.
 - In fact, however, there are many more points in Λ , but to show this fact will require a technical result about the topology of \mathbb{R} .

3.1.2 Nested Intervals, Itineraries, and Cantor Sets

- We can describe the set Λ more precisely using a result known as the nested intervals theorem:
- **Theorem** (Nested Intervals): If J_1, J_2, J_3, \dots is a sequence of finite closed intervals in \mathbb{R} and $J_{i+1} \subseteq J_i$ for each i , then the intersection $\bigcap_{i=1}^{\infty} J_i$ is nonempty. Furthermore, if the length of J_i tends to 0 as $i \rightarrow \infty$, then the intersection consists of a single point.
 - **Remark:** The assumption that the intervals are finite is necessary, because an infinite nested intersection of infinite closed intervals can be empty: for example, $\bigcap_{i=1}^{\infty} [i, \infty)$ is empty.
 - **Proof:** Let $J_i = [a_i, b_i]$ where by assumption $a_i \leq b_i$ for each i .

- Since $J_{i+1} \subseteq J_i$ we also have $a_i \leq a_{i+1}$ and $b_{i+1} \leq b_i$ for each i , so the sequence $\{a_i\}$ is increasing and bounded above (by b_1).
 - Thus, by the monotone convergence theorem (a bounded monotone sequence of real numbers has a limit), the sequence a_i has a limit L as $i \rightarrow \infty$. Similarly, $\{b_i\}$ is decreasing and bounded below, so the sequence b_i has a limit M as $i \rightarrow \infty$.
 - Then $[L, M]$ is contained in each interval J_i and hence also in the intersection $\bigcap_{i=1}^{\infty} J_i$. But no $x < L$ can be in the intersection: otherwise, it would be contained in every interval J_i and thus each a_i would be less than x , contradicting the assumption that the $a_i \rightarrow L$. Similarly, no $x > M$ is in the intersection.
 - Thus, $\bigcap_{i=1}^{\infty} J_i = [L, M]$. Furthermore, since $L \leq M$ (since M is an upper bound for the sequence $\{a_i\}$ and L is a lower bound for the sequence $\{b_i\}$), this interval is not empty.
 - For the second part, if the length of $J_i \rightarrow 0$ as $i \rightarrow \infty$, then $b_i - a_i \rightarrow 0$ as $i \rightarrow \infty$, so $L = M$, and thus the intersection is a single point.
- Using the nested intervals theorem, we can give a better description of Λ .

- Given $q_c^{-n}(I)$, we saw that $q_c^{-(n+1)}(I)$ is obtained by removing a piece from the middle of each of the intervals in $q_c^{-n}(I)$.
- Describing an interval in $q_c^{-n}(I)$ is therefore equivalent to recording whether we took the “left interval” or the “right interval” each time we applied f^{-1} .
- We can summarize this information using an n -digit binary string, with 0 meaning “left” and 1 meaning “right”, and thus label each of the intervals accordingly:




- To any infinite sequence of binary digits $\{d_0 d_1 d_2 \dots\}$, we can then construct elements of Λ : if we take J_n to be the interval $I_{d_0 d_1 d_2 \dots d_n}$, then the sequence J_1, J_2, J_3, \dots is a nested sequence of closed intervals of length tending to zero (by our earlier proposition), so by the nested intervals theorem, the intersection is a single point.
 - Conversely, if $x \in \Lambda$, then x necessarily lies in some subinterval of $q_c^{-n}(I)$ for every $n \geq 1$, so by writing down the labels of the sequences, we get an infinite sequence of binary digits.
 - An almost equivalent way of defining this sequence is to compute the iterates of x and determine which of I_0 and I_1 each iterate lands in:
- **Definition:** For $x \in \Lambda$, the itinerary of x is the infinite binary string $S(x) = \{d_0 d_1 d_2 \dots\}$, where $d_i = 0$ if $q_c^n(x) \in I_0$ and $d_n = 1$ if $q_c^n(x) \in I_1$.

- We will return to study this map later, but it serves as our primary motivation for studying sequence spaces.
- Our ultimate goal is to prove that the itinerary map is a homeomorphism (i.e., a continuous bijection with continuous inverse) when viewed in the appropriate context, and can be used to relate the dynamics of a simple map on the space of binary sequences to the dynamics of the quadratic map q_c on \mathbb{R} .

- Before we discuss sequence spaces, we will briefly mention Cantor sets, of which Λ is one example. The most famous Cantor set is the Cantor ternary set:

- **Definition:** The Cantor ternary set is the set $\bigcap_{n=0}^{\infty} C_n$, where $C_0 = [0, 1]$ and C_{n+1} is obtained by deleting the open middle third of each interval in C_n , for each $n \geq 0$.

- Thus, $C_1 = [0, \frac{1}{3}] \cup [\frac{2}{3}, 1]$, $C_2 = [0, \frac{1}{9}] \cup [\frac{2}{9}, \frac{3}{9}] \cup [\frac{6}{9}, \frac{7}{9}] \cup [\frac{8}{9}, 1]$, and so forth. Here is a picture:



- Observe that the picture is almost identical to the picture of the set Λ . Indeed, as we will show later, these two sets are homeomorphic (i.e., there exists a continuous function with continuous inverse mapping Λ to the Cantor ternary set).
- It is possible to establish many of the basic properties of the Cantor ternary set directly via the nested intervals theorem. However, there is a nicer approach arising from a convenient description of the points in the set using base-3 decimal expansions.
- Recall that the base-3 decimal expansion of a (nonnegative) real number has the form $n.d_1d_2d_3d_4\cdots_3 = n + \sum_{i=1}^{\infty} \frac{d_i}{3^i}$, where n is an integer and each d_i is 0, 1, or 2.
 - Notation: We will generally put a subscript of k when working with decimals in base k , when $k \neq 10$, although it should always be clear from the context what base the expansions should be considered in.
 - Any base-3 expansion always converges by comparison to the geometric series $0.2222\dots_3 = 1$.
 - By summing the series when feasible, we can evaluate base-3 expansions. For example, we have $1/3 = 0.1_3$ and $1/4 = 0.010101\dots_3 = 0.\overline{01}_3$, where the overline (bar) means that the indicated portion repeats indefinitely.
 - The real numbers of the form $m/3^n$, for m and n nonnegative integers, have two base-3 expansions: one ending in an infinite string of 0s, and another ending in an infinite string of 2s. (For example, $1/3 = 0.10_3 = 0.0\overline{2}_3$.) All other real numbers have a unique ternary expansion.
- Proposition (Base-3 Expansions and the Cantor Set): A point $\alpha \in [0, 1]$ lies in the Cantor ternary set if and only if it has a base-3 expansion containing only the digits 0 and 2.
 - Proof: Suppose α has a ternary expansion whose first digit is 1. Then α will be removed at the first stage of the construction of the Cantor set: the open middle third of $[0, 1]$ consists of all points whose first base-3 decimal digit is a 1, except for the points $1/3 = 0.1000\dots = 0.0222\dots$ and $2/3 = 0.1222\dots = 0.2000\dots$. But each of the endpoints has a representation containing no 1s, and they are both preserved since we only remove the open middle third.
 - In exactly the same way, if α has a ternary expansion whose n th digit is 1, then α will be removed at the n th stage of the construction of the Cantor set, unless α happens to be one of the endpoints (which both have a representation that does not contain a 1).
 - Conversely, if α has a representation containing no 1s, then it will never lie in the open middle third of any interval during the construction, so α is in the Cantor ternary set.
- As a corollary of the proposition above, we can see that the Cantor ternary set is uncountable: the elements of the Cantor set are the real numbers in $[0, 1]$ having a ternary expansion consisting of only 0s and 2s, and this set is in a bijection with the set of infinite binary sequences (namely, by replacing all of the 2s by 1s), which is uncountable.
 - Recall that a set is countable if it can be put in a one-to-one correspondence with some subset of the positive integers, and a set is uncountable otherwise.
 - A typical example of a countable set is the set of rational numbers.
 - The set of infinite binary sequences is uncountable, as originally proven by Cantor using his famous diagonal argument¹.

¹To summarize: Suppose by way of contradiction that the set of infinite binary sequences were countable. Arrange all of them into an infinite array:

$$\begin{array}{rcl}
 a_1 & = & d_{1,1}d_{1,2}d_{1,3}\cdots \\
 a_2 & = & d_{2,1}d_{2,2}d_{2,3}\cdots \\
 a_3 & = & d_{3,1}d_{3,2}d_{3,3}\cdots \\
 & \vdots & \vdots \\
 & \vdots & \vdots
 \end{array}$$

Now construct the binary sequence x whose i th digit is 1 if $d_{i,i} = 0$ and is 0 if $d_{i,i} = 1$. (These digits run down the diagonal of the array, whence the name of the argument.) Then x cannot be equal to any of the a_i , because it differs in at least one place from every element in the list. This is a contradiction, because we assumed all binary sequences were in the array.

- The set of real numbers, or even the real numbers in the interval $[0, 1]$, is also uncountable, as it can be put into a bijection with the set of infinite binary sequences: in essence², we associate each element in the interval with the infinite sequence of digits in its base-2 decimal expansion.
- From our arguments, we can conclude that there is actually a bijection between the points in the Cantor ternary set and the points in $[0, 1]$.
 - Somehow, the Cantor ternary set is still “large”, even though the sum of the lengths of the intervals at each stage tends to 0 exponentially rapidly!
 - There are many natural ways to alter the construction to create generalized Cantor sets, but we will postpone further analysis of these sets until the next chapter when we study fractals.

3.1.3 Metric Spaces and the Sequence Space

- Definition: The sequence space on two symbols is the set $\Sigma_2 = \{(d_0 d_1 d_2 \dots) : d_i = 0 \text{ or } 1 \text{ for each } i\}$.
 - We think of this space as the set of infinite binary sequences, although we could equally well use any symbols in place of 0 and 1.
 - Although we will not use it, the more general sequence space on n symbols is the set of infinite sequences $(d_0 d_1 d_2 \dots)$ where each d_i lies in the set $\{0, 1, \dots, n-1\}$.
- The sequence space, so far, has no structure: it is simply a set of sequences. In order to do anything with it, we need to specify some additional structure on the space.
 - We will do this by defining a distance function (or metric) that allows us to measure how far apart elements are.
 - First, however, we will outline some of the basic theory of metric spaces.
- Definition: If M is a set, a function $d : M \times M \rightarrow \mathbb{R}$ is called a metric, and the pair (M, d) a metric space, if it obeys the following three properties:
 1. (Nonnegativity) For any $x, y \in M$, $d(x, y) \geq 0$, with $d(x, y) = 0$ if and only if $x = y$.
 2. (Symmetry) For any $x, y \in M$, $d(x, y) = d(y, x)$.
 3. (Triangle Inequality) For any $x, y, z \in M$, $d(x, z) \leq d(x, y) + d(y, z)$.
- The prototypical example of a metric is the Euclidean distance function $d(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|$ on \mathbb{R}^n .
 - Explicitly, if $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_n)$, then $d(\mathbf{x}, \mathbf{y}) = \sqrt{(x_1 - y_1)^2 + \dots + (x_n - y_n)^2}$.
 - When $n = 1$, for example, this is the familiar $d(x, y) = |x - y|$ on \mathbb{R} .
 - The first two axioms are trivial, and the third is geometrically obvious, as it is the (actual) triangle inequality: namely, that the shortest distance between \mathbf{x} and \mathbf{z} is the straight line joining them.
 - The third axiom requires some actual work to prove algebraically, and ultimately, it reduces to the Cauchy-Schwarz inequality $\mathbf{x} \cdot \mathbf{y} \leq \|\mathbf{x}\| \|\mathbf{y}\|$, where $\mathbf{x} \cdot \mathbf{y}$ denotes the dot product. (Recall that $\|\mathbf{x}\|^2 = \mathbf{x} \cdot \mathbf{x}$.)
 - * Explicitly, if we set $\mathbf{a} = \mathbf{x} - \mathbf{y}$ and $\mathbf{b} = \mathbf{y} - \mathbf{z}$, then we want to show that $\|\mathbf{a} + \mathbf{b}\| \leq \|\mathbf{a}\| + \|\mathbf{b}\|$.
 - * We have $(\mathbf{a} + \mathbf{b}) \cdot (\mathbf{a} + \mathbf{b}) = \mathbf{a} \cdot \mathbf{a} + 2(\mathbf{a} \cdot \mathbf{b}) + \mathbf{b} \cdot \mathbf{b} \leq \|\mathbf{a}\|^2 + 2\|\mathbf{a}\| \|\mathbf{b}\| + \|\mathbf{b}\|^2$.
 - * Taking the square root immediately gives the desired result.

²This association is not quite a bijection since some real numbers have two base-2 expansions. However, there are only countably many such numbers, so it is straightforward to fix this issue.

Explicitly, remove all of the elements with two binary expansions from $[0, 1]$ and place them in a sequence a_1, a_2, a_3, \dots , and also remove all of the elements from the set of binary sequences corresponding to these elements of $[0, 1]$, and place them in a sequence b_1, b_2, b_3, \dots . Then define the bijection normally on the elements outside these sequences, and also identify a_i with b_i . The resulting map is a bijection.

- For completeness, we give a one-line proof of the Cauchy-Schwarz inequality in \mathbb{R}^n : observe that

$$\frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (x_i y_j - x_j y_i)^2 = \sum_{i=1}^n x_i^2 \sum_{j=1}^n y_j^2 - \left(\sum_{i=1}^n x_i y_i \right)^2 = \|\mathbf{x}\|^2 \|\mathbf{y}\|^2 - (\mathbf{x} \cdot \mathbf{y})^2$$

and since the left-hand side is a sum of squares, the right-hand side must be nonnegative.

- For posterity we will record the definitions of open and closed sets in a metric space:
- **Definition:** If (M, d) is a metric space, then if $x \in M$ and $r > 0$, the open ball $B_r(x)$ of radius r centered at x is defined to be the set of points within a distance r of x : namely, $B_r(x) = \{y \in M : d(x, y) < r\}$.
 - In the case where $X = \mathbb{R}^n$ with the Euclidean metric, the open ball is a literal ball (i.e., the points lying strictly inside the n -dimensional “sphere” of radius r centered at x).
- **Definition:** A subset $U \subseteq M$ is an open set if, for each $x \in U$, there is some open ball $B_\epsilon(x)$ of positive radius ϵ centered at x that is contained in U . A subset $C \subseteq M$ is a closed set if its complement $M \setminus C$ is open.
 - Ultimately, one can think of open sets as “sets that are the finite or infinite union of open balls”.
 - There is much that can be said about open and closed sets, even in general metric spaces, but it is not necessary to develop any more point-set topology for our purposes.
- There are many other examples of metric spaces, and they are one of the fundamental objects of study in analysis.
 - For example, if B is the set of all functions that are bounded on the interval $[a, b]$, then $d(f, g) = \max_{a \leq t \leq b} |f(t) - g(t)|$ defines a metric on B .
 - If L^2 is the set of integrable functions f on the interval $[a, b]$ such that $\int_a^b [f(t)]^2 dt$ is finite, then $d(f, g) = \left[\int_a^b [f(t) - g(t)]^2 dt \right]^{1/2}$ defines a metric on L^2 .
 - More generally, if $\langle \cdot, \cdot \rangle$ is any inner product on a real vector space V , then the function $d(x, y) = \sqrt{\langle x - y, x - y \rangle}$ is a metric on V .
 - The idea of a metric space provides a way to study general features of convergent sequences in an abstract and general way. In particular, one can apply the results to study spaces of functions, sequences of functions, the different types of convergence of sequences of functions in various settings, and so forth.
 - Our present goal is not to develop all of real analysis, so we will stop our discussion here. Instead, we use these examples as motivation for the definition of the metric on the sequence space:
- **Proposition (Sequence Space Metric):** If Σ_2 is the sequence space on two symbols, then the distance function $d(\mathbf{x}, \mathbf{y}) = \sum_{i=0}^{\infty} \frac{|x_i - y_i|}{2^i}$ is a metric on Σ_2 , where $\mathbf{x} = (x_0 x_1 x_2 \dots)$ and $\mathbf{y} = (y_0 y_1 y_2 \dots)$.
 - Observe that the series is bounded by the geometric series $\sum_{i=0}^{\infty} \frac{1}{2^i} = 2$, so it always converges.
 - **Examples:** If $\mathbf{x} = (111111\dots)$, $\mathbf{y} = (000000\dots)$, and $\mathbf{z} = (110110\dots)$, then

$$d(\mathbf{x}, \mathbf{y}) = 1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \frac{1}{32} = \sum_{i=0}^{\infty} \frac{1}{2^i} = 2,$$

$$d(\mathbf{x}, \mathbf{z}) = 0 + 0 + \frac{1}{4} + 0 + 0 + \frac{1}{32} + \dots = \sum_{i=0}^{\infty} \left(\frac{1}{2^{3i+2}} \right) = \frac{4}{7},$$
 and

$$d(\mathbf{y}, \mathbf{z}) = 1 + \frac{1}{2} + 0 + \frac{1}{8} + \frac{1}{16} + 0 + \dots = \sum_{i=0}^{\infty} \left(\frac{1}{2^{3i}} + \frac{1}{2^{3i+1}} \right) = \frac{12}{7}.$$
 - **Proof (of proposition):** Clearly, $d(\mathbf{x}, \mathbf{y}) \geq 0$ and equality occurs only when $\mathbf{x} = \mathbf{y}$, since all terms in the series are nonnegative, and they are all zero only when $\mathbf{x} = \mathbf{y}$.
 - It is also obvious that $d(\mathbf{x}, \mathbf{y}) = d(\mathbf{y}, \mathbf{x})$ since $|x_i - y_i| = |y_i - x_i|$.
 - Finally, if $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are in Σ_2 , then by the usual triangle inequality for real numbers we see $|x_i - z_i| \leq |x_i - y_i| + |y_i - z_i|$ for each i : then summing the appropriate terms yields $d(\mathbf{x}, \mathbf{z}) \leq d(\mathbf{x}, \mathbf{y}) + d(\mathbf{y}, \mathbf{z})$ as required.

- Intuitively, from the definition of the metric, we can see that the earlier terms in an element of Σ_2 matter much more than the later terms, similarly to the behavior of digits in the decimal expansion of a real number.
- **Proposition** (Nearby Sequences): If $\mathbf{x} = (x_0x_1x_2\cdots)$ and $\mathbf{y} = (y_0y_1y_2\cdots)$ are elements of Σ_2 such that $x_i = y_i$ for each $i \leq n$, then $d(\mathbf{x}, \mathbf{y}) \leq 2^{-n}$. Conversely, if $d(\mathbf{x}, \mathbf{y}) < 2^{-n}$, then $x_i = y_i$ for each $i \leq n$.
 - In other words: if the early terms of two elements of Σ_2 agree, then the distance between them must be small. Conversely, if two elements of Σ_2 are close together, then their early terms must agree.
 - **Proof:** If $x_i = y_i$ for each $i \leq n$ then $d(\mathbf{x}, \mathbf{y}) = \sum_{i=0}^{\infty} \frac{|x_i - y_i|}{2^i} = \sum_{i=n+1}^{\infty} \frac{|x_i - y_i|}{2^i} \leq \sum_{i=n+1}^{\infty} \frac{1}{2^i} = \frac{1}{2^n}$. Conversely, suppose that $d(\mathbf{x}, \mathbf{y}) < 2^{-n}$. Then each of the terms $\frac{|x_i - y_i|}{2^i}$ for $i \leq n$ must be zero, otherwise that term alone would cause the sum to be at least $2^{-i} \geq 2^{-n}$; so $x_i = y_i$ for $i \leq n$.
- In order to model the dynamics of the quadratic map q_c on the space Λ , we need to introduce the function that plays the analogous role in the sequence space.
- **Definition:** The shift map $\sigma : \Sigma_2 \rightarrow \Sigma_2$ is the map defined by $\sigma(x_0x_1x_2x_3\cdots) = (x_1x_2x_3\cdots)$.
 - In other words, σ is the map that deletes the first term in the sequence, thereby shifting the remaining terms one slot to the left. The k th iterate of σ is equally simple: it deletes the first k terms.
 - **Example:** If $\mathbf{x} = (101010\cdots)$ then $\sigma(\mathbf{x}) = (010101\cdots)$ and $\sigma^2(\mathbf{x}) = (101010\cdots) = \mathbf{x}$, so \mathbf{x} is a periodic point of period 2 for σ .
 - **Example:** If $\mathbf{y} = (011111\cdots)$ then $\sigma(\mathbf{y}) = (111111\cdots)$ and $\sigma^2(\mathbf{y}) = (111111\cdots) = \sigma(\mathbf{y})$, so \mathbf{y} is an eventually fixed point for σ .
 - Notice that σ is a surjective, two-to-one map on Σ_2 : for any given $\mathbf{y} \in \Sigma_2$, there are two \mathbf{x} satisfying $\sigma(\mathbf{x}) = (y_0y_1y_2\cdots)$, namely $\mathbf{x} = (0y_0y_1y_2\cdots)$ and $(1y_0y_1y_2\cdots)$.
- It is a simple matter to write down all of the periodic points for σ : they are the periodic sequences $\mathbf{s} = (s_0s_1\cdots s_{n-1}s_0s_1\cdots s_{n-1}\cdots) = (\overline{s_0s_1\cdots s_{n-1}})$.
 - Explicitly: it is obvious that any such sequence satisfies $\sigma^n(\mathbf{s}) = \mathbf{s}$. Conversely, if $\sigma^n(\mathbf{s}) = \mathbf{s}$ then each block of n terms must repeat in precisely the given manner.
 - Thus, we see immediately that there are 2^n sequences of period dividing n for the shift map σ .
 - It is a more difficult problem to determine exactly how many n -cycles there are, since the tally 2^n counts all of them n times, and also includes all of the cycles of period dividing n (and most of these are overcounted as well).
 - It is not hard to answer the question for small n simply by writing down all the cycles. For example, there are two 3-cycles given by $\{(\overline{001}), (\overline{010}), (\overline{100})\}$ and $\{(\overline{011}), (\overline{110}), (\overline{101})\}$, and there are three 4-cycles given by $\{(\overline{0001}), (\overline{0010}), (\overline{0100}), (\overline{1000})\}$, $\{(\overline{0011}), (\overline{0110}), (\overline{1100}), (\overline{1001})\}$, and $\{(\overline{0111}), (\overline{1110}), (\overline{1101}), (\overline{1011})\}$ – but explicit listing rapidly becomes cumbersome.
 - Ultimately, determining the answer in general requires a technique from number theory known as Möbius inversion.
 - Explicitly, if we define the Möbius function as $\mu(n) = \begin{cases} 0 & \text{if } n \text{ is divisible by the square of any prime} \\ (-1)^k & \text{if } n \text{ is the product of } k \text{ distinct primes} \end{cases}$, where $\mu(1) = 1$, then the number of n -cycles for the shift map σ is equal to $\frac{1}{n} \sum_{d|n} 2^{n/d} \mu(d)$, where the sum is taken over all divisors d of n .
 - **Example:** We have $\mu(1) = 1$, $\mu(2) = -1$, and $\mu(4) = 0$, so the number of 4-cycles is $\frac{1}{4}[2^{4/1}\mu(1) + 2^{4/2}\mu(2) + 2^{4/4}\mu(4)] = \frac{1}{4}[2^4 - 2^2] = 3$, as we can see from the list above.
 - In any case, the fact remains that it is still quite easy to write down periodic points for σ , in stark contrast to the situation for the other functions we have analyzed. (Note, however, that the set of periodic points is countably infinite, so they are still comparatively rare in Σ_2 .)

- A central property of the shift map is that it is continuous:
- **Proposition** (Continuity of Shift Map): The shift map $\sigma : \Sigma_2 \rightarrow \Sigma_2$ is continuous at every point of Σ_2 .
 - In general, if $f : X \rightarrow Y$ is a map from one metric space to another (with associated metrics d_X and d_Y), we say that f is continuous at a point $a \in X$ when, for any $\epsilon > 0$, there exists a $\delta > 0$ such for any $x \in X$ satisfying $d_X(x, a) < \delta$ it is true that $d_Y(f(x), f(a)) < \epsilon$.
 - This definition is simply a formalization of the intuitive notion of continuity saying that if x and a are close together in X , then $f(x)$ and $f(a)$ must also be close together in Y .
 - If $f : X \rightarrow Y$ is continuous at every point in X , then (as usual) we will simply say f is a continuous function.
 - In the case where $X = Y = \mathbb{R}$ with the Euclidean metric, this reduces down to the usual definition of a continuous real-valued function: f is continuous at a when, for any $\epsilon > 0$, there exists a $\delta > 0$ such that the statement $|x - a| < \delta$ implies $|f(x) - f(a)| < \epsilon$.
 - **Proof:** Suppose we are given $\epsilon > 0$ and a point $\mathbf{a} = (a_0 a_1 a_2 \dots)$.
 - Since $\epsilon > 0$, there exists a positive integer n such that $2^{-n} < \epsilon$: we then claim that $\delta = 2^{-(n+1)}$ will satisfy the definition of continuity.
 - By the “nearby sequences” proposition, if $\mathbf{x} = (x_0 x_1 x_2 \dots)$ and $d(\mathbf{x}, \mathbf{a}) < 2^{-(n+1)}$, then $x_i = a_i$ for $i \leq n + 1$.
 - But then $\sigma(\mathbf{x}) = (x_1 x_2 x_3 \dots)$ and $\sigma(\mathbf{a}) = (a_1 a_2 a_3 \dots)$, so these sequences agree from the 0th term up through the n th term, so again by the nearby sequences proposition, we see that $d(\sigma(\mathbf{x}), \sigma(\mathbf{a})) \leq 2^{-n} < \epsilon$, as required.
 - Since \mathbf{a} was arbitrary, σ is continuous everywhere.

3.1.4 Equivalent Dynamical Systems: Homeomorphisms and Conjugation

- Now that we have developed some basic properties of the shift map σ on the sequence space Σ_2 , we would like to show that the behavior of σ on Σ_2 is “the same as” the behavior of the quadratic maps $q_c(x) = x^2 + c$, for $c < -2$, on the set Λ of points whose orbits remain bounded under q_c .
 - The itinerary map provides a way to relate these two systems: recall that for each $x \in \Lambda$, we defined the itinerary of x as the infinite binary string $S(x) = \{d_0 d_1 d_2 \dots\}$, where $d_i = 0$ if $f^n(x) \in I_0$ and $d_n = 1$ if $f^n(x) \in I_1$, and I_0, I_1 were the two intervals in the inverse image of $[-p_+, p_+]$ under the map q_c .
 - However, before showing that the behavior of σ on Σ_2 is “the same as” the behavior of q_c on Λ , we need to define what it means for two dynamical systems to be equivalent.
- Our first step is to define an equivalence on metric spaces:
- **Definition:** If X and Y are metric spaces, then $F : X \rightarrow Y$ is a **homeomorphism** if it is a continuous bijection whose inverse F^{-1} is also continuous. If there exists a homeomorphism between two spaces, we say they are **homeomorphic**.
 - Recall that a bijection is a one-to-one (injective) map that is also onto (surjective).
 - **Example:** The function $F(x) = \tan(x)$ mapping $X = (-\frac{\pi}{2}, \frac{\pi}{2})$ to $Y = \mathbb{R}$ is a homeomorphism, since $\tan(x)$ is clearly a continuous bijection, and its inverse $\tan^{-1}(x)$ is also continuous.
 - **Example:** The function $F(t) = (\cos(t), \sin(t))$ mapping $X = \mathbb{R}$ modulo 2π to $Y = S^1$ (the unit circle, embedded in the Cartesian plane) is a homeomorphism. It is clearly continuous and (mostly) clearly bijective, and its inverse map is given by $G(x, y)$ = the angle formed by the vector $\langle x, y \rangle$ and the positive x -axis, considered modulo 2π .
- **Proposition:** The relation “ X is homeomorphic to Y ” is an equivalence relation on of metric spaces.
 - In other words, if $X \sim Y$ denotes “there exists a homeomorphism $h : X \rightarrow Y$ ”, then \sim obeys the three properties of an equivalence relation:

1. For any X , $X \sim X$,
 2. For any X and Y , $X \sim Y$ implies $Y \sim X$, and
 3. For any X, Y, Z , $X \sim Y$ and $Y \sim Z$ together imply $X \sim Z$.
- Proof: For (1), for any metric space X , the identity map $h(x) = x$ is a homeomorphism from X to itself.
 - For (2), if $h : X \rightarrow Y$ is a homeomorphism, then $h^{-1} : Y \rightarrow X$ is also a homeomorphism (since the inverse of a bijection is a bijection).
 - Finally, for (3), if $h_1 : X \rightarrow Y$ and $h_2 : Y \rightarrow Z$ are homeomorphisms, then $h_2 \circ h_1 : X \rightarrow Z$ is also a homeomorphism, since the composition of continuous functions is continuous and the composition of bijections is a bijection.
- Homeomorphisms provide a topologically natural equivalence for metric spaces because they preserve the topological structure of a metric space.
 - For example, if h is a homeomorphism and $U \subseteq X$ is any subset, then $h(U)$ is open if and only if U is open, and $h(U)$ is closed if and only if U is closed.
 - Similarly, if $h : X \rightarrow Y$ is a homeomorphism and $x, x_i \in X$, the statement that $\lim_{n \rightarrow \infty} x_n = x$ is equivalent to the statement that $\lim_{n \rightarrow \infty} h(x_n) = h(x)$.
 - In the dynamical realm, we will think of a homeomorphism $h : X \rightarrow Y$ as a change of coordinates from the space X to the space Y . We would now like to extend this idea to define an equivalence of dynamical systems. Recall our earlier definition of a discrete dynamical system:
 - Definition: A (discrete) dynamical system is a pair (X, f) where X is a metric space and $f : X \rightarrow X$ is a function on X .

- Now, suppose we have a dynamical system (X, f) along with a homeomorphism $h : X \rightarrow Y$. We would like to know: what function $g : Y \rightarrow Y$ corresponds to the function $f : X \rightarrow X$ once we change coordinates using the homeomorphism h ?
- In other words, given a homeomorphism $h : X \rightarrow Y$ and a map $f : X \rightarrow X$, how can we get a function $g : Y \rightarrow Y$?
- One way to do this is to observe that the inverse homeomorphism h^{-1} is a map from $Y \rightarrow X$, so to create a map $g : Y \rightarrow Y$ using f , we can first apply h^{-1} to obtain a point in X , then apply f , and then apply h to obtain a new point in Y .
- This suggests the correct definition is to take $g = h \circ f \circ h^{-1}$.
- A (slightly) more sensible way to organize this information is with a commutative diagram, where we draw each of the relevant spaces and maps:

$$\begin{array}{ccc} X & \xrightarrow{f} & X \\ \downarrow h & & \downarrow h \\ Y & \xrightarrow{g} & Y \end{array}$$

- The idea is that if we start with some $x \in X$, then following either path of arrows down to Y in the lower right should yield the same result. (In other words, the paths “commute”.)
- In other words, we want it to be the case that $h(f(x)) = g(h(x))$ for every $x \in X$.
- Since h is invertible (and a bijection) by setting $y = h(x)$ we can rephrase this as $g(y) = h(f(h^{-1}(y)))$, and this is equivalent to the expression we obtained before.
- We can see how this works more clearly if we reverse the direction of one of the arrows in the diagram:

$$\begin{array}{ccc} X & \xrightarrow{f} & X \\ \uparrow h^{-1} & & \downarrow h \\ Y & \xrightarrow{g} & Y \end{array}$$

Now it is easier to see that $g(y) = h(f(h^{-1}(y)))$ by comparing the two possible paths from the lower-left Y to the lower-right Y .

- **Definition:** Two dynamical systems (X, f) and (Y, g) are conjugate if there exists a homeomorphism $h : X \rightarrow Y$ such that $g = h \circ f \circ h^{-1}$. (Or equivalently, if $h(f(x)) = g(h(x))$ for every $x \in X$.)
 - If $X = Y$, then we will simply say that the maps f and g themselves are conjugate if there is a homeomorphism $h : X \rightarrow X$ such that $g = h \circ f \circ h^{-1}$.
 - **Remark** (for those who like linear algebra): Conjugation arises in linear algebra in the context of a change of basis for a vector space. Namely, if $T : V \rightarrow V$ is a linear transformation on a finite-dimensional vector space V , then T has a matrix representation that depends on which basis is chosen for V . If the representation is A in one basis and B in another basis, with change of basis matrix M , then $A = MBM^{-1}$.
 - **Remark** (for those who like group theory): Conjugation also plays an important role in the study of groups. Explicitly: if f and h are elements of a group G , then the conjugate of f by h is the element $g = hf h^{-1}$. The conjugation action of an element on a group is a central tool in elementary group theory.
- Conjugate dynamical systems behave essentially identically: the conjugating homeomorphism converts dynamical properties of f on X to properties of g on Y . In particular, the orbit structure of f on X is the same as the orbit structure of g on Y . More explicitly:
- **Proposition** (Conjugate Orbits): Suppose (X, f) is conjugate to (Y, g) via the homeomorphism h . Then $\{x_1, \dots, x_n\}$ is an n -cycle for f if and only if $\{h(x_1), \dots, h(x_n)\}$ is an n -cycle for g . Furthermore, if the underlying metric space is \mathbb{R} in both cases, both sets are n -cycles, and f, g, h are all differentiable with h' everywhere nonzero, then the cycles have the same (weakly) attracting/repelling behavior.
 - **Proof:** Suppose that (X, f) and (Y, g) are conjugate with $g = h \circ f \circ h^{-1}$.
 - First, by a trivial induction we can see that $g^k = h \circ f^k \circ h^{-1}$ for each $k \geq 1$. (There are $k - 1$ cancellations $h^{-1} \circ h$ in the resulting expression.) Equivalently, $g^k \circ h = h \circ f^k$ for each k .
 - * In particular, $f^k(x) = x$ holds if and only if $g^k(h(x)) = h(x)$ holds. Thus, x is periodic under f with period (dividing) k if and only if $h(x)$ is periodic with period (dividing) k for g , for each $k \geq 1$.
 - * So we see that $\{x_1, \dots, x_n\}$ is an n -cycle for f if and only if $\{h(x_1), \dots, h(x_n)\}$ is an n -cycle for g .
 - Now suppose f and g are both differentiable real-valued functions, that $f(x_0) = x_0$, and that h is differentiable and that $h' \neq 0$.
 - * Applying the chain rule to $h(f(x)) = g(h(x))$ yields $h'(f(x_0))f'(x_0) = g'(h(x_0))h'(x_0)$.
 - * Since $f(x_0) = x_0$ we obtain $h'(x_0)f'(x_0) = g'(h(x_0))h'(x_0)$, and since $h' \neq 0$ everywhere we may cancel to obtain $f'(x_0) = g'(h(x_0))$.
 - * Therefore, by the attracting fixed point criterion, we see that the behavior of x_0 as a fixed point of f is the same as the behavior of $h(x_0)$ as a fixed point of g .
 - For n -cycles, we can apply the fixed-point result to f^n (since it is conjugate to g^n as we also just showed) to see that the cycle $\{x_1, \dots, x_n\}$ has the same behavior as $\{h(x_1), \dots, h(x_n)\}$
 - Finally, suppose that x is a weakly attracting fixed point of f . (An analogous argument will cover the case where x is a weakly repelling fixed point of f , and these arguments trivially extend to cover the case of n -cycles as well.)
 - * Then there is an open interval $I = (x - \epsilon, x + \epsilon)$ such that every $y \in I$ has $f^n(y) \rightarrow x$.
 - * Since h is a homeomorphism, $h(I)$ is an open set containing $h(x)$, so (by definition) it contains some open interval $J = (h(x) - \delta, h(x) + \delta)$ containing $h(x)$.
 - * For any $b \in J$, since h is a homeomorphism there is some $a \in I$ such that $h(a) = b$.
 - * Then $\lim_{n \rightarrow \infty} f^n(a) = x$ so since h is a homeomorphism, we can apply h to see that $\lim_{n \rightarrow \infty} h(f^n(a)) = h(x) = y$, and $h(f^n(a)) = g^n(h(a)) = g^n(b)$.
 - * Thus, $\lim_{n \rightarrow \infty} g^n(b) = y$ so y is a weakly attracting fixed point of g .
- Showing that two dynamical systems are conjugate requires showing the existence of a homeomorphism that relates them. In some cases, we can find a simple homeomorphism that does the trick, but in other cases, it can be a nontrivial task:

- Example: Show that the dynamical systems (\mathbb{R}, f) and (\mathbb{R}, g) are conjugate, where $f(x) = x^2 - 2x + 2$ and $g(x) = x^2$.
 - We want to find a homeomorphism $h(x)$ such that $h(f(x)) = g(h(x))$.
 - If we search for linear functions $h(x) = ax + b$, we must have $a(x^2 - 2x + 2) + b = (ax + b)^2$, so that $ax^2 - 2ax + (2a + b) = a^2x^2 + 2abx + b^2$. Solving this as an identity in x produces $a = 1, b = -1$.
 - Thus, if $h(x) = x - 1$, it is true that $h(f(x)) = (x - 1)^2 = g(h(x))$, so these two maps are indeed conjugate.
- Example: Show that any quadratic map $p(x) = a_1x^2 + a_2x + a_3$ with $a_1 \neq 0$ is conjugate to a quadratic map of the form $q_c(x) = x^2 + c$ for some value of c .
 - Like in the previous example, we will search for linear functions of the form $h(x) = ax + b$.
 - We want to ensure that $h(p(x)) = q_c(h(x))$, so we get $a(a_1x^2 + a_2x + a_3) + b = (ax + b)^2 + c$, or $aa_1 = a^2, aa_2 = 2ab$, and $aa_3 + b = b^2 + c$.
 - Thus, we can take $a = a_1$ and $b = a_2/2$: then $h(p(x)) = q_c(h(x))$ where $c = a_1a_3 + \frac{a_2}{2} - \frac{a_2^2}{4}$.
 - So, every quadratic map is conjugate to one of the maps $q_c(x) = x^2 + c$.
 - As an immediate consequence of this calculation, we see that all of our previous analysis of the family q_c (e.g., our study of orbit diagram) actually extends to the set of all quadratic polynomials.
- Example: Show that $(\mathbb{R} \bmod 1, D)$ and (S^1, g) are conjugate, where $D = \begin{cases} 2x & \text{for } 0 \leq x < 1/2 \\ 2x - 1 & \text{for } 1/2 \leq x < 1 \end{cases}$ is the doubling map and $g(\cos t, \sin t) = (\cos 2t, \sin 2t)$ is the angle doubling map on the unit circle S^1 .
 - We claim that the homeomorphism $h(x) = (\cos(2\pi x), \sin(2\pi x))$ is a conjugation between f and g .
 - To see this, we first compute $h(D(x))$: for $0 \leq x < \frac{1}{2}$, this is $(\cos(4\pi x), \sin(4\pi x))$, and for $\frac{1}{2} \leq x < 1$ this is $(\cos(4\pi x - 2\pi), \sin(4\pi x - 2\pi)) = (\cos(4\pi x), \sin(4\pi x))$. So in either case, $h(D(x)) = (\cos(4\pi x), \sin(4\pi x))$.
 - We also compute $g(h(x)) = g(\cos(2\pi x), \sin(2\pi x)) = (\cos(4\pi x), \sin(4\pi x))$.
 - Thus, since, $h(D(x)) = g(h(x))$ for all x , the two systems are conjugate.
- Since conjugate dynamical systems have equivalent properties, we can show that two dynamical systems are not conjugate by finding a dynamical property that they do not have in common.
- Example: Show that the dynamical systems (\mathbb{R}, f) and (\mathbb{R}, g) are not conjugate, where $f(x) = x^2 - x$ and $g(x) = x^3 - 3x$.
 - Observe that f has two fixed points $x = 0$ and $x = 2$, while g has three fixed points $x = 0$ and $x = \pm 2$.
 - By our proposition above, conjugate systems must have the same orbit structure. Since f and g do not, we conclude that they are not conjugate.

3.1.5 Equivalence of $q_c(x) = x^2 + c$ and the Shift Map for $c < -2$

- Our main goal now is to show that when $c < -2$ the dynamical system $(\Lambda, q_c(x))$ is conjugate to the system (Σ_2, σ) via the itinerary map. This will allow us to fully understand the dynamics of the quadratic map $q_c(x)$ by studying the much simpler shift map σ .
 - Recall our earlier notation: if $q_c(x) = x^2 + c$, then p_+ denotes the larger fixed point of q_c , I denotes the interval $[-p_+, p_+]$.
 - We saw that $q_c^{-1}(I) = I_0 \cup I_1$ was the union of two closed intervals, and we defined the set Λ to be the the points in I whose orbits do not blow up to ∞ , equal to the infinite intersection $\Lambda = \bigcap_{n=1}^{\infty} q_c^{-n}(I)$.

- Finally, we defined the itinerary of $x \in \Lambda$ as the binary sequence whose i th digit records which of I_0 and I_1 the i th iterate of x lands in.
- We start by proving our earlier assertion that the set $q_c^{-n}(I)$ is a union of 2^n intervals whose lengths tend to 0 as $n \rightarrow \infty$, and also show that our labeling of these intervals was self-consistent:
- **Theorem** (Iterated Inverse Images of I Under q_c): Let $q_c(x) = x^2 + c$ for $c < -2$ and $I = [-p_+, p_+]$ for $p_+ = \frac{1 + \sqrt{1 - 4c}}{2}$. Then the following hold:
 1. The inverse image $q_c^{-n}(I)$ consists of 2^n disjoint subintervals of I , each of which is mapped homeomorphically onto I by q_c^n .
 - **Proof:** Induct on n .
 - The base case $n = 1$ is provided by the observation that $q_c^{-1}(I) = I_0 \cup I_1$, where $I_0 = [-p_+, -\sqrt{-c - p_+}]$ and $I_1 = [\sqrt{-c - p_+}, p_+]$ with $p_- = \sqrt{-c - p_+}$ the other fixed point of q_c . Since q_c is monotone on each interval with $q_c(-\sqrt{-c - p_+}) = q_c(\sqrt{-c - p_+}) = -p_+$ and $q_c(-p_+) = q_c(p_+) = p_+$, q_c maps each interval homeomorphically onto I .
 - For the inductive step, suppose that $q_c^{-n}(I)$ consists of 2^n disjoint subintervals of I each mapped bijectively onto I . Then $q_c^{-(n+1)}(I) = q_c^{-1}(q_c^{-n}(I))$ is the inverse image under q_c of these 2^n subintervals of $I_0 \cup I_1$. Since q_c is two-to-one and continuous on I , for each closed interval $J \in q_c^{-n}(I)$, we see that $q_c^{-1}(J)$ is the union of two closed intervals $J_0 \cup J_1$ (one in I_0 and the other in I_1) each of which is mapped bijectively onto J by q_c .
 - Since q_c^n maps J bijectively onto I , we conclude that q_c^{n+1} maps each of J_0 and J_1 bijectively onto I , as required.
 2. For each n -digit binary string $d_0 d_2 \cdots d_{n-1}$, if we let $I_{d_0 \cdots d_{n-1}} = \{x \in I : q_c^i(x) \in I_{d_i} \text{ for each } 0 \leq i \leq n-1\}$, then $q_c^{-1}(I_{d_0 \cdots d_{n-1}}) = I_{0d_0 \cdots d_{n-1}} \cup I_{1d_0 \cdots d_{n-1}}$. In particular, the intervals in $q_c^{-n}(I)$ are precisely the intervals $I_{d_0 \cdots d_{n-1}}$.
 - **Proof:** We show that $q_c^{-1}(I_{d_1 \cdots d_n}) = I_{0d_1 \cdots d_n} \cup I_{1d_1 \cdots d_n}$ by induction on n ; the second statement then follows immediately.
 - For the base case, we showed above that $q_c^{-1}(I) = I_0 \cup I_1$, and the notation for I_0 and I_1 is consistent with the given definition.
 - For the inductive step, we first observe that $q_c(I_{d_0 d_1 \cdots d_n}) = I_{d_1 \cdots d_n}$ for each choice of the digit d_0 : by definition, $I_{d_0 d_1 \cdots d_n}$ is the set of points $y \in I$ such that $y \in I_{d_0}$, $q_c(y) \in I_{d_1}$, \dots , $q_c^n(y) \in I_{d_n}$. By applying q_c , we see that $q_c(I_{d_0 d_1 \cdots d_n})$ is the set of points $x \in I$ such that $x \in I_{d_1}$, $q_c(x) \in I_{d_2}$, \dots , $q_c^{n-1}(x) \in I_{d_n}$: but this is simply $I_{d_1 \cdots d_n}$.
 - Furthermore, by (1) we know that $q_c^{-1}(I_{d_1 \cdots d_n})$ is the union of two closed intervals, one lying in I_0 and the other lying in I_1 , and that each of these intervals is mapped bijectively by q_c onto $I_{d_1 \cdots d_n}$. But this precisely describes the two intervals $I_{0d_1 \cdots d_n}$ and $I_{1d_1 \cdots d_n}$, so $q_c^{-1}(I_{d_1 \cdots d_n}) = I_{0d_1 \cdots d_n} \cup I_{1d_1 \cdots d_n}$ as claimed.
 3. As $n \rightarrow \infty$, then length of each interval in $q_c^{-n}(I)$ tends to zero.
 - We will prove the result under the additional assumption that $c < -(5 + 2\sqrt{5})/4 = -2.368$. (The claimed result is still true for general $c < -2$ but the proof is more intricate.)
 - **Proof:** For $c < -(5 + 2\sqrt{5})/4$, it is a straightforward computation that $|q_c'(x)| > 1$ on all of $q_c^{-1}(I)$. Now let λ be any constant larger than 1 for which $|q_c'(x)| > \lambda$ on all of $q_c^{-1}(I)$ and take $J = [a, b]$ to be any subinterval of $q_c^{-1}(I)$.
 - Then, since q_c is monotone on J (since q_c' is never zero), the endpoints of $q_c(J)$ are $q_c(a)$ and $q_c(b)$.
 - Now apply the mean value theorem on $[a, b]$: there is some constant $\gamma \in (a, b)$ for which $\frac{q_c(b) - q_c(a)}{b - a} = q_c'(\gamma)$, so $\frac{|q_c(b) - q_c(a)|}{|b - a|} = |q_c'(\gamma)| > \lambda$.
 - Equivalently, $|b - a| \leq \lambda^{-1} |q_c(b) - q_c(a)|$: thus, the length of J is at most λ^{-1} times the length of $f(J)$.
 - If we then take J to be any subinterval of $q_c^{-n}(I)$, applying this argument n times shows that the length of J is at most λ^{-n} times the length of $q_c^n(J) = I$. Since $\lambda > 1$, as $n \rightarrow \infty$ this quantity tends to zero, as claimed.

- We can now establish our result on the conjugacy of the shift map and the quadratic map.
 - Most of the hard work has already been done in proving the theorem that describes $q_c^{-n}(I)$: we just need to relate that information to the itinerary map.
 - The key observation is that if $S(x) = (d_0d_1d_2 \cdots d_n \cdots)$, then $x \in I_{d_0d_1 \cdots d_n \cdots}$.
- **Theorem** (Conjugacy of Shift Map and Quadratic Map): For any $c < -2$, let Λ be the set of points $x \in \mathbb{R}$ whose orbit under $q_c(x) = x^2 + c$ remains finite, and let $S : \Lambda \rightarrow \Sigma_2$ is the itinerary map on the sequence space Σ_2 . Then the following hold:
 1. For each $x \in \Lambda$ we have $S(q_c(x)) = \sigma(S(x))$.
 - Proof: Suppose that $x \in \Lambda$ has itinerary $S(x) = (d_0d_1d_2d_3 \cdots)$.
 - Then, by definition, $q_c^i(x) \in I_{d_i}$, so by reindexing we see that $q_c^i(q_c(x)) \in I_{d_{i+1}}$.
 - Thus, the itinerary of $q_c(x)$ is $(d_1d_2d_3 \cdots) = \sigma(S(x))$, as claimed.
 2. The map S is injective.
 - Proof: Suppose that $S(x) = S(y) = (d_0d_1d_2d_3 \cdots)$; we wish to show that $x = y$.
 - By definition, then $q_c^i(x) \in I_{d_i}$ for each $i \geq 0$. By the definition of the intervals in our earlier theorem about the iterated inverse images of I under q_c , we see that $x \in I_{d_0d_1 \cdots d_n}$ for each $n \geq 0$. Similarly, $y \in I_{d_0d_1 \cdots d_n}$ for each n .
 - Suppose $|x - y| = \epsilon > 0$. From our theorem, we know that the length of $I_{d_0d_1 \cdots d_n}$ tends to 0 as $n \rightarrow \infty$, so there is some n for which the length is less than ϵ . But this is impossible because an interval of length less than ϵ cannot contain two points of distance ϵ .
 - Thus, $|x - y| = 0$ and so $x = y$.
 3. The map S is surjective.
 - Proof: Given a digit string $(d_0d_1d_2 \cdots)$, we will construct an x such that $S(x) = (d_0d_1d_2d_3 \cdots)$.
 - Consider the infinite intersection $\bigcap_{i=0}^{\infty} I_{d_0d_1 \cdots d_i}$: this is an intersection of an infinite nested sequence of closed intervals of lengths tending to 0, so the intersection is a single point.
 - Take x to be this intersection point: then since $x \in I_{d_0d_1 \cdots d_i}$, we see that $q_c^i(x) \in I_{d_i}$. This holds for each i , so we conclude that $S(x) = (d_0d_1d_2d_3 \cdots)$, as claimed.
 4. The map S is continuous.
 - Proof: Suppose $\epsilon > 0$ and $x \in \Lambda$ are given. We need to show there exists δ such that for any y with $|x - y| < \delta$, it is true that $d(S(x), S(y)) < \epsilon$.
 - To do this, choose n such that $2^{-n} < \epsilon$, and let $S(x) = (d_0d_1 \cdots d_n d_{n+1} \cdots)$.
 - Now, $q_c^{-(n+1)}(I)$ consists of 2^{n+1} disjoint intervals, each pair of which is separated by a positive distance. Choose δ to be sufficiently small so that the interval $(x - \delta, x + \delta)$ does not intersect any interval in $q_c^{-(n+1)}(I)$ except for $I_{d_0d_1 \cdots d_n}$: this is possible because x lies in $I_{d_0d_1 \cdots d_n}$, and the endpoints of this interval are a positive distance away from each of the other intervals in $q_c^{-(n+1)}(I)$.
 - Then $|x - y| < \delta$ implies that $y \in I_{d_0d_1 \cdots d_n}$, so $S(y) = (d_0d_1 \cdots d_n e_{n+1} e_{n+2} \cdots)$ for some digits e_i .
 - By our nearby sequences proposition, we then have $d(S(x), S(y)) \leq 2^{-n} < \epsilon$, as required.
 5. The map S^{-1} is continuous.
 - Proof: First observe that S^{-1} is the map which maps the sequence $(d_0d_1d_2 \cdots)$ to the single point in the infinite intersection $\bigcap_{i=0}^{\infty} I_{d_0d_1 \cdots d_i}$.
 - Now suppose $\epsilon > 0$ and $\mathbf{x} \in \Sigma_2$ are given. We need to show there exists δ such that for any \mathbf{y} with $d(\mathbf{x}, \mathbf{y}) < \delta$, it is true that $|S^{-1}(\mathbf{x}) - S^{-1}(\mathbf{y})| < \epsilon$.
 - To do this, choose n such that the length of every interval in $q_c^{-n}(I)$ is less than ϵ . (This is possible because the lengths of these intervals tend to 0 as $n \rightarrow \infty$ as shown in the previous theorem.)
 - We claim that $\delta = 2^{-n}$ is sufficient. By our nearby sequences proposition, $d(\mathbf{x}, \mathbf{y}) < 2^{-n}$ implies that the 0th through n th terms of \mathbf{x} and \mathbf{y} agree: say they are $d_0d_1 \cdots d_{n-1}$.

- Thus, by the definition of S , we see that $S^{-1}(\mathbf{x})$ and $S^{-1}(\mathbf{y})$ both lie in $I_{d_0 d_1 \dots d_{n-1}}$. But now, by assumption, this interval has length less than ϵ : thus, $|S^{-1}(\mathbf{x}) - S^{-1}(\mathbf{y})| < \epsilon$ as required.
- 6. The map S is a homeomorphism and the dynamical systems (Λ, q_c) and (Σ_2, σ) are conjugate via S .
 - Proof: This is immediate from (1)-(5) above: (2)-(5) show that S is a homeomorphism and (1) establishes the conjugacy.
- This theorem is the origin of the name “symbolic dynamics”: it allows us to use the “symbolic” sequence space to understand the dynamics of the polynomial map q_c .
 - For example, we can write down all of the n -cycles of the shift map on Σ_2 . The fact that the shift map is conjugate to q_c then tells us that each n -cycle of σ gives rise to an n -cycle of q_c .
 - So, in particular, for any $c < -2$, the map q_c has exactly two 3-cycles and exactly three 4-cycles. (Try proving this fact analytically!)
 - More generally, since there are 2^n points of period dividing n , there is at least $2^n - 2^{n-1} - 2^{n-2} - \dots - 2 - 1 = 1$ point of period exactly n . (We gave an exact formula using Möbius inversion earlier: it is $\frac{1}{n} \sum_{d|n} 2^{n/d} \mu(d)$ where $\mu(d)$ is the Möbius function.)
 - Using the itinerary map and the results above, we can even describe how to compute these cycles numerically: a point with itinerary $(d_0 d_1 d_2 d_3 \dots d_n \dots)$ must lie in the interval $I_{d_0 d_1 d_2 d_3 \dots d_n}$. This interval can be found by computing, successively, the intervals $J_0 = I_{d_n}$, $J_1 = I_{d_{n-1} d_n}$, $J_2 = I_{d_{n-2} d_{n-1} d_n}$, \dots , $J_n = I_{d_0 d_1 \dots d_n}$, where $J_k = I_{d_{n-k}} \cap q_c^{-1}(J_{k-1})$. (Each successive interval is one of the two intervals lying in q_c^{-1} of the previous term.)

3.2 Chaotic Dynamical Systems

- In this section we will give Devaney’s definition of chaos and show that a number of the systems we have analyzed are chaotic.
 - We will note that there is no universally accepted definition for a chaotic system, but Devaney’s definition is one that is frequently used.

3.2.1 Motivation for Chaos: Properties of the Shift Map on the Sequence Space

- To motivate the definition of chaos, we will study some of the properties of the shift map σ on the sequence space Σ_2 .
- Our first observation is that, for any sequence $\mathbf{x} \in \Sigma_2$, there are periodic points of σ that are arbitrarily close to \mathbf{x} .
- Proposition: For any $\mathbf{x} \in \Sigma_2$, there is a sequence \mathbf{x}_i of periodic points for σ such that $\lim_{i \rightarrow \infty} \mathbf{x}_i = \mathbf{x}$.
 - Proof: We will construct the points explicitly, so suppose $\mathbf{x} = (d_0 d_1 d_2 \dots)$.
 - Now let $\mathbf{x}_i = (\overline{d_0 d_1 d_2 \dots d_i})$. Then since the 0th through i th terms of \mathbf{x} and \mathbf{x}_i agree, by our proposition on nearby sequences we immediately have $d(\mathbf{x}, \mathbf{x}_i) \leq 2^{-i}$.
 - Also, $\lim_{i \rightarrow \infty} \mathbf{x}_i = \mathbf{x}$, since the first n digits of \mathbf{x}_n do not change in any subsequent \mathbf{x}_i .
- An equivalent way to state the previous proposition involves the topological notion of “denseness”:
- Definition: If X is a metric space, then we say a subset S is dense in X if for any $x \in X$, there is a sequence of elements $s_i \in S$ such that $\lim_{i \rightarrow \infty} s_i = x$.
 - Example: The set of rational numbers \mathbb{Q} is a dense subset of \mathbb{R} , since every real number can be exhibited as the limit of a sequence of rational numbers. (For example, take the sequence of truncations of its decimal expansion in base 10, or any other base for that matter.)

- Example: The set of irrational numbers is also a dense subset of \mathbb{R} , since every real number x can be exhibited as the limit of a sequence of irrational numbers. (For example, take any sequence of rational numbers converging to $x - \sqrt{2}$, and add $\sqrt{2}$ to each of them.)
- Example: The set of periodic points for σ is a dense subset of Σ_2 , as we just showed.
- Non-Example: The set $\{1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}, \dots\}$ is not dense in the interval $[0, 1]$. For example, there are no points in this set that are within a distance 0.2 of $\frac{3}{4}$. Indeed, much more is true: the only limit points of this set (i.e., points that are the limit of some sequence of elements in the set) are 0 and the points already in the set.
- Although it may seem like a dense set must be “large” if it is to be close to every point in the metric space X , this is not necessarily the case. For example, even though \mathbb{Q} is a dense subset of \mathbb{R} , we can cover \mathbb{Q} with a union of open intervals whose total length is arbitrarily small.
 - Explicitly: let $\epsilon > 0$, list the rationals as r_1, r_2, r_3, \dots , and take I_i to be the open interval of length $2^{-i}\epsilon$ around r_i for each $i \geq 1$. Then the infinite union $\bigcup_{i=1}^{\infty} I_i$ contains \mathbb{Q} and is therefore dense in \mathbb{R} , but the sum of the lengths of all the intervals is $\sum_{i=1}^{\infty} 2^{-i}\epsilon = \epsilon$, which can be arbitrarily small!
 - Interestingly, however, it is not possible to cover the set of *irrational* numbers with a union of open intervals having a finite sum of lengths. (The problem is that the set of irrational numbers is uncountable, and an uncountable sum of positive numbers is necessarily infinite.)
 - A fuller discussion of this kind of phenomenon belongs to measure theory, which is a discipline of analysis concerned with assigning a “measure”, or size, to well-behaved subsets of a metric space in a way that is both consistent and analytically useful.
- There is another interesting property of σ involving denseness:
- Proposition: There exists an element $\mathbf{s} \in \Sigma_2$ such that the orbit of \mathbf{s} under σ is dense in Σ_2 .
 - Proof: We claim that the element $\mathbf{s} = (\underbrace{01}_{\text{length 1}} \underbrace{00011011000001 \dots 111 \dots}_{\text{length 2}} \dots)$, constructed by listing all sequences of length 1, then all sequences of length 2, then all sequences of length 3, and so forth, has a dense orbit.
 - To see this, observe that by applying an appropriate shift map σ^k to \mathbf{s} , we can obtain a sequence whose 0th through n th terms form any desired sequence (since any block of length $n+1$ appears in the expansion of \mathbf{s}).
 - Thus, for any $\mathbf{x} = (d_0 d_1 d_2 \dots)$ and any n , there is an integer a_n such that the 0th through n th terms of $\sigma^{a_n}(\mathbf{s})$ are $d_0 d_1 \dots d_n$: then $d(\mathbf{x}, \sigma^{a_n}(\mathbf{s})) \leq 2^{-n}$.
 - In other words: for any \mathbf{x} , there is a sequence of terms in the orbit of \mathbf{s} that approach \mathbf{x} . Since \mathbf{x} was arbitrary, this means the orbit of \mathbf{s} is dense in Σ_2 , as claimed.
- Related to the existence of a dense orbit is the idea of transitivity:
- Definition: A dynamical system (X, f) is transitive if, for every x and y in X and any $\epsilon > 0$ there is a $z \in X$ whose orbit contains a point within ϵ of x and another point within ϵ of y .
 - A dynamical system (X, f) with a dense orbit is necessarily transitive, since (by definition) the dense orbit contains points within ϵ of any point of X .
 - Thus in particular, (Σ_2, σ) is transitive.
- The converse of the above statement, that a transitive dynamical system necessarily contains a point with a dense orbit, is also true if X is a compact metric space. (We will not actually use this result, but it is useful as a general fact.)
 - The proof of this statement is not easy: it is essentially nonconstructive and invokes a technical result known as the Baire category theorem.

- A metric space is compact if any open covering possesses a finite subcover. Explicitly: X is compact if, for any collection of open sets $\{U_i\}_{i \in I}$ indexed by some set I with the property that $X = \bigcup_{i \in I} U_i$, then there is some finite collection U_{i_1}, \dots, U_{i_n} such that $X = U_{i_1} \cup \dots \cup U_{i_n}$.
 - The Heine-Borel theorem states that a subset of \mathbb{R}^n is compact if and only if it is closed and totally bounded (i.e., if the set is closed and contained in a ball of some finite radius). Thus, for example, the interval $[0, 1]$ and the Cantor ternary set are both compact.
- Example: Show that the doubling map $D(x) = \begin{cases} 2x & \text{for } 0 \leq x < 1/2 \\ 2x - 1 & \text{for } 1/2 \leq x < 1 \end{cases}$ is transitive on $[0, 1]$.
 - Suppose x and y are elements in $[0, 1]$ and $\epsilon > 0$ is given. We will find an element $z \in [0, 1]$ whose orbit comes within ϵ of both x and y .
 - To do this, choose n with $2^{-n} < \epsilon$, and write $x = 0.x_1x_2\dots x_nx_{n+1}\dots$ and $y = 0.y_1y_2\dots y_n\dots$, with each expression taken in base 2.
 - Then let $z = 0.x_1x_2\dots x_nx_{n+1}y_1y_2\dots y_n$.
 - Then $|x - z| \leq \sum_{i=n+2}^{\infty} \frac{|x_i - y_{i-n-1}|}{2^i} \leq \sum_{i=n+2}^{\infty} \frac{2}{2^i} = 2^{-n} < \epsilon$.
 - Also, notice that $D^{n+1}(z) = 0.y_1y_2\dots y_n$, so $|y - D^{n+1}(z)| = \sum_{i=n+1}^{\infty} \frac{y_i}{2^i} \leq \sum_{i=n+1}^{\infty} \frac{1}{2^i} = 2^{-n} < \epsilon$.
 - So the orbit of z comes within ϵ of both x and y , so D is transitive.
 - Non-example: Show that the dynamical system (\mathbb{R}, f) where $f(x) = x^2$ is not transitive.
 - We will show that there is no orbit that contains a point within 0.01 of 0.5 and of 0.4.
 - Suppose, by way of contradiction, that the orbit of $x \in \mathbb{R}$ did have this property.
 - Since $\left| \frac{f(x) - 0}{x - 0} \right| = x < 1$ for $x \in (0, 1)$, every orbit in $(0, 1)$ moves monotonically closer to 0. So in particular, the point within 0.01 of 0.5 must be before the point within 0.01 of 0.4.
 - But if $x \in (0.49, 0.51)$ then $f(x) \in (0.49^2, 0.51^2) = (0.2401, 0.2601)$, and all subsequent orbits will be even closer to 0.
 - In particular, the orbit of a point within 0.01 of 0.5 cannot contain a point within 0.01 of 0.4.
 - The last relevant property of the shift map on the sequence space is that it displays “sensitive dependence”: moving even a small distance away from a starting point will eventually introduce large changes in the orbit.
 - Definition: A dynamical system (X, f) depends sensitively on initial conditions if there is a $\beta > 0$ such that for any $x \in X$ and any $\epsilon > 0$, there is a $y \in X$ and an integer k such that $d(x, y) < \epsilon$ but $d(f^k(x), f^k(y)) > \beta$.
 - In other words, no matter which point $x \in X$ we start at, and no matter how close we start to x , there are points within that small distance of x that eventually move a distance at least as large as β from x . (Note that, importantly, the constant β is uniform and does not depend on the starting point x nor on the size of the “starting interval” ϵ around x .)
 - Even more roughly speaking, for any $x \in X$, there are points near x whose orbits eventually move far away from x . (The orbit does not have to stay far away: it may move back towards x at some later stage.)
 - Sensitive dependence is an extremely troublesome issue when performing numerical computations: if a system exhibits sensitive dependence, then small errors in the computation are likely to magnify as we compute iterates due to the phenomenon of orbits moving away from one another. As we iterate further, the errors can quite easily overwhelm the results of a computation to the point where it becomes completely inaccurate.
 - Example: Show that the shift map σ is sensitive to initial conditions at every point of Σ_2 .
 - We will show that the definition is satisfied with $\beta = 0.5$. Let $x = (d_0d_1d_2\dots)$ and choose any $\epsilon > 0$.

- Take any $y = (e_0 e_1 e_2 \dots)$ such that $y \neq x$ and $d(x, y) < \epsilon$. Then since $y \neq x$ there is some $i \geq 0$ for which $x_i \neq y_i$, as otherwise y and x would be identical.
- But then we have $d(\sigma^i(x), \sigma^i(y)) = \sum_{j=0}^{\infty} \frac{|x_{j+i} - y_{j+i}|}{2^j} \geq \frac{|x_i - y_i|}{2^0} = 1 > \beta$, meaning that the orbits of x and y eventually move at least a uniform distance β apart, as required.
- Note that we have actually proved much more than was required: we actually showed that *any* point near x has an orbit that eventually moves far away from x .
- Non-Example: Show that the map $f(x) = \frac{1}{2}x$ is not sensitive to initial conditions.
 - We easily see that $|f(x) - f(y)| = \frac{1}{2}|x - y|$, so any two points will be moved strictly closer to one another by f .
 - In particular, if $|x - y| < \epsilon$, then $|f^n(x) - f^n(y)| < \epsilon$ as well, so the distances between the orbits will never exceed any fixed bound.
- Some dynamical systems have sensitive dependence only near particular points.
- Example: Show that $f(x) = x^2$ has sensitive dependence at $x = 1$, but not at $x = 1/3$.
 - We know that any orbit starting in $(0, 1)$ is attracted to the fixed point $x = 0$, while any orbit starting in $(1, \infty)$ will diverge to $+\infty$. Thus, the map exhibits sensitive dependence at $x = 1$, since any point near 1 will eventually move a distance greater than $\beta = 0.5$ away from 1 (since it will approach either 0 or ∞).
 - The map does not exhibit sensitive dependence at $x = 1/3$, however, because all points $x, y \in (0, 1/2)$ have the property that $|f(x) - f(y)| < |x - y|$: this follows immediately from the mean value theorem and the fact that $|f'| < 1$ on this interval. Thus, any two orbits that begin within a distance ϵ from one another will never move further apart.
- More generally, if a dynamical system has a (weakly) repelling fixed point, then it has sensitive dependence near that repelling fixed point. On the other hand, there is no sensitive dependence anywhere inside the immediate attracting basin of a (weakly) attracting fixed point, since all orbits eventually approach the fixed point.

3.2.2 Devaney's Definition of Chaos

- We now give "Devaney's definition" of a chaotic dynamical system:
- Definition: A dynamical system (X, f) is chaotic if (i) has a dense set of periodic points, (ii) it is transitive, and (iii) it depends sensitively on initial conditions.
 - A chaotic system, in general, is computationally intractable (because of the sensitive dependence) and also cannot be broken down into a simpler systems (because of the transitivity, it cannot be split into two systems that can be analyzed separately). Nonetheless, it still contains some amount of regular and predictable behavior (namely, the dense set of periodic points).
 - Example: The shift map σ on Σ_2 is chaotic, since we already showed that it has a dense set of periodic points, that it is transitive, and that it depends sensitively on initial conditions.
 - Non-example: The map $f(x) = 2x$ on \mathbb{R} is not chaotic. Although it does depend sensitively on initial conditions (at every point in \mathbb{R}) it does not have a dense set of periodic points and it is also not transitive.
- Example: Show that the doubling map $D(x) = \begin{cases} 2x & \text{for } 0 \leq x < 1/2 \\ 2x - 1 & \text{for } 1/2 \leq x < 1 \end{cases}$ is chaotic on $[0, 1)$.
 - First, we check that D has a dense set of periodic points.
 - * As we have shown long ago, the periodic points for D are the rational numbers p/q with odd denominator, since D is a bijection on the set $\{0/q, 1/q, \dots, (q-1)/q\}$. Since the rationals with odd denominator are dense in $[0, 1)$, D has a dense set of periodic points.

- Second, we claim that D has a dense orbit, so (in particular) it is transitive.
 - * To do this, let $\alpha = 0.\underbrace{01}_{\text{length 1}}\underbrace{00011011000001\cdots111\cdots}_{\text{length 2}}\underbrace{\cdots}_{\text{length 3}}$ be the base-2 decimal constructed by listing all sequences of length 1, then all sequences of length 2, then all sequences of length 3, and so forth.
 - * Note that, in base 2, $D(\alpha)$ is obtained simply by deleting the first digit of the base-2 decimal expansion of α (i.e., it acts essentially as the shift map).
 - * Hence for any sequence of digits, there is a shift of α that begins with that sequence of digits.
 - * Now let $x = 0.d_1d_2d_3\dots$ and $\epsilon > 0$. We will show there is some shift of α within ϵ of x .
 - * Choose n with $2^{-n} < \epsilon$. Then there is a positive integer k such that $D^k(\alpha)$ begins as $0.d_1d_2\dots d_nd_{n+1}$, so that $D^k(\alpha)$ and x can only differ past the $n + 2$ nd decimal place.
 - * Then $|D^k(\alpha) - x| \leq \sum_{i=n+2}^{\infty} \frac{2}{2^i} < 2^{-n} < \epsilon$, as required.
- Finally, we show that D has sensitive dependence.
 - * We will show that the value $\beta = 1/3$ will satisfy the requirements of the definition.
 - * First, observe that if a, b are both in $[0, 1/2)$ or $[1/2, 1)$, then $|D(b) - D(a)| = 2|b - a|$.
 - * Also, if $a \in [0, 1/2)$ and $b \in [1/2, 1)$ then one of $|b - a|$ and $|D(b) - D(a)|$ is at least $1/3$, since if $b - a < 1/3$ then $|D(b) - D(a)| = |1 - 2(b - a)|$ is larger than $1/3$.
 - * Therefore, if x, y are any two distinct points, the value of $|D^k(y) - D^k(x)|$ will double at each stage until the points x and y land in opposite halves of $[0, 1)$, at which point either $|D^n(y) - D^n(x)|$ will exceed $1/3$ or $|D^{n+1}(y) - D^{n+1}(x)|$ will.
 - * Thus, for any two distinct points x and y , their orbits will eventually be a distance of at least $1/3$ apart after iterating some number of times.
- We will remark that the definition we have given is not the only possible definition of “chaos”, and there is no general definition of chaos that is universally accepted.
 - In colloquial usage, a system that exhibits sensitive dependence (but not necessarily the other properties) is often called “chaotic” due to its unpredictability.
 - The example of $f(x) = 2x$ on \mathbb{R} , which exhibits sensitive dependence but is nonetheless extremely predictable, suggests quite strongly that we want something stronger than sensitive dependence in order to call a system “chaotic”.
 - It is generally held that the most important aspects of a chaotic system are sensitive dependence and transitivity. Indeed, “Robinson’s definition” of chaos uses only these two conditions and discards the requirement for a dense set of periodic points.
- It has also been shown that, if X is an infinite, compact metric space, a transitive system that has a dense set of periodic points necessarily has sensitive dependence too.
- Theorem: If X is an infinite, compact metric space and $f : X \rightarrow X$ is a continuous function that has a dense set of periodic points and is transitive, then f also has sensitive dependence on initial conditions.
 - This theorem can be proven using topological arguments similar to the ones we have already given. The details are rather lengthy and not especially enlightening, however, so we will omit them.
- In fact, for continuous functions on an interval of the real line, even more is true: the existence of a transitive orbit actually implies both the existence of a dense set of periodic points and sensitive dependence, as first shown by Vellekoop and Berglund. Before showing this result we first establish a lemma:
- Lemma: Let I be an interval of \mathbb{R} , $f : I \rightarrow I$ be continuous, and suppose $J \subseteq I$ is a subinterval containing no periodic points of f . Then if $0 < m < n$ and $z, f^m(z), f^n(z)$ all lie in J , then either $z < f^m(z) < f^n(z)$ or $f^n(z) < f^m(z) < z$.
 - The idea of this lemma is that if there are no periodic points of f inside J , then the list of all of the iterates of z that land inside J must either be strictly increasing or strictly decreasing.
 - Proof: Clearly we cannot have any two of $z, f^m(z), f^n(z)$ equal, since this would yield a periodic point inside J .

- First suppose $f^m(z)$ is the largest of these three values. Then since $f^m(z) > z$, by the intermediate value theorem we must have $f^m(x) > x$ for all x in J , for otherwise $g(x) = f^m(x) - x$ would be zero somewhere inside J , contradicting the nonexistence of a periodic point. Iterating shows that $f^{km}(x) > x$ for all positive integers k . In particular, $f^{m+m(n-m)}(x) > f^m(x)$ for all x .
 - In the same way, since $f^m(z) > f^n(z)$, by the intermediate value theorem we must have $f^m(x) > f^n(x)$ for all x in J . Iterating yields $f^m(x) > f^n(x) > f^{m+2(n-m)}(x) > \dots > f^{m+k(n-m)}(x)$ for all positive integers k . In particular, $f^{m+m(n-m)}(x) < f^m(x)$ for all x .
 - But these two statements contradict one another. This means that $f^m(z)$ cannot be the largest value.
 - In the same way, by reversing all of the inequalities, we see that $f^m(z)$ cannot be the smallest value.
- **Theorem** (Transitivity Implies Chaos on Intervals): Let I be an interval of \mathbb{R} and $f : I \rightarrow I$ be a continuous function that is transitive on I . Then the periodic points of f are dense in I and f has sensitive dependence on initial conditions, so f is chaotic on I .
 - **Proof:** Suppose f is transitive but that the periodic points of f are not dense in I , so that there is an open interval J containing no periodic points.
 - By transitivity, there exists an orbit that lands at least twice in J : if we choose any two points in J that are a distance $> \epsilon$ away from the endpoints of J and from each other, for some $\epsilon > 0$, then transitivity requires some orbit to approach within ϵ of both points, which requires at least two values on the orbit to lie in J .
 - So now let $y \in J$ also have $f^m(y) \in J$: note that $y \neq f^m(y)$ since y is not periodic. Because f is continuous, there exists an open interval U containing y inside J such that U and $f^m(U)$ are disjoint. Intuitively, the reason is that as U shrinks in size around y , $f^m(U)$ also shrinks in size around $f^m(y)$ because f is continuous, so making U small enough will force U and $f^m(U)$ apart since they are intervals containing distinct points.
 - Formally, let $|y - f^m(y)| = \epsilon$. By continuity of f^m there exists δ' such that $|x - y| < \delta$ implies $|f^m(x) - f^m(y)| < \epsilon/3$. Taking $\delta = \min(\delta', \epsilon/3)$ we see that $|x - y| < \epsilon/3$ implies $|f^m(x) - f^m(y)| < \epsilon/3$. With $U = (y - \epsilon/3, y + \epsilon/3)$ we see that $f(U)$ is contained in $(f^m(y) - \epsilon/3, f^m(y) + \epsilon/3)$ and the distance between these intervals is $|f^m(y) - y| - 2\epsilon/3 = \epsilon/3 > 0$, so they are disjoint.
 - Now, by the same argument we used above applied to f^{m+1} , there is some orbit of f^{m+1} that lands at least twice in U : this yields some $z \in U$ such that $f^n(z) \in U$ for some $n > m$.
 - But both z and $f^n(z)$ are both in U but $f^m(z)$ is not: this means $f^m(z)$ cannot be between z and $f^n(z)$, contradicting the lemma above.
 - Therefore there cannot exist an interval J with no periodic points of f , so the periodic points are dense.
 - Finally, by the general result for infinite metric spaces, since f has a dense set of periodic points and a transitive orbit, f is chaotic.

3.2.3 Equivalence and Chaotic Systems, Examples

- If (X, f) is chaotic and conjugate to (Y, g) , it seems reasonable to hypothesize that (Y, g) is also chaotic. This turns out to be true if we assume some mild conditions on f and X .
 - Along the way we will prove that properties (i) and (ii) of a chaotic system (namely, having a dense set of periodic points and transitivity) are preserved by conjugation.
 - However, perhaps surprisingly, we can show right now that sensitive dependence is not necessarily preserved by conjugation!
 - As an explicit example, observe that the function $f(x) = 2x$ has sensitive dependence on $X = (0, \infty)$, but the function $g(x) = x + \ln(2)$ on $Y = \mathbb{R}$ does not have sensitive dependence. However, these maps are conjugate, via the homeomorphism $h(x) = \ln(x)$.
 - Ultimately, the issue in this example is that the space $X = (0, \infty)$ is not compact. If we require X be compact, then sensitive dependence will in fact be preserved by conjugation.
- We will require a proposition about dense sets, which is useful enough that we include it separately:

- Proposition (Continuous Image of a Dense Set): If $h : X \rightarrow Y$ is a continuous surjective map and D is a dense subset of X , then $h(D)$ is a dense subset of Y .
 - Proof: Suppose $y \in Y$. We will find a sequence of points in $h(D)$ with limit y .
 - First, since h is surjective, there exists $x \in X$ with $h(x) = y$.
 - Next, by the assumption that D is dense in X , there exists a sequence of elements $\{x_i\}_{i \geq 1} \in X$ with $\lim_{i \rightarrow \infty} x_i = x$.
 - Finally, since h is continuous, we can conclude that $\lim_{i \rightarrow \infty} h(x_i) = h(x) = y$. So $\{h(x_i)\}_{i \geq 1}$ is a sequence of points in $h(D)$ with limit y , as required.
- Theorem (Conjugacy and Chaos): If (X, f) is chaotic, $h : X \rightarrow Y$ is a continuous surjective map with $h(f(x)) = g(h(x))$ for all $x \in X$, g is continuous, and Y is infinite, then (Y, g) is also chaotic.
 - This result is stronger than merely saying that chaos is preserved by conjugation: chaos is actually preserved by any continuous surjective map that obeys the conjugation property $h(f(x)) = g(h(x))$. If h has the additional property that there is some n such that h is at most n -to-one, h is called a semi-conjugacy.
 - Proof: We will show each of the parts of the definition of a chaotic system. (Note that most of the individual sub-results do not require all of the hypotheses.)
 - If f has a dense set of periodic points then so does g :
 - * By hypothesis, f is chaotic on X so its set S of periodic points is dense in X .
 - * Since h is continuous and surjective, by the proposition on the continuous image of a dense set we see that $h(S)$ is dense in Y .
 - * Also, $h(S)$ is contained in the set of periodic points for g : by an easy induction, $h(f(x)) = g(h(x))$ implies $h(f^n(x)) = g^n(h(x))$, so if $f^n(x) = x$ then $g^n(h(x)) = h(f^n(x)) = h(x)$.
 - * Thus, g has a dense set of periodic points.
 - If f is transitive then g is transitive:
 - * Let $y_1, y_2 \in Y$ and $\epsilon > 0$. Take B_1 to be the open ball of radius ϵ around y_1 and B_2 to be the open ball of radius ϵ around y_2 . We want to show there is some point in Y whose orbit under g contains a point in B_1 and a point in B_2 .
 - * Since h is continuous and surjective, $h^{-1}(B_1)$ is an open set containing x_1 , so it contains some ball of positive radius r_1 around x_1 . Similarly, $h^{-1}(B_2)$ contains some ball of positive radius r_2 around x_2 .
 - * By the assumption that f is transitive, there is $z \in X$ such that the orbit of z contains points within a distance $\min(r_1, r_2)$ of x_1 and of x_2 .
 - * Thus, the orbit of z under f contains a point in $h^{-1}(B_1)$ and in $h^{-1}(B_2)$, so the orbit of $h(z)$ under g contains a point in B_1 and a point in B_2 , as required.
 - Finally, for sensitive dependence we invoke the theorem from earlier: by hypothesis, Y is an infinite metric space and $g : Y \rightarrow Y$ is a continuous function that has a dense set of periodic points and is transitive (as we just showed), so g also has sensitive dependence on initial conditions and is therefore chaotic.
- Using this theorem, we can quickly prove that a number of the systems we have already analyzed are chaotic:
- Example: For any $c < -2$, show that $q_c(x) = x^2 + c$ is chaotic on the set Λ .
 - We already showed that (Σ_2, σ) was chaotic and that (Σ_2, σ) is conjugate to (Λ, q_c) .
 - Also, Λ is clearly infinite, and σ and q_c are both continuous maps. So all of the requirements of the theorem hold, so (Λ, q_c) is also chaotic.
- Example: Show that the angle doubling map on the circle S^1 is chaotic.
 - The angle doubling map, as we saw earlier, is conjugate to the doubling map on \mathbb{R} modulo 1, which we already showed (directly from the definition) to be chaotic.

- Since \mathbb{R} modulo 1 is infinite and the doubling and angle doubling maps are continuous, the result is immediate.
- **Example:** Show that the map $q_{-2}(x) = x^2 - 2$ is chaotic on the interval $I = [-2, 2]$.
 - We will construct a “semi-conjugacy” between this map and the angle doubling map D on the circle S^1 .
 - Explicitly, we claim that $h(\cos t, \sin t) = 2 \cos t$ is a semi-conjugacy between (S^1, D) and (I, q_{-2}) .
 - To see this, first note that the map is continuous and surjective from the unit circle to $[-2, 2]$.
 - Also, $h(D(\cos t, \sin t)) = h(\cos 2t, \sin 2t) = 2 \cos 2t = 4 \cos^2 t - 2 = q_{-2}(2 \cos t) = q_{-2}(h(\cos t, \sin t))$.
 - So all of the conditions of the theorem are satisfied, meaning that $q_{-2}(x) = x^2 - 2$ is chaotic on the interval $I = [-2, 2]$ as claimed.
 - To motivate where this semi-conjugacy comes from, notice that the angle doubling map sends $(\cos t, \sin t) \mapsto (\cos 2t, \sin 2t)$. Now consider only what it does to the x -coordinate: it sends $\cos t \mapsto \cos 2t = 2 \cos^2 t - 1$, or, in other words, it sends the value x to the value $2x^2 - 1$. (And this is a quadratic map.)
- As a side-note, we can use the semi-conjugacy above to give explicit formulas for the periodic points of q_{-2} .
 - Explicitly, our calculations show that for $x = 2 \cos(t)$ we have $q_{-2}(x) = 2 \cos(2t)$ so by a trivial induction we see $q_{-2}^n(x) = 2 \cos(2^n t)$.
 - From basic trigonometry, we know that $\cos(a) = \cos(b)$ is equivalent to $a = \pm b + 2k\pi$ for some integer k .
 - Applying this shows that $2 \cos(t) = q_{-2}^n(x) = 2 \cos(2^n t)$ is equivalent to $2^n t = \pm t + 2k\pi$, so that $t = \frac{2k\pi}{2^n \pm 1}$ where k is an integer.
 - Thus, the points of period dividing n for q_{-2} are the values of the form $x = 2 \cos \frac{2k\pi}{2^n \pm 1}$. By periodicity, these consist of the 2^{n-1} values $x = 2 \cos \frac{2k\pi}{2^n + 1}$ for $1 \leq k \leq 2^{n-1}$ along with the 2^{n-1} values $x = 2 \cos \frac{2k\pi}{2^n - 1}$ for $0 \leq k \leq 2^{n-1} - 1$.
- **Example:** Explicitly calculate the two 3-cycles for $q_{-2}(x) = x^2 - 2$.
 - From above, the fixed points of q_{-2} are $2 \cos 0 = 2$ and $2 \cos \frac{2\pi}{3} = -1$, so the points of period 3 are $\{2 \cos \frac{2\pi}{9}, 2 \cos \frac{4\pi}{9}, 2 \cos \frac{8\pi}{9}\}$ and $\{2 \cos \frac{2\pi}{7}, 2 \cos \frac{4\pi}{7}, 2 \cos \frac{8\pi}{7}\}$, which are clearly each 3-cycles.
 - Indeed, using a computer we can factor $\frac{q_{-2}^3(x) - x}{q_{-2}(x) - x} = (x^3 - 3x + 1)(x^3 + x^2 - 2x - 1)$, and one can check that the roots of the first polynomial are $2 \cos \frac{2\pi}{9}, 2 \cos \frac{4\pi}{9}, 2 \cos \frac{8\pi}{9}$ while the roots of the second polynomial are $2 \cos \frac{2\pi}{7}, 2 \cos \frac{4\pi}{7}, 2 \cos \frac{8\pi}{7}$.

3.3 Sarkovskii’s Theorem and Applications

- In this section we will discuss Sarkovskii’s theorem, a result about the structure of periodic points of a continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$ that is both strong and surprising for its lack of hypotheses.
- We will frequently refer to periodic points of exact period n : for shorthand, we will call these period- n points.

3.3.1 The Period-3 Theorem

- In 1975, in a paper called “Period three implies chaos” (which was the first use of the word “chaos” to describe a dynamical system), Li and Yorke proved the following theorem:
- **Theorem** (“Period-3 Theorem”): Let $f : I \rightarrow \mathbb{R}$ be a continuous function defined on an interval I . If f has a period-3 point, then f has a period- k point for any $k \geq 1$.
 - We will remark first that this result requires f to be continuous: the function $f : \mathbb{R} \rightarrow \mathbb{R}$ defined as $f(x) = 1 - 1/x$ for $x \neq 0$ has the property that $f^3(x) = x$ for all $x \neq 0, 1$, so all points have period 3 (and there are no points of any other periods at all).

- Also, the result is heavily dependent on the underlying metric space being \mathbb{R} : it is not even true on the circle, as the map $(\cos t, \sin t) \mapsto (\cos(t + 2\pi/3), \sin(t + 2\pi/3))$, which is simply rotation by $1/3$ of a full circle, has every point of period 3, and thus has no points of any other periods.
- To prove this theorem, we will begin with a pair of observations:
- **Lemma 1:** If $f : I \rightarrow \mathbb{R}$ is continuous and I and J are closed intervals with $I \subseteq J \subseteq f(I)$, then f has a fixed point in I .
 - **Proof:** Let $I = [a, b]$.
 - Since $I \subseteq J \subseteq f(I)$, there is a point $c \in I$ such that $f(c) \leq a$, and there is also a point $d \in I$ such that $b \leq f(d)$. Since $c \in I$ we see that $f(c) \leq a \leq c$ and similarly $d \leq b \leq f(d)$.
 - Thus, $f(c) - c \leq 0$ and $f(d) - d \geq 0$. Applying the intermediate value theorem to $g(x) = f(x) - x$ shows that there is some point in $[c, d] \subseteq I$ where g is zero: this is a fixed point of f .
- **Lemma 2:** If $f : I \rightarrow \mathbb{R}$ is continuous and $J \subseteq f(I)$ is a closed and bounded interval, then there exists a closed and bounded interval $K \subseteq I$ with $f(K) = J$.
 - There are more “highbrow” proofs of this fact using topological properties of compact sets, but we will give a direct proof.
 - **Proof:** Let $J = [a, b]$: if $a = b$ the result is obvious so assume $a < b$.
 - By assumption, there exist $p, q \in I$ with $f(p) = a$ and $f(q) = b$. Also suppose that $p < q$ (if $p > q$ then the argument is essentially identical).
 - Let $c = \max\{x : p \leq x \leq q, f(x) = a\}$ be the point closest to q in I such that $f(c) = a$. (Such a point must exist: by the monotone convergence theorem, any monotone increasing sequence of points x_i with $f(x_i) = a$ has a limit L , and since f is continuous we obtain $f(L) = a$.)
 - Now let $d = \min\{x : \alpha \leq x \leq q, f(x) = b\}$ be the point closest to α in I such that $f(d) = b$.
 - We claim that $K = [c, d]$ has $f(K) = J$. By construction, $f(c) = a$ and $f(d) = b$ so $[a, b] \subseteq f(K)$ by the intermediate value theorem.
 - If there were some $w \in [c, d]$ with $f(w) > b$, then the intermediate value theorem would imply the existence of a point $e \in (c, w)$ such that $f(e) = b$, which is impossible. Similarly, there cannot be any $w \in [c, d]$ with $f(w) < a$. Thus, $f(K) = J$ as required.
- Now we can prove the period-3 theorem:
 - **Proof:** Suppose that $f : I \rightarrow \mathbb{R}$ has a 3-cycle $\{a, b, c\}$, where a is less than b and c : then either $a < b < c$ or $a < c < b$.
 - If $a < c < b$ then notice $g(x) = -f(-x)$ is conjugate to f (via $h(x) = -x$) and has the 3-cycle $\{-b, -c, -a\}$ where $-b < -c < -a$, so it is sufficient to treat the case where $a < b < c$. (The case $a < c < b$ is the “mirror image” of the case $a < b < c$.)
 - Let $I_0 = [a, b]$ and $I_1 = [b, c]$. Since $f(a) = b$, $f(b) = c$, and $f(c) = a$, we see that $f(I_0) \supseteq I_1$ and $f(I_1) \supseteq I_0 \cup I_1$ by the intermediate value theorem.
 - First, f must have a fixed point in I_1 by lemma 1, since $f(I_1) \supseteq I_1$.
 - Next, we show that f must have a period-2 point in I_0 .
 - * Since $f(I_0) \supseteq I_1$, by lemma 2 there is an interval B_1 in I_0 such that $f(B_1) = I_1$.
 - * Then $f^2(B_1) = f(I_1) \supseteq I_0 \supseteq B_1$, so since f^2 maps B_1 onto itself by lemma 2 we conclude that f^2 must have a fixed point y_0 in B_1 .
 - * If y_0 were a fixed point of f , then it would lie in $B_1 \cap f(B_1) = B_1 \cap I_1 \subseteq I_0 \cap I_1 = \{b\}$, but b has period 3. So y_0 must have period exactly 2.
 - Now suppose $n > 3$. We will construct a period- n point by invoking the interval mapping lemmas to show the existence of a point x_0 with $f^n(x_0) = x_0$ whose first iterate lies in I_0 but whose next $n - 2$ iterates each land in I_1 .
 - * Since $f(I_1) \supseteq I_1$, by lemma 2 there is a closed subinterval $A_1 \subseteq I_1$ such that $f(A_1) = I_1$.

- * Applying lemma 2 again, we see that there is a closed subinterval $A_2 \subseteq A_1$ such that $f(A_2) = A_1$.
 - * We can continue this procedure to construct a sequence of closed intervals $A_{n-2} \subseteq A_{n-1} \subseteq \cdots \subseteq A_2 \subseteq A_1 \subseteq I_1$ such that $f(A_k) = A_{k-1}$ for each $2 \leq k \leq n-2$, and also $f(A_1) = I_1$.
 - * Also by lemma 2, since $f(I_0) \supseteq I_1 \subseteq A_{n-2}$, there is a closed interval $A_{n-1} \subseteq I_0$ with $f(A_{n-1}) = A_{n-2}$.
 - * Again by lemma 2, since $f(I_1) \supseteq I_0 \supseteq A_{n-1}$, there is a closed interval $A_n \subseteq I_1$ with $f(A_n) = A_{n-1}$.
 - * If we put all of this together, we see that $f^n(A_n) = f^{n-1}(A_{n-1}) = \cdots = f(A_1) = I_1$.
 - * But because $A_n \subseteq I_1$, by lemma 1 we conclude that there is a fixed point x_0 of f^n lying in I_1 .
 - * Furthermore, one can check that $x_0 \in I_0$, $f(x_0) \in I_1$, $f^2(x_0) \in I_0$, $f^3(x_0) \in I_0$, \dots , $f^{n-1}(x_0) \in I_0$.
 - * If x_0 had exact period $k < n$, then we would have $f^{k+1}(x_0) = f(x_0)$, but $f^{k+1}(x_0) \in I_0$ and $f(x_0) \in I_1$, and the only intersection point of these two intervals is $I_0 \cap I_1 = \{b\}$, so it would necessarily be true that $f(x_0) = b$. But then $f^2(x_0)$ would equal c , which is not in I_0 . This is impossible, so x_0 has period exactly n .
- Example: Show that the function $f : [-1, 1] \rightarrow [-1, 1]$ defined by $f(x) = \begin{cases} x+1 & \text{for } -1 \leq x \leq 0 \\ \cos(\pi x) & \text{for } 0 < x \leq 1 \end{cases}$ has a period- n point for every $n \geq 1$.
 - First, observe that f is continuous on I , since each of the component functions is continuous and they are equal at the transition point $x = 0$.
 - The desired result then follows from the period-3 theorem, provided we demonstrate that f has a 3-cycle.
 - A very short amount of experimentation reveals that $f(-1) = 0$, $f(0) = 1$, and $f(1) = -1$, so $\{-1, 0, 1\}$ is a 3-cycle. Thus, by the period-3 theorem, f has a point of exact period n for every $n \geq 1$.
 - Example: For any $c \leq -7/4$, show that the quadratic map $q_c(x) = x^2 + c$ has a period- n point for each $n \geq 1$.
 - Since $q_c(x)$ is continuous on \mathbb{R} , by the period-3 theorem it is sufficient to show that $q_c(x)$ has a period-3 point.
 - To do this, we will show that $q_c^3(x)$ has a saddle-node bifurcation at $c = -7/4$.
 - Some algebra shows that $\frac{q_{-7/4}^3(x) - x}{q_{-7/4}(x) - x} = \frac{1}{64} (8x^3 + 4x^2 - 18x - 1)^2$, and the cubic polynomial $8x^3 + 4x^2 - 18x - 1$ has three real roots x_0 , x_1 , and x_2 given approximately by -1.747 , -0.055 , and 1.302 .
 - If we rearrange the expression as $q_{-7/4}^3(x) = x + (x - x_0)^2 \cdot \left[\frac{1}{64} (q_{-7/4}(x) - x)(x - x_1)^2(x - x_2)^2 \right]$, then it is straightforward to see that $q_{-7/4}^3(x_0) = x_0$ and $(q_{-7/4}^3)'(x_0) = 1$, and that $(q_{-7/4}^3)''(x_0) \neq 0$.
 - Finally, we can also compute $\left. \frac{\partial q_c^3}{\partial c} \right|_{c=-7/4}(x) = \frac{1}{16} (64x^6 - 304x^4 + 364x^2 - 89)$ and verify that it is nonzero at each of x_0 , x_1 , and x_2 .
 - So by the saddle-node criterion we conclude that q_c^3 has a saddle-node bifurcation at x_0 , x_1 , and x_2 when $c = -7/4$, and the bifurcation opens in the direction of negative c .
 - It can also be shown using some polynomial algebra that q_c^3 has no other fixed points where the derivative is equal to 1 (which would be the only places these 3-cycles could disappear), so these two 3-cycles will persist for all $c \leq -7/4$.
 - The example above does shed some light on the behavior of $q_c(x)$ on the interval $[-2, -7/4]$: it shows that there are infinitely many periodic points inside the interval $[-p_+, p_+]$.
 - Having infinitely many periodic points inside a finite interval does not guarantee they will be dense, but it is at least suggestive of the chaotic behavior we saw experimentally in the orbit diagram of $q_c(x)$.
 - Our computations above also help explain the “period-3 window” near $c = -7/4$: when the 3-cycle appears, it is attracting for a brief window before undergoing a period-doubling bifurcation.

3.3.2 The Sarkovskii Ordering and Sarkovskii's Theorem

- Given the period-3 theorem, it is natural to wonder how far the result extends: for example, does the existence of a 2-cycle, or a 5-cycle, or a 6-cycle, also guarantee the existence of cycles of other orders?
- It turns out that, unbeknownst to Li and Yorke, this question had already been answered in a paper of Sarkovskii published (in Russian) more than a decade earlier in 1964.
- To state the theorem, we first need to define the Sarkovskii ordering of the positive integers:
 - First, we list all of the odd integers starting with 3, followed by 2 times the odd integers starting with $2 \cdot 3$, followed by 2^2 times the odd integers starting with $2^2 \cdot 3$, and so forth, and finishing with the powers of 2 in descending order.
 - Explicitly, the ordering is

$$3 \triangleright 5 \triangleright 7 \triangleright \dots \triangleright 2 \cdot 3 \triangleright 2 \cdot 5 \triangleright 2 \cdot 7 \triangleright \dots \triangleright 2^2 \cdot 3 \triangleright 2^2 \cdot 5 \triangleright 2^2 \cdot 7 \triangleright \dots \triangleright 2^n \triangleright 2^{n-1} \triangleright \dots \triangleright 2^2 \triangleright 2 \triangleright 1.$$
 - Thus for example, under the Sarkovskii ordering we have $7 \triangleright 22$, $13 \triangleright 2$, $20 \triangleright 24$, and $14 \triangleright 32$.
 - The Sarkovskii ordering is a total ordering on the positive integers: any two distinct integers a, b either satisfy $a \triangleright b$ or $b \triangleright a$.
- Theorem (Sarkovskii): Suppose $f : I \rightarrow \mathbb{R}$ is continuous and has a period- k point. If n is any integer with $k \triangleright n$, then f also has a period- n point.
 - As an immediate corollary of Sarkovskii's theorem, we see that if f is continuous and has any cycle whose length is not a power of 2, then f in fact has infinitely many periodic points each of whose periods is a power of 2.
 - Thus, in particular, if f has only finitely many periodic points, then the periods of each of these points must be a power of 2.
 - In particular, whenever an orbit diagram indicates the existence of an attracting n -cycle when n is not a power of 2, there are actually infinitely many other periodic cycles present as well (but we cannot see them, because they are all repelling).
 - The fact that the powers of 2 appear at the end of the Sarkovskii ordering also helps provide an explanation for the period-doubling bifurcations we saw in the orbit diagrams: the emergence of an n -cycle with $n \neq 2^d$ at a parameter value λ requires 2^d -cycles already to be present for each $d \geq 1$, and the most natural way for these to arise is via a sequence of period-doubling bifurcations of a fixed point.
- We will not give the full proof of Sarkovskii's theorem, but instead give a detailed outline. The missing pieces are all elementary, in the sense of not requiring anything more than the intermediate value theorem, but the arguments are rather technically involved.
 - Proof (outline): The first step is to prove a smaller number of more specific cases, each of which can be done using arguments similar to the ones we gave in the period-3 theorem:
 - * If f has a period- m point with $m \geq 3$, then f has a period-2 point.
 - * If f has a period- m point with $m \geq 3$ odd, then f has a period- $(m + 2)$ point.
 - * If f has a period- m point with $m \geq 3$ odd, then f has a period- $2m$ point.
 - * If f has a period- m point with $m \geq 3$ odd, then f has a period-6 point.
 - Then, by applying these results to f and its iterates in an appropriate way, one can obtain all of the implications in Sarkovskii's theorem.
 - Explicitly, if f has a period- m point with $m \geq 3$ odd, repeatedly applying the second statement shows that f also has points of each odd period larger than k , and also has a point of period 6.
 - If f has a period- $2m$ point with $m \geq 3$ odd, then f^2 has a point of period m . Thus, by the above, f^2 also has a period- k point for each odd $k \geq m$, so f has a period- k or period- $2k$ point. But by the third statement, having a period- k point implies the existence of a period- $2k$ point, so in either case f has a period- $2k$ point. Also, f^2 has a period-6 point, meaning that f has a period-12 point.

- Proceeding inductively in a similar way, we can show that if f has a period- $2^d m$ point with $m \geq 3$ odd, then f also has a period- $2^d k$ point for any odd $k \geq m$, and also has a period- $2^{d+1} 3$ point. This gives all of the implications in Sarkovskii's theorem except for the ones involving points whose period is a power of 2.
- For these, suppose first that f has a period- $2^d m$ point where $m \geq 3$ is odd. Then for any k , f^{2^k} has a point of period $2^{d-k} m$ if $k < d$ and m if $k \geq d$, so in any case we see that f^{2^k} has a period-2 point, whence f has a period- 2^{k+1} point.
- Finally, suppose f has a period- 2^d point with $d \geq 2$. Then $f^{2^{d-2}}$ has a period-4 point, so by the first statement $f^{2^{d-2}}$ has a period-2 point, and therefore f has a period- 2^{d-1} point. Repeatedly applying this result gives all of the remaining implications.
- A natural followup question to Sarkovskii's theorem is: are there any implications that are missing?
 - In other words, is it possible that, even if $a \triangleright b$ in the Sarkovskii ordering, a continuous $f : I \rightarrow \mathbb{R}$ having a period- b point must necessarily also have a period- a point?
 - It turns out that the answer is no!
- Theorem (Sarkovskii Converse): For any integer $k \geq 1$, there exists a (bounded) continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$ having a period- k point, but no period- n points for any n with $n \triangleright k$. There also exists a bounded continuous f having a period- 2^d point for every $d \geq 0$ but no periodic points of any other period.
 - To prove this theorem, it suffices to construct a single example for each of the possible cases: a function with a period- b point but no period- a point for each consecutive pair of terms $a \triangleright b$ in the Sarkovskii ordering, along with a function having a period- 2^d point for every $d \geq 0$ but no periodic points of any other period.
 - We will construct the required examples using the "truncated tent maps" $T_h : [0, 1] \rightarrow [0, 1]$, which, for a parameter $h \in [0, 1]$, are defined as $T_h(x) = \begin{cases} 2x & \text{for } 0 \leq x \leq 1/(2h) \\ h & \text{for } 1/(2h) \leq x \leq 1 - 1/(2h) \\ 2 - 2x & \text{for } 1 - 1/(2h) \leq x \leq 1 \end{cases}$.
 - Proof: Let $T(x) = T_1(x)$, and observe first (by an easy induction) that $T^n(x)$ maps each interval of the form $[\frac{k}{2^n}, \frac{k+1}{2^n}]$ monotonically onto the interval $[0, 1]$. Thus, T^n has 2^n fixed points (one in each such interval) and so T has at least one period- n point for each $n \geq 1$.
 - If $\{x_1, x_2, \dots, x_m\}$ is an m -cycle for T , and $a = \max(x_1, \dots, x_m)$, then because the only points x for which $T_a(x) \neq T(x)$ are those where $T(x) > a$, we see that $\{x_1, x_2, \dots, x_m\}$ remains an m -cycle for T_a , but is not an m -cycle for T_b for any $b < a$ since $\max(x_1, \dots, x_m)$ is no longer in the range of T_b .
 - Now fix m , and, for any m -cycle \mathcal{O} of T , let $h_{\mathcal{O}}$ denote the largest element in that m -cycle. Define $a = \min\{h_{\mathcal{O}}\}$, where the minimum is taken over all m -cycles of T . We claim that the map T_a has the property that if $k \triangleright m$, then T_a has no k -cycle.
 - * For example, if $m = 3$, one can compute that T has the two 3-cycles $\{\frac{2}{9}, \frac{4}{9}, \frac{8}{9}\}$ and $\{\frac{2}{7}, \frac{4}{7}, \frac{6}{7}\}$, so the corresponding value of a here would be $\min(\frac{8}{9}, \frac{6}{7}) = \frac{6}{7}$.
 - First observe that T_a has a unique m -cycle, since if there were a second m -cycle for T_a then either it would have a smaller maximum element (impossible, because by assumption a was minimal) or a larger maximum element (which is also impossible because $\{x_1, \dots, x_m\}$ stops being an m -cycle for T_b when $b < \max(x_1, \dots, x_m)$).
 - Suppose the elements of this unique m -cycle are $\{z_1, \dots, z_m\}$, in increasing order. Then, by construction, $z_m = a$, and it is also straightforward to see that $z_1 = T_a(a)$.
 - We also observe that the interval $[z_1, z_m] = [T_a(a), a]$ is invariant under the map T_a : clearly the maximum value is a (which is achieved at the point preceding z_m in the cycle C , and is the largest value in the range of T_a) and the minimum value is $\min(T_a(z_1), T_a(a)) = T_a(a)$, again since the minimum value of T_a on an interval will occur on one of the endpoints.
 - Now, suppose that T_a also has a k -cycle with $k \triangleright m$ in the Sarkovskii ordering. Then all points of the k -cycle must lie in the interval $[T(a), a]$, since successively iterating T_a eventually maps all values larger than 0 into this interval. Since the k -cycle contains neither of these endpoints (which are both part of the unique m -cycle), the k -cycle is strictly contained in some smaller interval $[c, d]$.

- Now, since $k \triangleright m$, we may apply Sarkovskii's theorem on the interval $[c, d]$ to obtain the existence of an m -cycle lying in this smaller interval. This is a contradiction, however, because T_a only has one m -cycle, and it contains the points a and $T(a)$ which do not lie in the interval $[c, d]$.
- Our final task is to construct a map having a 2^d -cycle for each $d \geq 1$ but no other cycles. If we let a_{2^n} denote the value $a_{2^n} = \min\{h_{\mathcal{O}}\}$ over all 2^d -cycles \mathcal{O} and let $L = \lim_{n \rightarrow \infty} a_{2^n}$, then T_L contains a 2^d -cycle for each $d \geq 1$. It can then be shown using an argument similar to the above that T_L does not contain any other cycles. This finishes the proof.
- **Example:** Suppose $f : \mathbb{R} \rightarrow \mathbb{R}$ has a period-18 point. List all other n for which f must necessarily have a period- n point.
 - By Sarkovskii's theorem and its converse, the desired list is precisely those values which follow 18 in the Sarkovskii ordering: these are the integers of the form $2m$ with $m \geq 11$ odd, $2^d m$ with $d \geq 2$ and m odd, and all the powers of 2.
 - Equivalently, these are the integers $2m$ with $m \geq 11$ odd and the multiples of 4, along with 1 and 2.
- Sarkovskii's theorem has various interactions with our discussions of chaotic functions on an interval $[a, b]$.
 - For example, suppose that a continuous function $f : [a, b] \rightarrow [a, b]$ has a periodic cycle of any length n that is not a power of 2.
 - Then by Sarkovskii's theorem, f must possess cycles of all possible lengths following n in the Sarkovskii ordering. As there are infinitely many such values, f must have infinitely many different periodic points.
 - Since the interval $[a, b]$ is finite, the periodic points must have an accumulation point P (a point equal to a limit of a sequence of distinct periodic points) somewhere in $[a, b]$. Then any open interval containing P necessarily contains periodic points of arbitrarily large period, so the function f displays sensitive dependence near P .
 - Moreover, if the periodic points for f are actually dense in $[a, b]$, which is not entirely unreasonable given that there are infinitely many of them and they have infinitely many different possible orders, then f would have sensitive dependence everywhere and would also have a dense set of periodic points (two of our three criteria for chaotic behavior).
 - We can see, then, that the seemingly quite weak requirement of having a periodic cycle of length not a power of 2 actually forces a continuous function to have quite a lot of chaotic behavior.
- The Sarkovskii ordering also helps explain the phenomenon of repeated period-doubling bifurcations leading to chaotic behavior that we have observed in orbit diagrams.
 - Explicitly, since the Sarkovskii ordering on powers of 2 is $\dots \triangleright 2^n \triangleright 2^{n-1} \triangleright \dots \triangleright 2^2 \triangleright 2 \triangleright 1$, in order for a function in a continuous one-parameter family $f_\lambda(x)$ to have a 2^n -cycle, it must also have cycles of lengths for all lower powers of 2.
 - We can see, therefore, that as we vary the parameter continuously from a value where f_λ has no fixed points to a value where f_λ has a 2^n -cycle, cycles of lengths 1, 2, 2^2 , ..., 2^{n-1} , 2^n must appear, and they must be created in that order. The most natural way for this to occur is via a sequence of period-doubling bifurcations.
- We finish with a more philosophical remark: namely, that Sarkovskii's theorem illustrates the limitations of studying the orbit behavior of functions using purely local properties (e.g., continuity, differentiability).
 - Although the existence of cycles of various lengths appears to be a purely local phenomenon, requiring only information about the function f at various places, Sarkovskii's theorem demonstrates that understanding orbit behaviors requires understanding global properties of f (e.g., continuity everywhere along with the the existence of a 3-cycle).

Well, you're at the end of my handout. Hope it was helpful.

Copyright notice: This material is copyright Evan Dummit, 2015-2025. You may not reproduce or distribute this material without my express permission.