An analysis on the iteration of four interesting functions.

Alexander V. Gheorghiu

Thanks to

Professor P. J. Giblin, Aneurin Jackson and Eleanor Bell.

Abstract

We investigate the convergence to values and p-cycles of iterative sequences based on the Möbius transformation, the general exponential function, the quadratic function and the Newton-Raphson method. In this analysis we demonstrate some particular and general theorems applicable to iterative sequences based on functions. We further look at the types of diagrams useful for understanding sequences. Lastly we briefly discuss the nature of complex iterative sequences and the understanding they bring to the subject.

The Nuffield Foundation
The University of Liverpool

1 The Möbius Sequence

The Möbius function can be interpreted many ways, it is a prime example of how wonderful iterative sequences can be. In our analysis of it we will see how a small change in the set up can result in very different results.

Let us begin by considering the Möbius function for real numbers:

$$M(x) = \frac{ax+b}{cx+d} \tag{1}$$

Let us now define what an iterative sequence based on M is. By using the recurrence relation with an arbitrary x_0 we have the following

$$\begin{aligned} x_{n+1} &= M(x_n), \qquad \text{for} \quad n = 0, 1, 2, 3... \\ \Leftrightarrow x_{n+1} &= \frac{ax_n + b}{cx_n + d} \end{aligned}$$

In the case of $c = 0$ it is trivial to see that the iterative sequence grows unimpeded, simply being the iteration of an arithmetic sequence which is to say that every term gets bigger at a constant rate. Consequently, there is no limit to the sequence and so we are not interested in the case of linear functions, in other words we will always take $c \neq 0$. Note that it is therefore always possible to take $c = 1$ simply by dividing through by c for this reason.

The sequence is somewhat curious in its behavior. It is certainty not clear from the set up (that is the configuration of a, b, c, and d) what will happen as we iterate. For example, suppose $a = 1, b = 1, c = 1, d = 1$ with $x_0 = 1$, then it s not hard to see that $x_n = 1$ for all n (though this case is quite boring). Furthermore, if we chose another starting point, say $x_0 = -1$, every term is now -1.

However, if we make a small change to the set up, say let $b = 2$, then the first 5 terms of the sequence are:

1.500000000
1.400000000
1.416666667
1.413793103
1.414285714

The last number is recognizable as the decimal expansion of $\sqrt{2}$, in fact we can quite easily deduce that as we iterate infinitely we do end up with $\sqrt{2}$ as the continued fraction is $\sqrt{2} = [1, 2, ..., 2]$. (For a detailed explanation of continued fractions and number theory see [1]). We call sequences like this, where the terms get arbitrarily close to some value, *converging*. They will form a major part of our analysis.

More formally a convergent sequence is one in which all the terms beyond a certain point have a specified difference from a number L which would be the *limit* of the sequence. If a sequence is convergent, then for any number $\epsilon > 0$ there is an N dependent on ϵ such that for all $n \geq N$ we have $|x_n - L| < \epsilon$.

We will meet convergent sequence and other types of convergence later on, but for now suppose we simply make $a = 3$ in our set up. As we iterate we find that x_n keeps getting larger. In fact we find that $x_n \to \infty$ as $n \to \infty$. When this happens we call it a *diverging* sequence as it does not get closer to any value; an example of a family of diverging sequences based on M is the case of $c = 0$. What is surprising is that to go from converging to diverging we only changed one parameter by a little amount.

Another very simple configuration, which is similar to the first case, is $(a, b, c, d) = (1, 1, -1, 1)$ with $x_0 = 3$ and this has the really strange result

$$\begin{array}{r} 3 \\ -\ 2 \\ -\ \frac{1}{3} \\ \frac{1}{2} \\ 3 \end{array}$$

This is a cycle, more precisely a 4-cycle as the 4 successive value are different but the 5^{th} term is the same as the first, or if preferred $x_0 = x_4 = 3$. This too will be a significant part of our analysis. Once more we have a very simple set up, but an entirely different result occurs.

The different cases presented here are very intriguing and so we seek to analyse why the small differences in their configurations give such fundamentally different outcomes. We could not tell from the set up what would happen. This analysis on the Möbius sequence will attempt to deduce what the relationships between the parameters have to be so that we get the different cases seen.

1.1 Stationary Values

We saw that for some configurations of M, with certain starting values, the function never changes. That is to say the sequence is stationary at that value, we therefore name these values *stationary values*.. (We may also call them fixed points). These values for the Möbius sequence have the definition,

$$x = M(x) = \frac{ax + b}{cx + d} \tag{2}$$

So to find the stationary values of a particular set up, we must simply solve the above equation. We begin by rearranging to get the following quadratic,

$$cx^2 - (a - d)x - b = 0 \tag{3}$$

and solve it like any other. As with all quadratics there are three cases that can arise for different coefficients; either there are 0, 1 or 2 roots. It is also possible to distinguish between these cases simply by looking at the discriminant of the quadratic,

$$\Delta = (a - d)^2 + 4bc$$

From the above we can easily see that if $(a - d)^2 + 4bc \geq 0$, then there are one or two roots to the equation. These roots we will call α and β and from (1) we see that $c\alpha + d,\ c\beta + d \neq 0$ for any set up. Of course we can now easily compute these stationary values simply by solving (3). For now we shall restrict ourselves to the case where $\Delta \geq 0$ as we wish to only use real numbers.

We have now deduced the conditions and the value(s) for which the sequence is stationary. However, it should be remarked that it is only stationary for particular values of x_0, that is to say when $x_0 = \alpha, \beta$. What happens if we chose another x_0 is the subject of the next subsection.

This algebraic understanding of the values is extremely useful, but to get a complete view of what these values are we shall look at their geometric significance too. It is easy to see that if we were to graph the Möbius function in (x, y) space we would use the equation

$$y = \frac{ax + b}{cx + d} \tag{4}$$

and from this it is quite easy see that, from the definition of the stationary values given in (1), the values α and β are where the lines 4) and $y = x$ intersect;

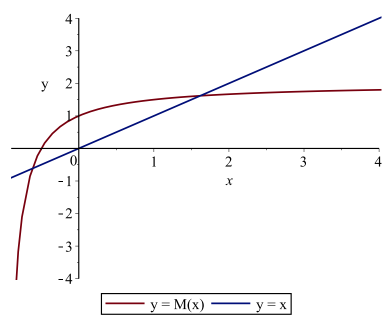

Figure 1: M has the configuration $(a, b, c, d) = (2, 1, 1, 1)$

1.2 Convergence of the Möbius Sequence

The stationary values of an iterative sequence have a very powerful relationship with what it converges to. In fact it happens that if the sequence converges, then it converges to one of its stationary values. The question we must now answer is under what conditions the sequences converges? and when it does,to which of the stationary values does it converge?

Since we know that the limit of the sequence is related to the fixed points it might be useful to relate the terms directly to the stationary values and see what happens. In fact from our definition of convergence we see that what is really important is that the value $|x_n - \alpha|$ gets very small, we can use this since if

$$\begin{aligned}
\alpha = \frac{a\alpha + b}{c\alpha + d} \quad \text{and} \quad x_{n+1} = \frac{ax_n + b}{cx_n + d},\ \text{then} \\
\Leftrightarrow x_{n+1} - \alpha &= \frac{ax_n + b}{cx_n + d} - \frac{a\alpha + b}{c\alpha + d} \\
\Leftrightarrow x_{n+1} - \alpha &= \frac{(ax_n + b)(c\alpha + d) - (a\alpha + b)(cx_n + d)}{(cx_n + d)(c\alpha + d)} \\
\Leftrightarrow x_{n+1} - \alpha &= \frac{cax_n\alpha + bd + bc\alpha + adx_n - (cax_n\alpha + bd + bcx_n + ad\alpha)}{(cx_n + d)(c\alpha + d)} \\
\Leftrightarrow x_{n+1} - \alpha &= \frac{bc\alpha + adx_n - bcx_n - ad\alpha}{(cx_n + d)(c\alpha + d)} \\
\Leftrightarrow x_{n+1} - \alpha &= \frac{ad(x_n - \alpha) - bc(x_n - \alpha)}{(cx_n + d)(c\alpha + d)} \\
\Leftrightarrow x_{n+1} - \alpha &= \frac{(ad - bc)(x_n - \alpha)}{(cx_n + d)(c\alpha + d)} \qquad (5)
\end{aligned}$$

The exact same manipulation would be true if we replaced all the αs with βs. So assuming that $x_0 \neq \beta$, as this would a stationary sequence, we divide the above by its β version.

$$\begin{aligned}
\frac{x_{n+1} - \alpha}{x_{n+1} - \beta} &= \frac{\frac{(ad-bc)(x_n-\alpha)}{(cx_n+d)(c\alpha+d)}}{\frac{(ad-bc)(x_n-\beta)}{(cx_n+d)(c\beta+d)}} \\
\Leftrightarrow \frac{x_{n+1} - \alpha}{x_{n+1} - \beta} &= \frac{\frac{(x_n-\alpha)}{(c\alpha+d)}}{\frac{(x_n-\beta)}{(c\beta+d)}} \\
\Leftrightarrow \frac{x_{n+1} - \alpha}{x_{n+1} - \beta} &= \left(\frac{c\beta + d}{c\alpha + d}\right)\left(\frac{x_n - \alpha}{x_n - \beta}\right)
\end{aligned}$$

This equation is hugely significant because it is very simple and manageable whilst telling us something very interesting about the sequence. We see that since it relates the x_n term to the x_{n+1} term it is a recursively deified sequence, however because of its form we may manipulate it further such that we get

$$\begin{aligned} \frac{x_1-\alpha}{x_1-\beta} &= \left(\frac{c\beta+d}{c\alpha+d}\right)\left(\frac{u_0-\alpha}{u_0-\beta}\right) \\ \Leftrightarrow \frac{x_2-\alpha}{x_2-\beta} &= \left(\frac{c\beta+d}{c\alpha+d}\right)\left(\frac{x_1-\alpha}{x_1-\beta}\right) \\ \Leftrightarrow \frac{x_2-\alpha}{x_2-\beta} &= \left(\frac{c\beta+d}{c\alpha+d}\right)^2\left(\frac{x_0-\alpha}{x_0-\beta}\right) \\ &\vdots \\ \Leftrightarrow \frac{x_n-\alpha}{x_n-\beta} &= \left(\frac{c\beta+d}{c\alpha+d}\right)^n\left(\frac{x_0-\alpha}{x_0-\beta}\right) \end{aligned} \tag{6}$$

This stunning result is the n^{th} term of the sequence which is an incredibly useful tool. (Sadly it is not one we will be able to use again in any of the other sequences investigated henceforth.) This is so wonderful because first we see that it has turned the Möbius function into a geometric sequence which we can understand better and second it is a very good way of analysing what happens as n gets very large which is what we are interested in as we want to know when it converges.

It is obvious from (6) that if $\left|\frac{c\beta+d}{c\alpha+d}\right| < 1$ and $n \to \infty$, then the RHS will disappear. Further we know that if $RHS \to 0$, then it *must* be true that $x_n \to \alpha$ by looking at the LHS so the sequence converges. (Taking the reciprocal and applying the same argument works just as well for β). Therefore we now know that if this weird constant ratio has size less than one, then the sequence converges to whichever stationary value is in the denominator of it.

The danger of the formula, is that we must observe that if we have implicitly taken $\alpha \neq \beta$ as this would cancel 6 down completely and give us no information. So as we progress, for now we will take the stationary values to be distinct.

We know what condition will give us a limit, but we wish to know when that condition is satisfied. For this we will leave our wonderful new formula briefly and look at the derivative of the Möbius function at the stationary value as this might illuminate what the function is doing to the sequence

$$M'(x) = \frac{ad-bc}{(cx+d)^2} \tag{7}$$

$$\Rightarrow M'(\alpha) = \frac{ad-bc}{(c\alpha+d)^2} \tag{8}$$

Already we see something interesting as we observe that the denominator of (8)is the same as the square of the denominator of the constant ratio in the formula for the n^{th} term. The numerators are very different though and so it is interesting to look at when $(c\alpha+d)(c\beta+d) = ad-bc$ as we know this means the sequence converges.

Let's return to equation (3) where we got α and β from in the first place; we find that the coefficient of the constant term and the linear term of this equation can be written as

$$\alpha\beta = \frac{-b}{c} \text{ and } \alpha+\beta = \frac{a-d}{c}$$

Using the above identities in the quadratic we may see that,

$$\begin{aligned}(c\alpha + d)(c\beta + d) &= c^2(\alpha\beta) + cd(\alpha + \beta) + d^2 \\ &= c^2(\frac{-b}{c}) + cd(\frac{a-d}{c}) + d^2 \\ &= -cb + d(a-d) + d^2 \\ &= ad - bc\end{aligned}$$

This result is truly wonderful! We have shown is that if α and β exists, this was the only assumption, then $(c\alpha + d)(c\beta + d) = ad - bc$. This means that we can rewrite the derivative in the following way,

$$\begin{aligned}M'(\alpha) &= \frac{ad - bc}{(c\alpha + d)^2} \\ &= \frac{(c\alpha + d)(c\beta + d)}{(c\alpha + d)^2} \\ &= \frac{(c\beta + d)}{(c\alpha + d)}\end{aligned} \tag{9}$$

We have now found that the constant ratio of the formula was in fact just the derivative at one of the stationary values, thus we see that a sufficient condition for the sequence to converge when there are two stationary values is that $|M'(\alpha)| < 1$, and we know this means α is the limit. Moreover, the arguments have all been independent of x_0 which tells us that if the condition on the derivative is true, and there are two distinct stationary values, the function will converge for all starting values.

Further we can see that the same arguments will work for β and that using (9) we may conclude that

$$M'(\alpha)M'(\beta) = \frac{(c\beta + d)}{(c\alpha + d)} \frac{(c\alpha + d)}{(c\beta + d)} = 1$$

This tells us that at any one point only one of the roots can have the condition for convergence. So when α and β exist and $\alpha \neq \beta$, the function will converge for all x_0 to one, but never the other, stationary value.

To complete the analysis on the convergence of the sequence based on the Möbius function we must ask what happens if α and β exist and $\alpha = \beta$? Of course this renders the wonderful formula given in (6) useless. If we review the geometric interpretation of the stationary values from the last subsection we see that this means that the line $y = x$ meets the line $y = M(x)$ only once and it is not hard to see that this happens if $y = x$ is tangent to $y = M(x)$, which is to say $M'(\alpha) = 1$ something we have not looked at before. (Note that the curious thing is that when $|M'(\alpha)| \neq 1$ it was quite simple to deduce what happens, but the case of $|M'(\alpha)| = 1$ is very strange. In fact the two cases of $M'(\alpha) = 1$ and $M'(\alpha) = -1$ are vastly different, so different that the latter case in fact belongs to section 1.3 where we discuss another kind of limit.) This time we will commence by creating a new sequence with the definition below, which using (5) can be simplified in the following way,

$$\begin{aligned} u_n &= \frac{1}{x_n - \alpha} \\ \Leftrightarrow u_{n+1} &= \frac{1}{x_{n+1} - \alpha} \\ &= \frac{(cx_n + d)(c\alpha + d)}{(ad - bc)(x_n - \alpha)} \end{aligned}$$

We see that we have some familiar factors on the RHS so by substituting (8) when $M'(\alpha) = 1$ we get that

$$\begin{aligned} u_{n+1} &= \frac{cx_n + d}{(c\alpha + d)(x_n - \alpha)} \\ &= \frac{c}{(c\alpha + d)} + \frac{1}{(x_n - \alpha)} \\ &= \frac{c}{(c\alpha + d)} + u_n \end{aligned}$$

This result is impressive because previously after some manipulation we turned the Möbius function into something very well behaved, a geometric sequence, this time we have turned it into a arithmetic progression. In fact as can be done with all arithmetic sequence we could turn this into a formula for the n^{th} term, but this is unnecessary. We have already discussed arithmetic progression and it was easy to see that they always diverges as n goes to infinity. However, from the definition of u_n we see that this can only be the case if $|x_n - \alpha| \to 0$ and this is only true if $x_n \to \alpha$ which means that the sequence converges. Once more this argument has been entirely independent of x_0 which tells us that it converges for all starting values.

Therefore, we may now conclude that the Möbius function is convergent whenever it has two stationary values or the line $y = x$ is a tangent to it. Further, we know that in both these cases it will converge for any $x_0 \in \mathbb{R}$ to its stationary value(s).

1.3 Möbius Sequence Cycles

So far when we have discussed convergence we have done so with the definition that a sequence x_n is convergent if $x_n \to k$ where k is some finite constant. However, being convergent simply means that the sequence is approaching something, this something does not have to be a number it could be a set of numbers. When a sequence converges a set of numbers we call this a cycle as we saw in the introduction. Suppose a sequence converges to a 2-cycle, then the sequence has the properly that $x_{2n} \to k_1$ and $x_{2n+1} \to k_2$ as $n \to \infty$ where $k_1 \neq k_2$. To generalise this definition, let

$$\begin{aligned} M_n(x) &= \underbrace{M(M(M(M(M(...)))))}_{\text{n-times}} \\ (&= M \circ M \circ M \circ M \circ ... \circ M) \end{aligned}$$

Then, a cycle with period length p has the property that

$$\begin{aligned} M_p(k) &= k \\ M_i(k) &\neq k \; for \; 0 < i < p \end{aligned}$$

We will take α to represent the first term of the cycle, β represents the second term, γ represents the third term of the cycle... and so on such that in a p-cycle we have $x_0 = \alpha = M_p(\alpha) = x_p$, $x_1 = \beta = M_{p+1}(\beta) = x_{p+1}$...

Function of the form $f(g(x))$ are called composite functions and they can be very cumbersome to handle. For example, trying to write $M_3(x)$ out is possible after some effort, but then to attempt to manipulate it afterwards would be extremely difficult. Therefore to understand cycles we will think about the Möbius function in a somewhat different way than previously. So far we have thought of it analytically as a function mapping one value of the sequence to the next, now we will think of it as a matrix transforming one vector to another where the vectors represent values in the sequence. This is a useful thing to do as unlike functions an n by n matrix always stays an n by n matrix after however many iterations are wanted. In fact when we think about it like this we see that the Möbius function is a transformation.

Before we begin looking into cycles, we will return briefly to convergence to specific values as we wish to understand how the matrix interpretation works. In doing this we will also gain some useful relationships and formulae for later on. We have the Möbius matrix defined as

$$M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

We then apply this matrix to a vector based on the intimal term of the sequence get a vector representing the next value,

$$\begin{aligned} M \begin{pmatrix} x_0 \\ 1 \end{pmatrix} &= \begin{pmatrix} \mu_0 \\ \mu_1 \end{pmatrix} \\ &= \begin{pmatrix} ax_0 + b \\ cx_0 + d \end{pmatrix} \end{aligned}$$

From this it is easy to see how we may compute the next value,

$$\begin{aligned} x_1 &= \frac{\mu_0}{\mu_1} \\ \Leftrightarrow M \begin{pmatrix} x_0 \\ 1 \end{pmatrix} &= \mu_1 \begin{pmatrix} x_1 \\ 1 \end{pmatrix} \end{aligned}$$

It is obvious that μ_1 is a real number dependent on x_0, but let's generalise the above statement and see that there exists a real number μ_n, dependent on x_{n+1} such that

$$M \begin{pmatrix} x_n \\ 1 \end{pmatrix} = \mu_{n+1} \begin{pmatrix} x_{n+1} \\ 1 \end{pmatrix}, \text{ where } \mu_{n+1} = cx_n + d$$

The above is quite obvious from the definition of the Möbius sequence, but let us take it even further. Suppose there was a way of short-cutting μ, with some value ν_k which has the property that

$$M^k \begin{pmatrix} x_0 \\ 1 \end{pmatrix} = \nu_k \begin{pmatrix} x_k \\ 1 \end{pmatrix} \tag{10}$$

The proof of the existence of the number ν is fairly straight forward.

Proof. Let's assume that statement (10) is true for some $n = k$ and prove that if this is the case, then it is true for $n = k + 1$.

$$\begin{aligned} M^{k+1}\begin{pmatrix} x_0 \\ 1 \end{pmatrix} &= MM^k\begin{pmatrix} x_0 \\ 1 \end{pmatrix} \\ &= M\nu_k\begin{pmatrix} x_k \\ 1 \end{pmatrix} \\ &= \nu_k M\begin{pmatrix} x_k \\ 1 \end{pmatrix} \\ &= M\nu_k\mu_k\begin{pmatrix} x_{k+1} \\ 1 \end{pmatrix} \\ &= M\nu_{k+1}\begin{pmatrix} x_{k+1} \\ 1 \end{pmatrix} \end{aligned}$$

Since ν_k and μ_k are both real numbers, we have shown that if the proposition (10) is true for k, then it is true for $k + 1$. We know that it is true for $k = 1$ as this is identical to taking $\nu_1 = \mu_1$ when we defined the transformation, hence it is true for $k = 1, 2, 3....$ QED. □

The existence of ν has remarkably let us deduce n^{th} term of the Möbius transformation in to a quite manageable formula which will come to be very helpful to us.

We will now proceed and use the matrix form to our advantage in another way, namely to find a relationship between the eigenvalues, the characteristic values, of the matrix and the stationary values of the Möbius transformation. The eigenvalues of a matrix are found by solving the characteristic equation,

$$\begin{aligned} & \det \quad (M - \lambda I) = 0 \\ \Leftrightarrow \quad & \det\left[\begin{pmatrix} a & b \\ c & d \end{pmatrix} - \begin{pmatrix} \lambda & 0 \\ 0 & \lambda \end{pmatrix}\right] = 0 \\ \Leftrightarrow \quad & \det\begin{pmatrix} (a-\lambda) & b \\ c & (d-\lambda) \end{pmatrix} = 0 \\ \Leftrightarrow \quad & (a-\lambda)(d-\lambda) - bc = 0 \\ \Leftrightarrow \quad & \lambda^2 - (a+d)\lambda + ab - dc = 0 \\ \Leftrightarrow \quad & \lambda = \frac{(a+d) \pm \sqrt{(a+d)^2 + 4bc - 4ad}}{2} \end{aligned} \tag{11}$$

We learnt in section 1.2 that to compute the fixed points of the sequence we must simply use the quadratic formula to solve (3),

$$\begin{aligned} cx^2 + (d-a)x - b &= 0 \\ \Leftrightarrow x &= \frac{(a-d) \pm \sqrt{(d-a)^2 + 4bc}}{2c} \end{aligned}$$

Note that at this stage we do not know how many stationary values there are as we have not restricted the parameters in anyway other than to say that there are real stationary values. Both the equations we have found have a similar structure, they are both cases of the quadratic formula, and they both use the same numbers. Therefore to link them there are three identities we may wish to consider,

$$\begin{aligned}(d-a)^2 &\equiv (a-d)^2\\ (a-d)^2 &\equiv (a+d)^2-4ad\\ a-d &\equiv a+d-2d\end{aligned}$$

These identities appear in our formulae and so after substituting we find that,

$$\begin{aligned}x &= \frac{(a+d-2d)\pm\sqrt{(a+d)^2+4bc-4ad}}{2c}\\ \Leftrightarrow x &= \frac{1}{c}\left(-d+\frac{(a+d)\pm\sqrt{(a+d)^2+4bc-4ad}}{2}\right)\\ \lambda &= \frac{(a+d)\pm\sqrt{(a+d)^2+4bc-4ad}}{2}\\ \Leftrightarrow x &= \frac{\lambda-d}{c} \qquad (12)\end{aligned}$$

This is really interesting, it tells us that the eigenvalues of M have a very simple relationship with the limit of the iterative sequence based on the Möbius transformation. We can certainly use this to our advantage.

We now have a sufficient understanding of the matrix interpretation of the sequence and so can finally proceed to the problem of cycle. We will begin this by addressing the promised case of $M'(\alpha)=-1$. When looking at (8) we can see that the condition on the derivative is true if and only if $ad-bc=-(c\alpha+d)^2$. The very strange thing about this is that by rearranging 12 we see that the RHS of it may be written in terms of λ fully,

$$\begin{aligned}bc-ad &= \lambda^2\\ \Leftrightarrow \lambda &= \pm\sqrt{bc-ad}\end{aligned}$$

This above tells us that there are two eigenvalues and that they are the negatives of each other, we then use this information and deduce using (11) that $a+d=0$. This is very useful because it allows us to eliminate one of the terms in matrix M, either a or b, and therefore get a simpler transformation from which we see that,

$$\begin{aligned}\begin{pmatrix}a & b\\ c & d\end{pmatrix} &= \begin{pmatrix}a & b\\ c & -a\end{pmatrix}\\ \Leftrightarrow \begin{pmatrix}a & b\\ c & -a\end{pmatrix}^2 &= \begin{pmatrix}a^2+bc & 0\\ 0 & a^2+bc\end{pmatrix}\end{aligned}$$

This result is so shocking because it says that M^2 is a *scalar* matrix meaning it has the same effect as multiplying by a constant. This property together with (10) gives us,

$$\begin{aligned}M^2\begin{pmatrix}x_0\\ 1\end{pmatrix} &= k\begin{pmatrix}x_0\\ 1\end{pmatrix}\\ &= \nu_2\begin{pmatrix}x_2\\ 1\end{pmatrix}\end{aligned}$$

It is trivial that if $x_0 = \alpha$, then this is just a 1-cycle so we will assume that is not true. In which case above we have demonstrated that $x_2 = x_0$ which is the defining property for a 2-cycle! The surprising thing here is that the only condition we set is $M'(\alpha) = -1$ which is sufficient to make a 2-cycle. Furthermore, this last bit of argument works for higher powers of M too, which is to say in order to form p-cycle we must just have the configuration of M satisfy the property

$$M^p = kI \tag{13}$$

where I is the identity matrix. As an example, let $p = 4$ which would give us the following matrix expansion

$$M^3 = \begin{pmatrix} (a^2+bc)^2 + (ab+db)(ac+dc) & (a^2+bc)^2(ab+db) + (d2+bc)(ab+db) \\ (a^2+bc)(ac+dc) + (d^2+bc)(ac+dc) & (d^2+bc)^2 + (ab+db)(ac+dc) \end{pmatrix}$$

We then use this formula to get the system of equations which we solve for a, b, c and d by setting it equal to a scalar matrix,

$$\left.\begin{aligned} (a^2+bc)^2 + (ab+db)(ac+dc) &= k \\ (d^2+bc)^2 + (ab+db)(ac+dc) &= k \\ (a^2+bc)(ab+db) + (d^2+bc)(ab+db) &= 0 \\ (a^2+bc)(ac+dc) + (d^2+bc)(ac+dc) &= 0 \end{aligned}\right\}$$

To solve this system let us first combine the last two equations and get the following formula which proves to be very exciting,

$$(a+d)(a^2+2bc+d^2) = 0$$

The first factor is the same as the condign we saw for a 2-cycle. It shows up here because a 2-cycle is an improper 4-cycle; the 0^{th} value and 4^{th} value are the same, but the 2^{nd} value is the same as the other two which contradicts the definition of a p-cycle with $p = 4$ given at the start of this section. This problem shows up whenever we try to formulate non-prime p-cycles as the factors of p always form improper p-cycles, therefore to find the condition for a p-cycle we must always deduce which factor we want first. Here we are interested in the second factor only.

When we solve for d, substitute the result into the first two equations and subtract them we find that both sides become 0. This is important as it tells us that this is a sufficient condition to satisfy the system. Let's, therefore, use this to construct a configuration which has a 4 cycle. Solving for d we get the formula

$$d = \pm\sqrt{-a^2 - 2bc}$$

This is clearly satisfied by $a = 1, b = 1, c = -1 \Rightarrow d = 1$, which gives us the cycle below when we arbitrarily chose $x_0 = 3$,

$$\begin{array}{r} 3 \\ -\ 2 \\ -\ \frac{1}{3} \\ \frac{1}{2} \\ 3 \end{array}$$

The above cycle is familiar, in fact it is the 4-cycle we met when we introduced the curious properties of the Möbius sequence. (The reason we jumped straight into the cycle rather than having to converge to it is because the stating value $x = 3$ is a solution to the system $M_4(x) = 0, M(x) \neq 0)$.

Hence, we have completed the demonstration of the one and only necessary condition (13) under which we form a p-cycle. We have applied this in an example to demonstrate how it could be used for a 4-cycle and lastly we have seen that the mysterious case we met in 1.2, where $M'(\alpha) = -1$, is in fact the only condition to form a 2-cycle.

1.4 Complex Möbius Function

We purposefully restricted ourselves to $a, b, c, d, x \in \mathbb{R}$ at the start of the investigation. In fact much more curiosity and splendor can be found in the complex plane, the Möbius transformation is one of the sequences that want to be complex. What we will see is that when we renamed it 'transformation' earlier, we were in some manner premature and in fact it is from the treatment in this section where we will really appreciate what it is doing. To begin, however, we will first introduce the concept of conjugate maps.

Suppose we had two sequences defined on different function such that

$$x_{n+1} = x_n(1 - x_n) \text{ (Sequence 1)} \tag{14}$$

$$u_{n+1} = u_n^2 + \frac{1}{4} \text{ (Sequence 2)} \tag{15}$$

These two sequences have an interesting property which is that to get from sequence 1 to sequence 2 we must simply rotate the function by π radians and then translate $(0, 0)$ to $(\frac{1}{2}, \frac{1}{2})$. This transformation if applied to the defining function of the sequences. The appropriate substitution to achieve this is $x = -u + \frac{1}{2}$;

$$\begin{aligned}
x_n &= -u_n + \frac{1}{2} \\
x_{n+1} &= -u_{n+1} + \frac{1}{2} \\
&= (-u_n + \frac{1}{2})(1 - (-u_n + \frac{1}{2})) \text{ (sequence 1)} \\
&= \frac{1}{4} - u_n^2 \\
\Leftrightarrow u_n + 1 &= u_n^2 + \frac{1}{4} \text{ (sequence 2)}
\end{aligned}$$

We have shown that by stating with the first sequence we can end up with the second sequence via this substitution, what this means is that the sequence they are *conjugate*. It follows that any information deduced about the sequence 1 can be applied to sequence

2 and vice versa. This relates to function by a simple change of notation,

$$\begin{aligned} x_{n+1} &= f(x_n), \text{ (sequence 1)} \\ T(u) &= -u + \frac{1}{2} \\ T(u_{n+1}) &= f(T(u_n)) \\ \Leftrightarrow u_n + 1 &= T^{-1}(f(T(u_n))) \\ &= g(u_n), \text{ (sequence 2)} \end{aligned}$$

This is exactly the same argument with the function notation. (Note that for the argument to valid in both cases the transforming function T, or substitution, has to be invertible and both sides of it must be continuous).

Whenever two functions f and g have this type of relationship we call them *conjugate*, they are considered to be the same as far as iterative sequences are concerned because whatever is true for one is true for the other. That is to say whatever is true about sequence 1 of function f is also true about sequence 2 of function g. A very important property to note is that $(T^{-1}fT)^n = T^{-1}f^nT$ and pacifically for our use of conjugate maps we have that $f_n = e \Leftrightarrow g_n = e$ where e is the identity and subscript n means repeated application of the function.

To continue without investigation of the Möbius transformation we will use this idea of conjugant maps. We will also restrict ourselves to the case where b is negative, $b = -p^2$ for some real constant p and secondly $a = d$ such that,

$$M(z) = \frac{az + b}{z + a}$$

The reason we apply these restriction is so that (3)only has complex roots because we force $\Delta < 0$. These roots are still the stationary values and can be computed in the same way resulting in the result that,

$$\alpha = \pm ip$$

Let us now create a transforming function T with the definition,

$$T(z) = \frac{z - ip}{z + ip}$$

Note that this transforming function is continuous as we are in $\mathbb{C} \cup \{\infty\}$ and so it does not matter that we are sending one of the stationary values to infinity. Therefore it meets all the conditions for the transform to be valid. Now to perform the transformation, after a short calculation we find that,

$$T \circ M \circ T^{-1} = \frac{a - \alpha}{a^2 + p^2} z$$

This is incredible. Once we see that the constant has modulus one we understand that the conjugate map to M is in fact just a rotation in the complex plane. Another curious thing is that the constant has the strange property that

$$M'(\alpha) = \frac{a - \alpha}{a^2 + p^2}$$

The conjugate map being a rotation is very interesting because rotations are easy to handle in the complex plane. We can apply Euler's equation with the quick result

$$\begin{aligned} e^{i\theta} &= \frac{a-\alpha}{a^2+p^2} \\ \Leftrightarrow \cos\theta &= \frac{a^2-p^2}{a^2+p^2}, \\ \Leftrightarrow \sin\theta &= -\frac{2ap}{a^2+p^2} \end{aligned}$$

At this stage we have three unknowns, θ, α and β, and two equations. However, we can chose what θ is so that we end up with a system of two equation with two variables.

For example, suppose that we want the 6^{th} iteration of M on some z to be equal to z, in functional notation $M_6(z) = z$. We see that via the conjugate map that this is identical to choosing $\theta = \frac{\pi}{3}$ as this will result in 6 iterations transforming, or turning, z through 2π radians back to where it started.

In fact this is a 6-cycle as no fewer iterations will return to z, the identity. We can now solve for the constants,

$$\begin{aligned} \cos\theta = \frac{1}{2} &= \frac{a^2-p^2}{a^2+p^2} \\ \sin\theta = \frac{\sqrt{3}}{2} &= -\frac{2ap}{a^2+p^2} \\ p = 1 &\Leftrightarrow b = -1 \\ a &= \pm\sqrt{3} \end{aligned}$$

Hence, we have found the specific conditions where we get a 6-cycle in the complex plane when $a = d$. Naturally, we could get a p-cycle in the exact same way by choosing an appropriate value for θ, this value obviously being $\theta = \frac{2\pi}{p}$.

1.5 Conclusion on Möbius Function

Overall, the Möbius sequence is an incredibly interesting. Its very nature is threefold; it can be thought of as a function mapping one value to the next; we may consider it a matrix transforming one vector to the next which represent values in the sequence; or it can be a rotation in the complex plane.

This last nature is the one which could use much further exploration than we had time for. We have so far seen that it is in some manner *easy* to make Möbius sequence make a complex cycle. In reality T simply moves the stationary values to 0 and ∞, yet the result we get is so astounding we can't help but be impressed with the ease with which this transformation results in such a simple conjugate map to the rather complicated M.

There is much further study to be carried out here which is left up to the reader. For example, what is the relationship between the rotation in the complex plane and the gradient? The sequence has shown so much promise already and so we expect that there are some truly remarkable properties yet to come.

2 The Exponential Sequence

In this section we will be looking at an iterative sequence defined on the exponential function which has base 'a'. The actual sequence we are interested in is defined recursively as

$$\begin{aligned} E(x) &= a^x \\ x_{n+1} &= a^{x_n} \text{ for } n = 0, 1, 2... \end{aligned} \tag{1}$$

In this investigation we will restrict ourselves to $a \in \mathbb{R}^+$ as some of this sequence wonder lies in its unprecedented simplicity. In general exponents can cause some real problems, especial when considering no-integer exponents as is the majority of our investigation. Yet despite expectations this sequence drops out to be extremely well ordered and aesthetic. For example, if we simply take $x_0 = 0$ we get the very predictable sequence,

$$0,\ 1,\ a,\ a^a,\ a^{a^a}... \tag{2}$$

Although one should not get the impression that this means there is a formula for the n^{th} term of the sequence, this is a privilege we are not afforded thus making it slightly more difficult to handle than the Möbius sequence.

To understand better the behaviour of this sequence let's plot a graph of a against x_n, as $n \to \infty$, so that we can see how changing a yields some quite interesting features in the limit of the sequence;

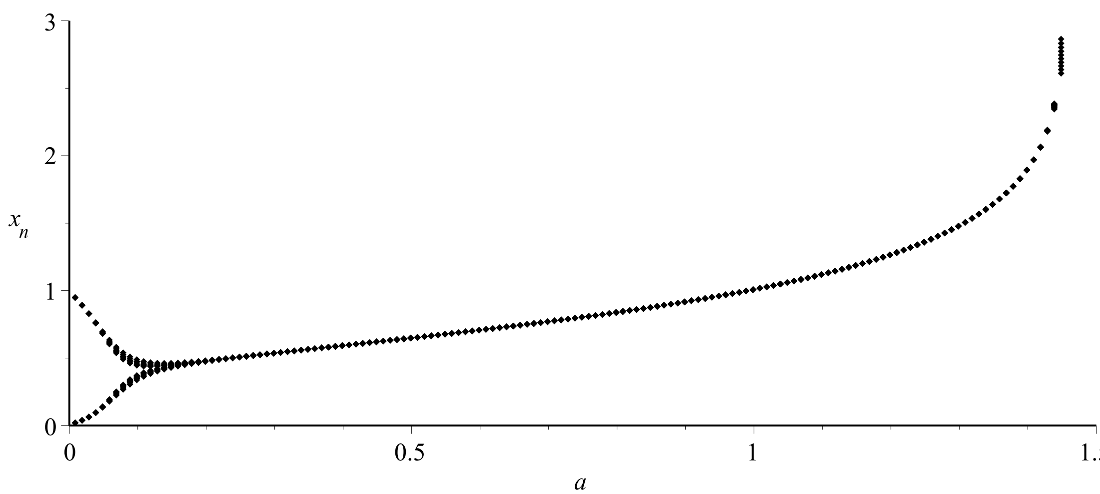

Figure 2: (Appendix A) How the limit of x_n as $n \to \infty$ changes as the parameter a changes.

Note that this graph was in fact not obtained by taking the limit but rather by doing a few hundred iterations of sequence and plotting the last 10 values. This is not analytically perfect, as can be seen by the vertical stretches of points formed in the diagram, but it is sufficiently accurate to get a good idea for what is happening (any further would go beyond our computational limit).

When we look at the above graph we see that there are three interesting regions which we should remark on;

for	$0 < a < 0.1$	there is a 2-cycle;
for	$0 < a < 1.4$	the sequence $x_n \to \alpha$;
for	$1.4 < a < 2$	the sequence $x_n \to \infty$.

The numbers on the bounds for a are approximate and so in this investigation we will be focusing on computing what they actually are: when does x_n converge? and when does it go to a cycle or a value?

The observations are useful to us as it suggests an approach for the analysis. It may seem as though we should move left to right for values of a until we find these strange changes, but in fact we see that it would be wiser to move from the right to the left as this means we wont have to convert ourselves with divergence too much later on. As it happens we will be splitting the sequence up into two different ranges of a not based on what the sequence is doing, directly, but rather on what the function $E(x)$ is doing.

2.1 Convergence

The function $E(x)$ is entirely dependent on a and has quite different properties as this parameter is varied. So much so that there are three very distinct cases;

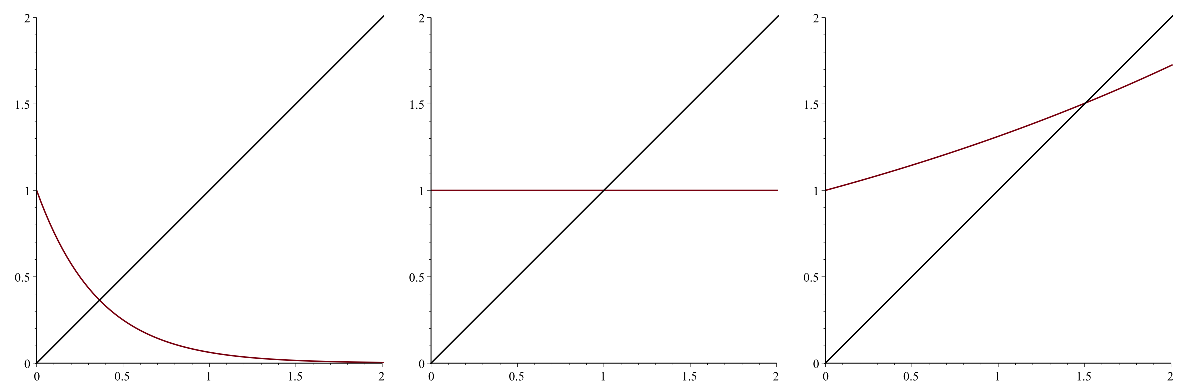

Figure 3: (A) $a < 1$ (B) $a = 1$ (C) $1 < a$

Case (B) where $a = 1$ is trivial as the sequence x_n is stationary at 1 for all $x_0 \in \mathbb{R}$ so we will ignore this value. However, the other two cases need some analysis as they exhibit some quite interesting properties. So since we can see in figure 2 that the case of (A) has the added complication of a cycle, we will begin by looking at the last case, where $a > 1$, first.

2.1.1 The range $a > 1$

Let us first observe that from image (C) we see that the function is strictly increasing as the graph is pointing upwards. Indeed, we may deduce this simply by taking the first and second derivative of the function given in (1). What we see is that not only is the function strictly increasing, but it is also convex because as it is easy to see that:

$$\begin{aligned} E'(x) &= a^x(\ln a) < 0 \\ E''(x) &= a^x(\ln a)^2 > 0 \end{aligned} \tag{3}$$

When we consider the graph of this function we see that there are three possible scenarios for the stationary values; there are two values α and β; there is only one value α; or there are no stationary values.

In this last case it is obvious that the sequence will diverge because if a sequence converges it converges to a stationary value and so by the contrapositive if there are no stationary values then the sequence does not converge. Therefore, we propose that there exists a number A with the property that

$$\begin{aligned} x_n &\to \alpha \qquad \text{for } 1 < a \leq A \\ x_n &\to \infty \qquad \text{for } a > A \end{aligned}$$

To calculate A we must simply ask what is the greatest possible stationary value the sequence can have? The corresponding parameter a would have to be A from the above definition. So we simply need to maximise A in the below equation, which defines the stationary value of the sequence x_n based on the exponential function,

$$A^x = x \tag{4}$$

This is an optimization problem: to find A we must just compute the stationary points, or rather the local maxima, the bounds of a and then compare which gives the greatest value. To find the local maxima we differentiate and solve for the stationary points,

$$\begin{aligned} A' = \frac{x^{\frac{1}{x}}}{x^2}(1 - \ln x) &= 0 \\ \Leftrightarrow \ln x &= 1 \\ \Leftrightarrow x &= e \end{aligned}$$

We see here that A only has one stationary point and when we look at the graph below we see quite clearly that this value must be a local maximum.
However, we wish to prove this analytically and to do so we must simply observe that

$$\begin{aligned} &\text{if} \quad x \to e^+ \quad, \text{ then } A' < 0 \;\; \text{as } \ln(x) \to 1^+ \text{ and,} \\ &\text{if} \quad x \to e^- \quad, \text{ then } A' > 0 \qquad \text{as } \ln(x) \to 1^- \end{aligned}$$

Another way to prove that $x = e$ is a local maximum is by taking the second derivative, but this gives a very complicated equation and is quite unnecessary. The actual value of A corresponding to he local maximum is $e^{\frac{1}{e}}$.

Now it is simply necessary to check the bounds of (4) to make sure that $x = e$ is a *global* maximum. So we ask two question: is A larger when $x \to 0^+$? and is A larger

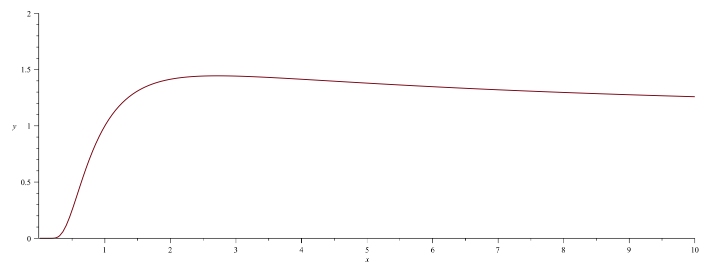

Figure 4: The function $A = x^{1/x}$

when $x \to \infty$? The first can be answered by a simple substitution,

$$\begin{aligned} u &= \frac{1}{x} \\ \Rightarrow A &= \lim_{x \to 0^+} x^{\frac{1}{x}} \\ &= \lim_{u \to \infty} \frac{1}{u}^u \\ &= \lim_{u \to \infty} \frac{1}{u^u} = 0 \end{aligned}$$

Which means that the limit is 0. This is obviously less than the value we got for the local maximum, hence the first bound is checked and the question is answered is answered in the negative.

The second case of $x \to \infty$ is slightly more difficult to work out. To answer this question we will first manipulate it to a much more manageable limit;

$$\begin{aligned} A &= \lim_{x\to\infty} x^{\frac{1}{x}} \\ &= \lim_{x\to\infty} e^{\frac{\ln x}{x}} \\ &= e^{\lim_{x\to\infty} \frac{\ln x}{x}} \end{aligned}$$

The last line is interesting because we can finally take the limit using L'Hôpital's rule as we have an indeterminate form,

$$\begin{aligned} \lim_{x\to\infty} \tfrac{\ln x}{x} &= \frac{\infty}{\infty} \\ \Rightarrow \lim_{x\to\infty} \tfrac{\ln x}{x} &= \lim_{x\to\infty} \frac{\frac{1}{x}}{1} = 0 \\ \therefore A = \lim_{x\to\infty} x^{\frac{1}{x}} &= e^0 = 1 \end{aligned}$$

This limit once again is below the value of the local maximum, hence the value $A = e^{\frac{1}{e}}$ is the global maximum. Therefore we have found the largest value of A which satisfies (4), this is the last a for which there exists a stationary value of the sequence x_n. Meaning, we have found what the largest a such that sequence possibly converges. It follows therefore that we must only look at 0< $a < e^{\frac{1}{e}}$ a much smaller range than from 0 to infinity.

We know that if x_n converges, then it converges to a stationary value. However, we do not in fact know if the sequence $x_{n+1} = E(x_n)$ converges at all.

Suppose we wanted to graphically *see* how successive values of the sequence related to each other. It is obvious that both the lines $y = x$ and $y = E(x)$ are important, and naturally we must mark all the points of the sequence on this diagram with the coordinate (x_n, x_{n+1}). Doing this would give us a plane with two lines, a collection of points and disorder. So to see the relationship between successive iterations we will connect the corresponding points to successive with lines such that we get the above diagrams.

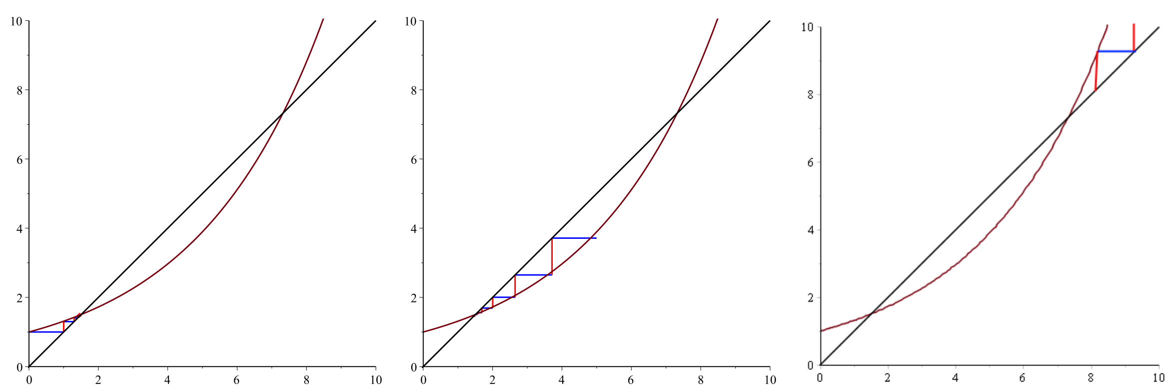

Figure 5: (Appendix B) (A) $x_0 < \alpha$ (C) $\alpha < x_0 < \beta$ (B) $x_0 < \alpha$

This type of diagram is called a *cobweb* diagram and it can be very revelling. For example we can see not only if the sequence approaches a limit, but how it does so. Certainly from the above it looks as though the sequence always converges to the lower stationary value and that to deduce this we must just consider the line $y = x$ with respect to $y = E(x)$. We see that the relative positions of these two lines results in the different cobwebs, namely one the sequence ascends and the other it descends. However, we may simplify the question a little if we observe that the behaviour in the range of $0 \leq x_0 < \alpha$ is the same as the range $\beta < x_0$. So we must just ask what happens when $0 \leq x_0 < \alpha$ and when $\alpha < x_0$.

Firstly, suppose $0 \leq x_0 < \alpha$, we see that this means that the line $y = x$ is below the line $y = E(x)$ meaning that $E(x > x \Leftrightarrow x_{n+1} > x_n$ from the definition of the sequence. This is true for all a in the range we know that function is strictly increasing and $y = E(x)$ crosses the y-axis at $(0, \alpha)$. Furthermore, we see that if $0 < x_n < \alpha$, then $1 < x_{n+1} < \alpha$ for $n > 2$ by substituting the bounds into the function. This means that the sequence is strictly increasing, but is always below the upper bound which implies it converges to the upper bound in this case that is α.

Secondly, suppose that we had $\alpha < x_0$. We see that this means $y = x$ is above the line $y = E(x)$ meaning that $x_{n+1} < x_n$. Additionally we see again that the range of x_n is the same as the range of x_{n+1}. Therefore we know the sequence will converge to the lower bound, α in this case.

We have now shown that in both cases the sequence will converge and we know what it will converge to. However, in the argument we never referred to anything that defines the function or the sequence in anyway other than to say the function was strictly increasing. Therefore we may call this a general theorem.

General Theorem 1. *Increasing Functions*
Let f be a strictly increasing function in the range $r_0 \leq x \leq r_1$, and let an iterative sequence be based on f such that $x_{n+1} = f(x_n)$. The values r are either stationary values of the sequence or are the boundary of where the function is an increasing function. The sequence x_n will converge to r_0 if the line $y = x$ is above the function and to r_1 if the lines is below. Further, x_n will always stay on the same side of the limit as x_0.

Let us note that from the above theorem if x_0 is larger than the last stationary value and the line $y = x$ is above the function f, then the sequence must diverge which is indeed in accordance with (C) from the earlier diagram.

We now fully understand what happens in the range $a > 1$. If x_0 is below the last root, then

$$\begin{array}{lll} x_n \to \alpha \text{ as } n \to \infty & \text{for} & 1 \leq a \leq e^{\frac{1}{e}} \\ x_n \to \infty \text{ as } n \to \infty & \text{for} & a > e^{\frac{1}{e}} \end{array}$$

2.1.2 The range of $0 \leq \alpha < 1$

The difference between this range and last is that here the function $E(x)$ is strictly descending, but still convex. We can easily see this because the first and second derivatives are

$$\begin{aligned} E'(x) &= a^x(\ln a) < 0 \\ E''(x) &= a^x(\ln a)^2 > 0 \end{aligned}$$

The problem we have is that General Theorem 1 no longer applies. So to understand the sequence we will need quite a different analysis than before.

Let's begin by looking at the cobweb diagram of the sequence in this range, This is

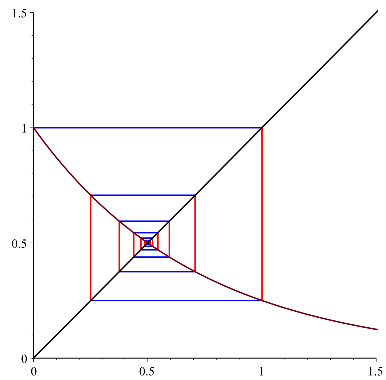

Figure 6: (Appendix B) A cobweb diagram for the case of $a < 1$.

very different from what we have seen before, the sequence is oscillating above and below the stationary value closing in as it does. In fact, it's from pictures like this that we get the name 'cobweb', the diagrams we have seen so far are sometimes called 'staircases'. In the diagram it looks as though the sequence has the property that,

$$x_0 < x_2 < x_4 ... < \alpha < ... < x_5 < x_3 < x_2 \tag{5}$$

If we can prove that these inequalities hold, then we will have shown that the sequence x_n does converge.

We will begin by looking at (2) which quite quickly gives us the very useful subsequence,

$$x_0 < x_2 < x_1 \tag{6}$$

This is important because we can use it to show that repeated application of the E(x) on it gives us the full sequence inequality. It is important to see that (5) is actually composed of three rules which define the whole; the even subsequence, the odd subsequence and their relationship:

$$\begin{aligned} x_{2n+2} &= E_2(x_{2n}) \\ x_{2n+3} &= E_2(x_{2n+1}) \\ x_{2n} & \quad \text{compared to} \quad x_{2n+1} \end{aligned}$$

Where $E_2(x)$ means the composite function $E(E(x))$. The conjectures we wish to demonstrate are these:

$$\begin{aligned} x_{2n} &< x_{2n+2}, \\ x_{2n+1} &> x_{2n+3} \text{ and} \\ x_{2n} &< x_{2n+1}. \end{aligned}$$

We will prove each of these statements in turn.

Theorem 2.1. *: The Even Subsequence*
If $E(x) = a^x$ and x_n is an iterative sequence defined such that $x_{n+1} = E(x_n)$, then for $0 < a < 1$, the even subsequence has the property that: $x_{2k} < x_{2k+2}$, while the sequence is not in a cycle.

Proof. We begin by assuming that the theorem is true for some $n = k$,

$$x_{2k} < x_{2k+2}$$

We will prove that if this is the case, then the statement it is true for $n = k + 1$;

$$x_{2k+2} < x_{2k+4}$$

To begin we write the two sides of the above inequality in terms of the case $n = k$ by using the definition of the sequence:

$$\begin{aligned} x_{2k+2} &= E_2(x_{2k}) = a^{a^{x_{2k}}} \text{ and,} \\ x_{2k+4} &= E_2(x_{2k+2}) = a^{a^{x_{2k+2}}} \end{aligned}$$

Since the sequence is not in a cycle either it is true that $x_{2k+2} < x_{2k+4}$ or it is true that $x_{2k+2} > x_{2k+4}$. If we assume that second is true we get,

$$\begin{aligned} & a^{a^{x_{2k}}} && > a^{a^{x_{2k+2}}} \\ \Leftrightarrow\ & a^{x_{2k}} \ln(a) && > a^{x_{2k+2}} \ln(a) \\ \Leftrightarrow\ & a^{x_{2k}} && < a^{x_{2k+2}} \\ \Leftrightarrow\ & x_{2k} && > x_{2k+2} \end{aligned}$$

which is a contradicts the statement of the case $n = k$. Thus if there exists such case, then it must be true that $x_{2k+2} < x_{2k+4}$. Hence, we have shown that if the conjecture is true for some $n = k$, then it must be true for $n = k + 1$. We know from (6) that it is true for $n = 0$, so it must be true for n=0, 1, 2, 3... QED. □

Theorem 2.2. *The Odd Subsequence*
If $E(x) = a^x$ and x_n is an iterative sequence defined as $x_{n+1} = E(x_n)$, then for $0 < a < 1$ the odd subsequence has the property that: $x_{2k+3} < x_{2k+1}$, if the sequence is not in a cycle.

Proof. We assume that the sequence is true for some $n = k$,

$$x_{2k+1} > x_{2k+3}$$

We will prove that if this is true, then the theorem is true for $n = k + 1$;

$$11x_{2k+3} > x_{2k+5}$$

Since there isn't a cycle either it is true that $x_{2k+3} > x_{2k+5}$ or it is true that $x_{2k+3} < x_{2k+5}$. We will assume the latter and, recalling the fact that $a < 1$, by using the definition of the sequence we get;

$$\begin{aligned} & x_{2k+3} && < x_{2k+5} \\ \Leftrightarrow\ & a^{a^{x_{2k+1}}} && < a^{a^{x_{2k+3}}} \\ \Leftrightarrow\ & a^x_{2k+1} && > a^x_{2k+3} \\ \Leftrightarrow\ & x_{2k+1} && > x_{2k+3} \end{aligned}$$

This is a contradiction of the case $n = k$. Thus, the opposite must be true such that $x_{2k+3} > x_{2k+5}$. Hence, we have shown that if the theorem is true for some $n = k$, then it must be true for $k + 1$. Since we know that it is true for $n = 0$, we know it is true for $n = 0,\ 1,\ 2,\ 3...$ QED. □

Theorem 2.3. *The Relationship between Odd and Even terms*
If $E(x) = a^x$ and x_n is an iterative sequence defined such that $x_{n+1} = E(x_n)$, then every odd terms is greater than every even term: $x_{2k} < x_{2k+1}$. If the sequence does not enter a cycle.

Proof. We assume the statement for some $n = k$:

$$x_{2k} < x_{2k+1}$$

We will prove that if this is true, then the theorem is true for $n = k + 1$,

$$x_{2k+2} < x_{2k+3}$$

Either the above statement is true, it its opposite is true. If we assume the latter then we get that,

$$\begin{aligned} & x_{2k+2} && < x_{2k+3} \\ \Leftrightarrow\ & a^{a^{x_{2k}}} && < a^{a^{x_{2k+1}}} \\ \Leftrightarrow\ & a^{x}_{2k} && > a^{x}_{2k+1} \\ \Leftrightarrow\ & x_{2k} && > x_{2k+1} \end{aligned}$$

This contradicts the statement for some $n = k$, therefore the opposite must be true. Hence, we have shown that if it is true for $n = k$, it is true for $n = k + 1$. We know it is true for $n = 0$ because of (6) and so we know that it is true for $n = 0,\ 1,\ 2...$ QED. □

One may notice that all three proofs follow the exact same form and in fact the first two cross through each other. For example at one point in the first theorem we compare odd terms and in one case of the second theorem we compare even terms. In reality these are the same proofs shifted by one iteration, they are presented here separately for clarity though.

We have now succeeded in proving all three features which define (5) and therefore we have proved the statement itself: the sequence x_n converges for $0 < a < 1$. There is only one problem with our proof which is that we assumed that there were no cycles. However, from figure 2 we can see that one forms for very small a so we must now deduce what he limit is.

2.2 Exponential Sequence Cycles

There are two possible cases, either the sequence converges to a cycle or to a value. To continue our investigation, therefore, we must understand what at is it about a stationary value which means that a sequence will converge to it rather than to a cycle around it.

We shall introduce another general theorem here which will help us solve this problem.

General Theorem 2. *Attracting Stationary Values*
If an iterative sequence x_n is defined recursively as $x_{n+1} = f(x_n)$, for some differentiable function f, and the sequence is convergent such that $x_n \to \alpha$ as $n \to \infty$. Then, the size of the derivative at α is $|f'(\alpha)| \leq 1$.

Proof. From the definition of the derivative at some point we have that

$$f'(\alpha) = \lim_{x \to \alpha} \frac{f(x) - f(\alpha)}{x - \alpha}$$

Now since we have the sequence x_n and its limit is the stationary value $f(\alpha) = \alpha$ we can write

$$f'(\alpha) = \lim_{x \to \alpha} \frac{x_{n+1} - \alpha}{x_n - \alpha}$$

If we assume that $|f'(\alpha)| > 1$ we get that

$$|x_{n+1} - \alpha| > |x_n - \alpha|$$

which contradicts the fact that $x_n \to \alpha$. Hence we have that $f'(\alpha) \leq 1$. QED. □

This theorem is really fascinating because we can create similar theorems to show that for any recursive sequence based on a differentiable function f with a stationary value at α we have the following statements:

$$\alpha \text{ is attracting if } |f'(\alpha)| < 1$$
$$\alpha \text{ is indifferent if } |f'(\alpha)| = 1$$
$$\alpha \text{ is repelling if } |f'(\alpha)| > 1$$

These statement are quite revealing about sequences, they tell us that if we know a sequence converges to a value then we know which value it converges to. For example, suppose that we had a convergent sequence x_n based on a function f with two stationary values, α and β, and we could compute the derivative at these values. Suppose now that $f'(\alpha) < 1$, but $f'(\beta) > 1$, then we know that it must converge to α by the above General Theorem. This is so useful because f might be incredibly difficult to manipulate.

In fact we already used these theorems: we saw that if $|M'(\alpha)| < 1$, then the iterative sequence defined on M will converge to α. Furthermore, we saw that if the sequence x_n based on M converges and $|M'(\alpha)| = 1$, then the sequence is indifferent as we cannot tell if it converges to a cycle or a value without further analysis. We also saw that if $|M'(\alpha)| < 1$, then $|M'(\beta)| > 1$ which means that β is a repelling stationary value and indeed we concluded that the sequence never converges to β when we did the analysis in the last section.

We know that in the range $0 < a < 1$ the sequence converges, but let's suppose that for some a we have that $|E'(\alpha)| \geq 1$ and $|E'(\beta)| \geq 1$ such that the stationary value is repelling or indifferent. Meanwhile $|E_2'(\alpha)| < 1$ and $|E_2'(\beta)| < 1$ so that the even and odd subsequences. The values α and β are stationary values o$f E_2(X)$.

In reality this argument works for a p-cycle too and is in fact a corollary of the above theorem.

Corollary 2.1. *Let an iterative sequence be based on a function f with the property that $|f_p'(\alpha)| < 1$ and $|f_k'(\alpha)| > 1$ for all $0 < k < p$. Then the p-subsequences each converge to a different values, such that the sequence forms a p-cycle.*

So using this we may propose that there exists an $a_{min} = A$ such that the following conditions are satisfied;

$$\begin{aligned} x_n &\to \alpha \text{ as } n \to \infty \text{ for } A \leq a < 1 \\ x_n &\to \text{2-cycle as } n \to \infty \text{ for } 0 < a < A \end{aligned}$$

To deduce where the above is true we need to ask

$$\begin{aligned} \text{for what } a = A \text{ is} \quad & |E_2'(\alpha)| > 1 \quad \text{and,} \\ \text{if} \quad & |E_2'(\alpha)| \leq 1 \quad \text{does } x_n \to \alpha? \end{aligned}$$

The first question can be answered by looking at where the derivative of the second iteration of $E(x)$ at α satisfies the condition set. However, it is easier for us to ask when is the opposite true? That is to say, we will look at when the sequence converges to a value rather than a 2-cycle. If the sequence converges to a value, then the 2-subsequence must also do so such that by the above General Theorem

$$\begin{aligned} |E_2'(\alpha)| &= \ln(\alpha)^2 \leq 1 \\ &\Leftrightarrow \frac{1}{e} \leq \alpha < 1 \\ &\Leftrightarrow (\frac{1}{e})^e \leq a < 1 \end{aligned}$$

The last step is true because section 2.1 told us that the function $x^{\frac{1}{x}}$ is strictly increasing in the range $0 < x < 1$ for that range of a. Hence, from this we can see that $A = (\frac{1}{e})^e$ as this is the minimum bound for which the condition on the derivative is true. We may notice that the remark in 2 suggesting that the cycle starts at $a \approx 0.1$ is supported here as $A \approx 0.065988...$ which is very close.

We must now answer the second question which asks; can we prove that for the range $A < a < 1$, the sequence x_n converges to a value? This may seem obvious, and it may seem like we already proved it in the last section, however we need to be careful not to confuse convergence with convergence to a value. We do not in fact know whether it converges to a 2-cycle or a value in this range.

To prove this we will show that there is only one stationary value in each of the even and odd subsequences and these values are one and the same: this means that there cannot be a cycle and since the sequence converges, it must therefore converge to a number. Therefore, we must simply ask how many times the line $y = x$ and the line $y = E_2(x)$ intersect for the range $(\frac{1}{e})^e \leq a < 1$ as these are the stationary values. We begin by looking at the inflexion point(s) t of E_2, which we deduce from the definition has the property that

$$\begin{aligned}
E_2''(x) &= a^{a^x}(a^x)^2\ln(a)^4 + a^{a^x}a^x\ln(a)^3 \\
\Leftrightarrow \quad & a^{a^t}(a^t)^2\ln(a)^4 + a^{a^t}a^t\ln(a)^3 = 0 \\
\Leftrightarrow \quad & t = \frac{\ln(\frac{-1}{\ln(a)})}{\ln(a)} \\
\Leftrightarrow \quad & a^t = \frac{-1}{\ln(a)} \\
\Leftrightarrow \quad & a^{a^t} = \frac{1}{e}
\end{aligned}$$

From this we understand that if $(\frac{1}{e})^e \leq a < 1$, then t must be on the left of, or equal to, the stationary value as $E_2(x)$ is strictly increasing. Indeed, we may see that if $(\frac{1}{e})^e \leq a < 1 \Leftrightarrow \frac{1}{e})^e \leq \alpha < 1$, then we get the result that

$$a^{a^t} = \frac{1}{e} \leq \alpha = a^{a^\alpha}$$

The reason this is so useful is that since there are no points of inflection on the right of α the function E_2(x) must be strictly decreasing for all $x > \alpha$ and since we computed earlier that $E'(\alpha) \leq 1$, we see that the lines $y = E(x)_2$ and $y = x$ are getting further away from each other as the gradient of $y = x$ is one, hence they never meet again. So we may conclude that there is only one stationary value for $E_2(x)$ which tells us that in the limit the even and odd subsequence must converge to the same value. We have thus succeeded in answering both the question we set out to do.

2.3 Conclusion on the Exponential Sequence

We have now completed the analysis on the entire valid range of a for the sequence based on E(x). To sum up our results we have shown that as $n \to \infty$,

$$\begin{aligned}
&\text{for} && a \leq 0, && x_n \notin \mathbb{R}; \\
&\text{for} && 0 < a \leq (\frac{1}{e})^e, && x_n \to \text{2-cycle} \\
&\text{for} && (\frac{1}{e})^e < a \leq e^{\frac{1}{e}}, && x_n \to \alpha \text{ or } \beta; \\
&\text{and for} && e^{\frac{1}{e}} < a, && x_n \to \infty
\end{aligned}$$

Which means that for all $a \in \mathbb{R}^+$ we understand what the sequence is doing. We have seen here that even though a function may appear to be complicated for solving equations, modelling, calculus etc. when it comes to sequences it can be very well behaved and understandable despite expectations.

However, the analysis presented is in fact incomplete. In the last part of 2.2 it was only proved that x_n does not converge to 2-cycle for $(\frac{1}{e})^e \leq a < 1$. We have not shown that it does not form a p-cycle for odd $p > 2$. This last statement we believe can be proved in a similar way as for the 2-cycle because the first and second derivatives of the iterated function follow a very clear pattern:

$$
\begin{aligned}
E'_p(x) &= \ln(a)^p \underbrace{a^x a^{a^x} \dots}_{\text{until height p}} \\
&= \ln(a)^p \prod_{i=1}^{p} E_i(x) \\
E''_p(x) &= E'_p(x) \sum_{i=1}^{p} \frac{E'_i(x)}{E_i(x)}
\end{aligned}
$$

The task of deducing that there are no p-cycles in the range is left up to the reader using the above or otherwise. Indeed, it looks as though it may be *easier* to prove it without the above, though the formulae are very interesting and aesthetic. General Theorem 3 which we meet later on may also be useful.

On a similar note, it would indeed be interesting to complete an full analysis of $x_0, a \in \mathbb{C}$. That is to say we have only looked at a very restricted range of these values when in fact we may ask what happens when $x_0 = 1$ and $a = i$ where $i = \sqrt{(-1)}$. Though this would certainly be a very long and technical investigation, it would be very wonderful to do.

3 Quadratic Sequence

Suppose that there was a pond in which there were some fish. We would expect the size of the population to have a maximum; after which the food would begin to run out as the demand was too high. This would mean that the population would be reduced as there are not enough resources for them to survive.

We could model such a situation with a graph of the percentage of maximum population on the vertical axis and the ratio of food to fish on the horizontal,

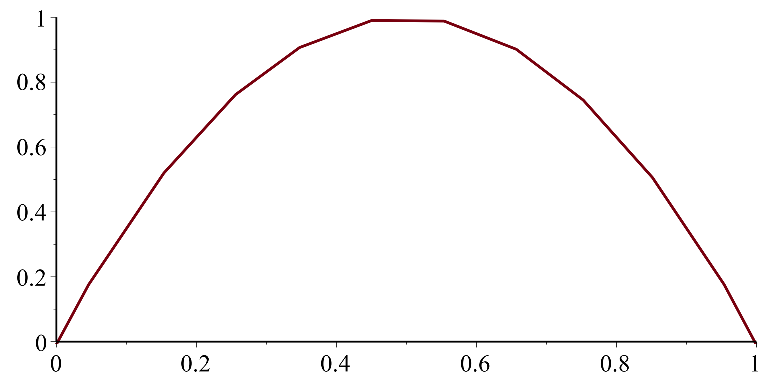

Figure 7: A model for population rise and decline.

Suppose in month zero there was a food to population ratio x_0, we then apply the model to see what the ratio x_1 will be after one month. The next question we may ask is what happens as we apply the model again and again such that we find the food to population ratio x_2 after two months of monitoring the pond and then the ratios $x_3, x_4, x_5...$ Let x_n be the percentage of the maximum possible population after n months, we may compute this figure by the recursion relation,

$$\begin{aligned} Q(x) &= \lambda x(1-x) \\ x_{n+1} &= Q(x_n)\ for\ n = 1, 2, 3... \\ \Leftrightarrow x_{n+1} &= \lambda x_n(1-x_n) \end{aligned}$$

Where λ is some parameter to configure the model. The question we want to ask ourselves is what is the long term population, how many fish can we expect to eventually have? This model was in fact invented through a problem such as the one above and is sometimes called the *'logistic map'* and it was invented through a problem like the one presented.

Clearly this is a problem in iterative sequences, the reason it is so interesting is because it is actual conjugate to all quadratic sequences of the form

$$x_{n+1} = ax_n^2 + bx_n + c$$

The proof of this will not be reproduced here as an explanation of it was already presented in section 1.4. However, reader may want to do the demonstration for themselves with the substitutions,

$$\begin{aligned} x &= -\left(\frac{2a\alpha + b}{a}\right)u + \alpha \\ \Rightarrow \lambda &= 2a\alpha + b \end{aligned}$$

Therefore if this map is conjugate to the general quadratic map, then if we understand what happens if we iterate using the logistic map, we understand what happens when we iterate all quadratic functions.

We saw in section 2 that complicated function can be quite simple as iterative sequences, this time we shall see the converse as though second order polynomials are quite easy to manipulate, the sequence produced is a prime example of chaos. It so happens that changing the only parameter λ by a little bit can result in impossibly different outcomes of the sequence, and therefore of our fish population.

To demonstrate this more clearly observe the diagram below which has λ on the horizontal and x_∞ on the vertical starting with $x_0 = \frac{1}{2}$ as this is the maximum possible population.

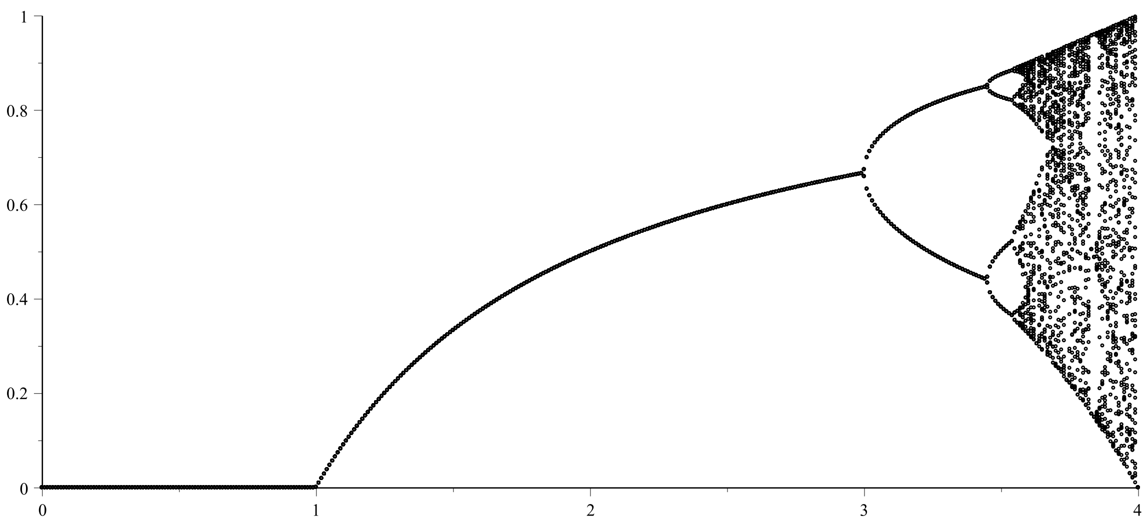

Figure 8: (Appendix C) A period doubling bifurcation diagram for the quadratic sequence.

This diagram has some very interesting features;

$$\begin{aligned} \text{the graph converges to 0} \quad &\text{for} \quad 0 \leq \lambda \leq 1 \text{ and then} \\ \text{to some non-0} \quad &\text{for} \quad 1 < \lambda \leq 3 \end{aligned}$$

After which, however, it does something very strange indeed called period-doubling bifurcation where it splits into;

$$\begin{aligned} \text{a 2-cycle for} \quad &3 < \lambda \leq 3.45; \\ \text{a 4-cycle for} \quad &3.45 < \lambda \leq 3.55. \end{aligned}$$

As λ grows there is chaos, the period of cycles doubling and coinciding giving unpredictable even and odd cycles.

As if this was not bizarre enough for such a small range of λ not even reaching $\lambda = 4$, it 'crystallises' in a strange way such that amid the long-cycles we get a beautiful 3-cycle at about $\lambda \approx 3.8$. We will need to investigate what this mysterious figure is as it is such a wonderful point of order. It would appear that though this sequence should be well behaved, it is in fact very chaotic and hard to understand.

Furthermore, it should be noted that where we stated that the vertical axis represented x_∞ this was an exaggeration. In fact as can be seen in the appendix, the vertical values are found after a few hundred iterations of the function and the last 50 of these are plotted. We must not rely too much on the diagram therefore as it can do no more than give us a good idea of what is happening as we iterate.

In this analysis we will not concern ourselves so much with 'why' it has this strange behaviour, but rather with 'when' it happens. That is to say we want to be able to compute, and prove, where we get the convergence to numbers and p-cycles culminating in the remarkable 3-cycle. Lastly, not that as we cannot have negative or complex fish or food in the pond, we shall take $\lambda, x \in \mathbb{R}^+$ in the whole investigation.

3.1 Convergence in the range $0 < \lambda \leq 3$

We have the theorem which says that if a function converges, then it must converge to its stationary values. Let's therefore begin by deducing what the stationary values are for particular values of λ, it is obvious that the origin is always a stationary value, the other one may be computed as

$$\begin{aligned} \alpha &= Q(\alpha) \\ \Leftrightarrow \alpha &= \lambda\alpha(1-\alpha) \\ \Leftrightarrow \alpha &= 1 - \frac{1}{\lambda} \end{aligned} \tag{7}$$

The formula for α is very interesting because it tells us that there are no valid values for α in the range $0 < \lambda \leq 1$ as the formula gives only negative values or 0 in this range, so the only stationary value is at the origin. Now let's look at the derivative,

$$Q'(x) = \lambda(1-2x) \tag{8}$$

It is easy to see that for all $0 < x < \frac{1}{2}$ when we restrict the range of λ in this way we may apply General Theorem 1 to see that the sequence converges to the origin as we expected from the diagram. We end up with no fish. We may apply the theorem because function is increasing and the line $y = x$ is above the lines $y = Q(x)$ as there are not stationary values other than $(0, 0)$. However, we must now investigate what happens for the range of $\frac{1}{2} \leq x_0 \leq 1$ to get a complete picture of the function.

When we look at 7 we see that for all x in this range, the function return a smaller value. Indeed by a similar argument to the above we see that this is true as $y = x$ is above the function and there are no stationary value other than the origin. Eventually, since x_n is getting smaller, the sequence will be in the range we have already proved with the General Theorem.

We may interpret this therefore as saying that when the population has a maximum of less than a quarter the threshold of the pond, as $\lambda \leq 1 \Rightarrow Q_{max} \leq \frac{1}{4}$ by completing the square, then the population cannot sustain itself and dies out.

There rest of this analysis is similar, indeed it is just an application of the general theorems in the specific case of the logistic map except the case of $\lambda = 3$ which is quite unique. As can be seen in the diagrams at the end of this section there are several different cases which may arise, some we have already taken care of, and the subtle difference means individual analysis is necessary.

Firstly, we will restrict ourselves to the range $1 < \lambda < 2$, we now know from (7) that this means there exists a stationary value α:

1) Suppose $0 < x < \alpha$, we see that this means that the line $y = x$ is below the line $y = Q(x)$ and if we take $\alpha \leq \frac{1}{2}$ then the function is strictly increasing again so the General Theorem tells us the sequence converges α as this s the upper bound.

2) Suppose instead that we were to look at the case $\alpha \leq x \leq \frac{1}{2}$, then we see that $y = x$ is above the line $y = Q(x)$ meaning that we once more can apply General Theorem 1 to see we converge to the lower bound α as the function is stall strictly increasing.

Secondly, we will restrict ourselves to the range $2 < \lambda \leq k$ where we understand that α is on the right hand side of the function. We have left in k as we do not know where the bifurcation starts. Let us first deduce the range of α using equation (7) which is strictly decreasing as it is a hyperbola, so that we get

$$\frac{1}{2} < \alpha \leq 1 - \frac{1}{k}$$

Now we ask what the last value for which the stationary value is attracting is, as this is the last time we know the function definitely converges to a value. Such that by using General Theorem 2 we may deduce the following,

$$\begin{aligned} |Q'(\alpha_{max})| &< 1 \\ \Leftrightarrow |k(-1 + 2/k)| &< 1 \\ \Leftrightarrow |2 - k| &< 1 \\ \Leftrightarrow 2 < k &< 3 \end{aligned}$$

This is a very interesting result, we have shown after this short calculation that for $2 < \lambda < 3$ the sequence converges to a value. The population of our fish stabilises.

This is in fact not surprising though, as we saw in the period-doubling diagram the sequence now splits into a 2-cycle. We must, however, not be fooled into thinking that we have found the whole range where the sequence is convergent to value as when the stationary value is indifferent it could still be the limit of the sequence. It is not clear what happens when $\lambda = 3$, since α is indifferent, so some extra argument is needed here. Roughly, we will simply have to prove that all p-cycles have a single stationary value for this parameter, and therefore all subsequences converge to the same limit. What makes the following argument so difficult is that when $\lambda = 3$ the sequence bifurcates for the first time, and so the analysis becomes much harder.

We will begin by looking at the case of potential 2-cycles as we may generalise these results. To do this we want to begin by familiarising ourselves with some of the features of the second iteration of the map, the very first statement we wish to see is that $Q_2(x)$

is symmetric about the line $x = \frac{1}{2}$, we may deduce this as

$$\begin{aligned}
Q_2\left(\frac{1}{2} - x\right) &= \lambda(\frac{1}{2} - x)(1 - (\frac{1}{2} + x)) \\
&= \lambda(\frac{1}{2} - x)(\frac{1}{2} + x) \\
&= \lambda(\frac{1}{2} + x)(1 - (\frac{1}{2} - x)) \\
&= Q_2\left(\frac{1}{2} + x\right)
\end{aligned}$$

This will half our investigation as everything we know about one side of the function we may apply to the other. Now we wish to find the maximum values s of this function, after a short calculation we achieve the following,

$$\begin{aligned}
Q_2'(x) = -4\lambda^3x^3 + 6\lambda^3x^2 - 2\lambda^3x - 2\lambda x + \lambda^2 &= 0 \\
\Rightarrow -\lambda^2(-1 + 2x)(2\lambda x^2 - 2\lambda x + 1) &= 0 \\
\Leftrightarrow s_1 &= \frac{l + \sqrt{\lambda^2 - 2\lambda}}{2\lambda} \\
\Leftrightarrow s_2 &= 1 - s_1 \\
\Leftrightarrow Q_2'(s_1) &= \frac{\lambda}{4} \\
\Leftrightarrow Q(s_1) &= \frac{1}{2}
\end{aligned}$$

Using the above statements we may see that for all λ in the range,

$$\begin{aligned}
\alpha < \quad Q(\tfrac{1}{2}) &< s_1 \\
Q_2(\tfrac{1}{2}) &> \frac{1}{2}
\end{aligned} \tag{9}$$

The last thing we wish to observe is that there are no other stationary values than α in this range of λ, we may simply deduce this from the definition,

$$Q_2(x) = x$$

Though this is a quantic, we can solve it without too much difficulty as we already know that $x = 0, \alpha$ are roots of it. This tells us that the function is below the line $y = x$ for the entire range of λ up to $x = \alpha$.

We will now use a wonderfully interesting trick to effectively solve for 2-cycles, from this we can confirm that $\lambda = 3$ is not a 2-cycle and we may generalise this to see that it does not form a p-cycle. The principle of this trick is hat a 2-cycle has the property that for some value $Q_2(\alpha) = \alpha$ and $Q(\alpha) \neq \alpha$, to eliminate this second case we simple send it to infinity,

$$\begin{aligned}
\frac{Q_2(x) - x}{Q(x) - x} &= q_2(x) \\
&= \lambda^2x^2 - (\lambda^2 + \lambda)x + (\lambda + 1)
\end{aligned} \tag{10}$$

This new function has a brilliant property which is that whatever stationary value it has is a 2-cycle, it cannot be anything else. We can therefore quite easily see that in the special case of $\lambda = 3$ there are no real solutions and so this is definitely not a 2-cycle. It is now clear that,

$$x < Q_2(x) < \alpha, \qquad \text{for } \frac{1}{2} \leq x < \alpha,$$
$$x > Q_2(x) > \alpha, \qquad \text{for } \alpha x \leq s_1,$$

It follows that the sequence $x_{2n} = Q_{2n}(x_0) \to \alpha$ for all $\frac{1}{2} \leq x_0 \leq s_1$, and if we take $x_0 = \frac{1}{2}$ as does $x_1, x_2 x_3 ...$ because from (9) we get $x_1 = Q(\frac{1}{2})$.

The next task is to generalise this statement and show that $x_n = Q_n(x_0) \to \alpha$ for all x_0 in $(0, 1)$. The bounds have been excluded because they are trivial. This task relies on one feature, which is to show that for some n depending on x the function $Q_n(x)$ lies in $[\frac{1}{2}, s_1]$ as we know it must converge once it is in this range by the above argument.

Firstly, let $x < \frac{1}{2}$ then by General Theorem 1 we know $Q(x) > x$ such that when iterating $x \to \frac{1}{2}$ which is in the range we know the sequence converse for, hence it converges for all $0 \leq x_0 \leq \frac{1}{2}$.

Secondly, suppose now that $\frac{1}{2} < x < 1$, it is not hard to see that in this case effectively the same thing happens. This time the function $Q(x)$ will first map to the case where $x < \frac{1}{2}$ and then the same process applies so that after n iterations we are in the case of the previous argument where w know the sequence converges.

Hence, we have seen that when $\lambda = 3$ the sequence conversed for the whole of the valid range of x and so we have proved that the sequence x_n is convergent for the range $0 \leq \lambda \leq 3$ and is non-zero for $2 \leq \lambda \leq 3$. The fish are happy.

On the following page are all the cases we have looked at to demonstrate the convergence of the sequence; from left to right are the cases of $x_0 < \frac{1}{2}$, $x_0 = \frac{1}{2}$ and $x_0 > \frac{1}{2}$. (The script used to generate all of the images may be found in Appendix D).

Convergence to a value of the sequence $x_{n+1} = Q_{(}x_{n})$ for various starting points:

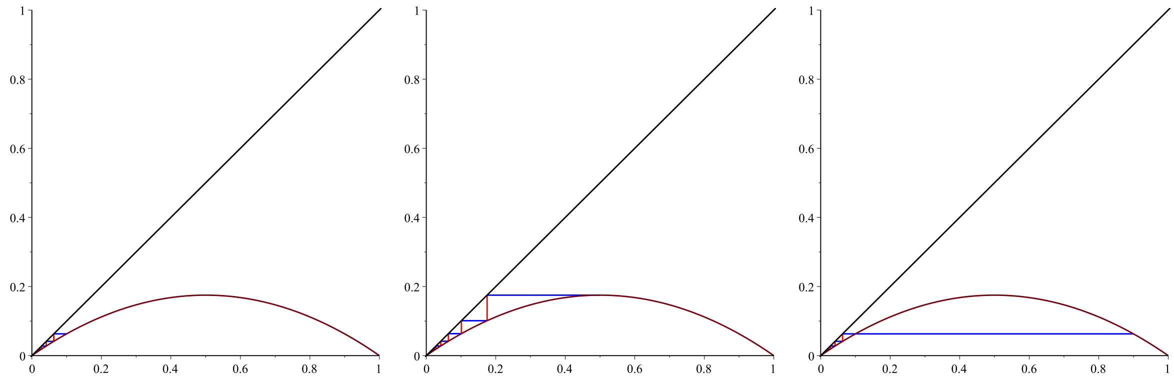

Figure 9: The case of $0 < \lambda < 1$.

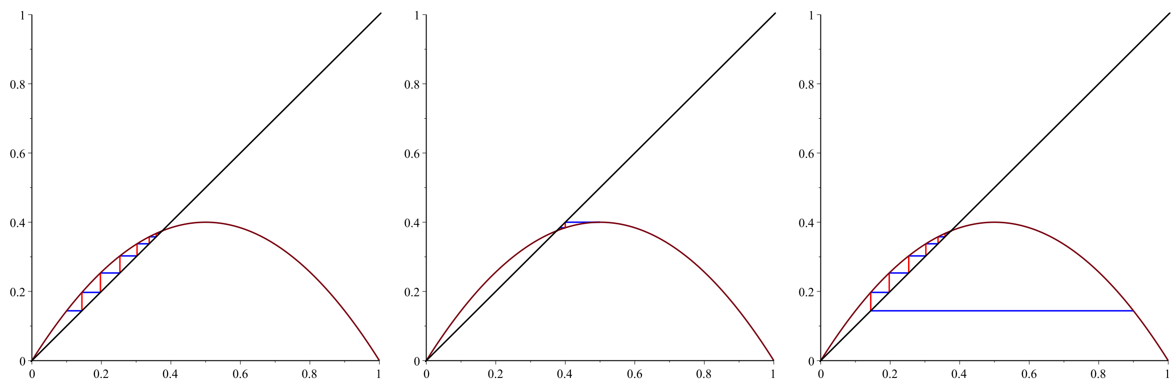

Figure 10: The case of $1 < \lambda < 2$.

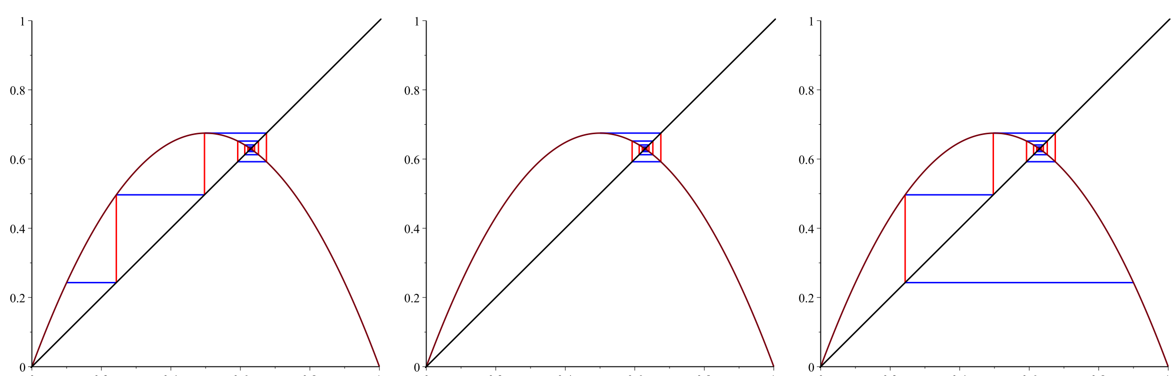

Figure 11: The case of $2 < \lambda \leq 3$.

3.2 Converging to 2-cycles

In the last section we already discovered many of the conditions and properties of the second iteration of the logistic map, and from the bifurcation diagram we approximated that for λ in (3,3.45] the sequence x_n converges to a 2-cycle. We will now look more in detail at what is happening in this range and what the significance of the upper bound is.

The first thing to note is that the graph of $y = Q_2(x)$ being a quartic looks quite different than the normal parabola, but many of the things deduced in the last section can be seen on the diagram quite easily. So the analyst may be easier if we have a picture of what we are investigating:

There are now 4 fixed points to consider, we shall call them $0, \alpha, \beta$ and γ in increasing order, recalling that this α is not the same as in the last section. Indeed, we already know the value of two of them as they are the stationary values of the sequenced based on $Q_(x)$, we may remove these happily so that we are left with the two new roots α and γ. These values are the solutions to the equation $q_2(x) = 0$,

$$x_{\alpha,\gamma} = \frac{\lambda + 1 \pm \sqrt{(\lambda+1)(\lambda-3)}}{2\lambda} \tag{11}$$

Now since we know the values of α and γ we may observe that they form a 2-cycle. This can be checked through substitution or by the argument that since $Q(\alpha)(\neq 0, \beta)$ is a fixed point of $Q_2(x)$, then $Q_(\alpha) = \gamma$ and hence $Q(\gamma) = Q_2(\alpha) = \alpha$.

To investigate this 2-cycle we may propose that the nature of the stationary values is the same, that is to say attracting, repelling or indifferent. We may even propose that the value of their slopes is exactly identical, this is the theorem below.

General Theorem 3. *Multiplier of a Periodic Point*
Let x_n be the n^th term of a sequence based on a function f such that $x_{n+1} = f(x_n)$ with some arbitrary starting value. Then the derivative of f_p takes the same value at every point in the cycle. That is to say if $\alpha, f(\alpha), ..., f_{p-1}(\alpha)$ form a p-cycle, then the slope at each one of these points is identical.

Proof. By the chain rule we may differentiate $f_p(x)$,

$$\begin{aligned}
f_p'(x) &= (f(f_{p-1}(x)))' \\
&= f'(f_{p-1}(x))f_{p-1}'(x) \\
&= f'(f_{p-1}(x))(f(f_{p-2}(x)))' \\
&= f'(f_{p-1}(x))f'(f_{p-2}(x))f_{p-2}'(x) \\
&\vdots \\
&= f'(f_{p-1}(x))...f'(f(x))f'(x)
\end{aligned}$$

First substitute $x = \alpha$,

$$f_p'(\alpha) = f'(f_{p-1}(\alpha))...f'(f(\alpha))f'(\alpha)$$

then substitute $x = f(\alpha)$,

$$\begin{aligned}
f_p'(f(\alpha)) &= f'(f_{p-1}(f(\alpha))...f'(f(f((\alpha)))f'(f(\alpha)) \\
&= f'(f_p(\alpha))...f'(f_2(\alpha))f'(f(\alpha))
\end{aligned}$$

Since by definition $f_p(\alpha) = \alpha$, we deduce from the last two equations that $f_p'(f(\alpha)) = f_p'(\alpha)$. We then use this to see that,

$$f_p'(\alpha) = f_p'(f(\alpha)) = f_p'(f_2(\alpha)) = ... = f_p'(f_{p-1}(\alpha))$$

Which says that the derivative at every point in the cycle is identical. QED. □

The title comes from the convention that if α is a periodic point of f with period p, then $f_p'(\alpha)$ is called the *multiplier* of α. Thus we may restate the theorem as saying that the multiplier of all values in a p-cycle is the same. We may use this in conjunction with General Theorem 2 to see that we can classifies periodic values as

$$\begin{array}{rl} \text{attracting if} & |f_p(\alpha)| < 1 \\ \text{indifferent if} & |f_p(\alpha)| = 1 \\ \text{repelling if} & |f_p(\alpha)| > 1 \end{array}$$

Now we may wish to know what the multiplier of the stationary values in our 2-cycle is so we may determine if its attracting or not. By the above general theorem it is sufficient to compute the value at one of the values in the cycle,

$$\begin{aligned} f_2'(\alpha) &= f'(f(\alpha))f'(\alpha) \\ &= f'(\gamma)f'(\alpha) \\ &= \lambda^2(1-2\gamma)(1-2\alpha) \\ &= \lambda^2(1-2(\alpha+\gamma)+4\alpha\gamma) \\ &= \lambda^2\left[1-2\left(\frac{\lambda+1}{\lambda}\right)+4\left(\frac{\lambda+1}{\lambda}\right)\right] \\ &= 4+2\lambda-\lambda^2 \end{aligned}$$

As expected the awkward case of $\lambda = 3$ is also indifferent in this case as $f_(\alpha) = 1$. However, as λ increases beyond 3 the multiplier decreases so that we enter an attracting 2-cycle for $3 < \lambda < 1+\sqrt{6} \approx 3.44948974$, the upper bound being recognisable as our earlier estimate when looking at the bifurcation diagram. At this upper bound the sequence bifurcates again and we expect it to be similar to the first period-doubling event at $\lambda = 3$ and ultimately that the sequence x_n converges to a 2-cycle for $2 < \lambda \leq 1+\sqrt{6}$. For a much more general investigation than wham will be presentable see [3].

It is important to stress that being attracting is not the same as being the limit, the proof of General Theorem 2 assumes that the sequence converges when in fact we do not know it does. So f0r the range of $\lambda \in (3, 1+\sqrt{6})$ we must simply prove that the sequence is convergent.

General Theorem 4. *Convergence to an Attracting stationary value*
There exists a costive number δ such that if $|x_0 - \alpha| < \delta$, then sequence $x_{n+1} = f(x_n) \to \alpha$ as $n \to \infty$. Where x_0 is the intimal value of the sequence, α is an attracting stationary value and f is some real differentiable function.

Proof. Chose a number r such that $|f'(\alpha)| < r < 1$, then since

$$f'(\alpha) = \lim_{x\to\alpha}\frac{f(x)-\alpha}{x-\alpha}$$

there is a positive number δ such that for $0 < |x - \alpha| < \delta$,

$$\begin{aligned} \left|\frac{f(x) - f(\alpha)}{x - \alpha}\right| &\leq r \\ \Rightarrow |f(x) - f(\alpha)| &\leq r|x - \alpha| \\ \Rightarrow |f(x) - \alpha| &\leq r|x - \alpha| \end{aligned}$$

since $f(\alpha) = \alpha$. Now suppose that $|x_0 - \alpha| < \delta$, then

$$|x_1 - \alpha| = |f(x_0) - \alpha| \leq |x_0 - \alpha|$$

so that $|x_1 - \alpha| < \delta$. If we apply the general statement for r repeatedly we obtain the statement

$$|x_n - \alpha| \leq r^n |x_0 - \alpha|$$

From this it is easy to see that since $0 < r < 1$, the sequence $x_n \to \alpha$ as $n \to \infty$ because $r^n \to 0$. QED. □

Notice that the smaller $|f'(\alpha)$ is the smaller we may chose the value r to be, which means that the sequence converge much faster. Of course this means that the convergence is extremely fast when $f'(\alpha) = 0$ and therefore call such values *superattracting*.

From this theorem on values we may gain a similar one for p-cycles which are what we are in fact interested in, the below theorem does just that.

General Theorem 5. *Convergence to an Attracting p-cycle*
If α is an attracting periodic value of a function f, then there exists a positive number δ such that whenever $|x_0 - \alpha| < \delta$,

$$(f_p(x_0))_n = f_{pn}(x_0) \to \alpha \text{ as } n \to \infty$$

Using this theorem we need just to deduce the value of δ and then we have found our basin of attraction, that is the range of x_0 for which α is the limit, and we have demonstrated that the sequence is convergent. Certainly we know that there is a δ which satisfies the condition.

To compute δ we need to first maximise r from the proof as this will give us the largest possible range, we therefore take $r \to 1^-$ ¡ 1. Knowing this we need to find the largest $x - \alpha$ which satisfies the condition that for some α dependent on λ,

$$\left|\frac{Q_2(x) - Q_2(\alpha)}{x - \alpha}\right| < 1 \tag{12}$$

This is an optimization problem in multivariable calculus as we have two variables x and λ with the constraint being the above statement. However, to do the actual calculation for this is unrealistic as the functions are very large and complicated.

Instead we will approach it as a non-linear programming problem with the objective function $\delta = |x - \alpha|$ and the goal to maximise it. Plotting x_0 against λ we end up with the graph on the top of the next page. We will not prove analytically that the feasible region in the graph is true, that the region has no holes, but the reader is invited to do this for themself using the information found in section 3.1 about the function.

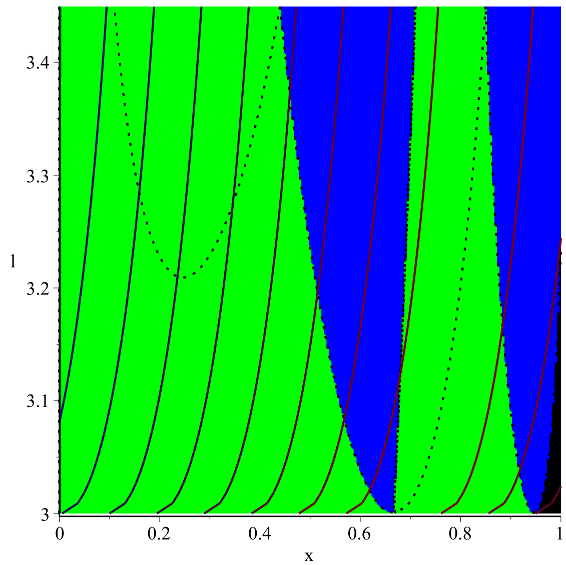

Figure 12: (Appendix E) Feasible region for constraint (12) with δ as the objective function.

Unlike most graphs of this type in this one there are three colours rather than the two that represent feasible and unfeasible regions. The colours are in fact basins, range of values, for which the statement has different behaviour; it is blue if it is true for α; green if it is true for γ; and black if it is unfeasible for both. The value δ is the largest value in the green or blue regions, however we need not enumerate it as we wish to prove that the whole valid plane is convergent.

To see why this unfeasible region arises we will look at the cobweb diagram of $y = Q_2(x)$ for some representative x and λ in the region,

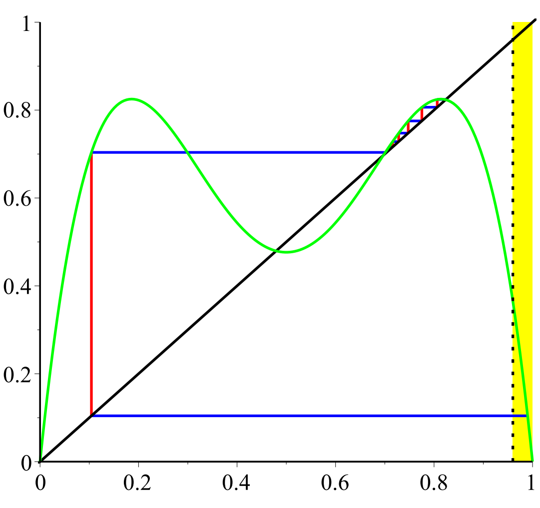

Figure 13: The Cobweb Diagram for a 2-cycle with unfeasible region of x_0 marked in yellow.

Therefore to complete the analysis of the attracting 2-cycle we must just show that the black basin maps to either of the other two regions after n iterations of the function, since we know by the above general theorem that values in this region will converge.

The cobweb diagram is exactly what we expected it to be, first x_0 maps to x_1 which is further away from either root but is in the feasible region. Indeed this graph is quite familiar; it exhibits the same behaviour as the last set of the cobweb diagrams for attracting stationary values and so we may use the same kind of analysis to demonstrate the map.

First we must find the range of the unfeasible region. Looking at the diagram it seems as though anything on the right of γ will have the map we are interested in therefore if we can show that the largest possible value of γ is not in the unfeasible region, then we take the region to be $\gamma_{max} < x < 1$ and we prove that $Q_2(x) \in (0, \gamma_{max})$.

If we differentiate 11 to find the local maxima we find that the stationary points are as follows,

$$\begin{aligned} x'_{\alpha,\gamma}(\lambda) = \frac{\lambda + 3 - \sqrt{\lambda^2 - 2\lambda - 3}}{2\lambda^2\sqrt{\lambda^2 - 2\lambda - 3}} &= 0 \\ \lambda + 3 - \sqrt{\lambda^2 - 2\lambda - 3} &= 0 \\ \lambda &= -\frac{3}{2} \end{aligned}$$

This is very useful as we now know that there is no local maxima in the valid range of λ, and so we must just check the bounds,

$$\begin{aligned} x_{\alpha,\gamma}(3) = \gamma = \frac{2}{3} &\approx 0.6666666666... \\ x_{\alpha,\gamma}(1+\sqrt{6}) = \gamma = \frac{\sqrt{2}\sqrt{3} + \sqrt{2} + 2}{2 + 2\sqrt{2}\sqrt{3}} &\approx 0.8499377795... \end{aligned}$$

Clearly the upper bound for λ gives the largest γ, and this value is outside the unfeasible region so we may set this is as out upper bound. What we must now show is that $Q_2 : (\gamma_{max}, 1) \to (0, \gamma_{max}]$. To do this we will consider the stationary values of $Q_2(x)$ which we saw earlier were s_1 and s_2 though we are only interested in the former as it is the greatest. We want to maximise the stationary values and so we once more treat it like and optimization problem,

$$s'_1 = \frac{1}{2\lambda\sqrt{\lambda^2 - 2\lambda}}$$

There another stationary values of s_1 and so we must just check the bounds of λ and see which one is larger,

$$\begin{aligned} s_1(3) = \frac{1}{2} + \frac{1}{6}\sqrt{6} &\approx 0.7886751347 \\ s_1(1+\sqrt{6}) = \frac{1 + \sqrt{6} + \sqrt{5}}{2 + 2\sqrt{6}} &\approx 0.8241157600 \end{aligned}$$

The upper bound gave the largest value of s_1, but this is on the left of γ_{max} and so we know the curve is strictly decreasing in the domain $(\gamma_{max}, 1)$. Further, since we know that the line $y = x$ is above the curve for the black basin, we know that it must map to smaller values. So that even if it maps to itself in the first iteration, it will after n iteration map to at most $gamma_{max}$ which is in the feasible region.

Hence we have shown that for all $x_0 \in (0,1)$ and $\lambda \in (3, 1+\sqrt{6})$ the function converges to an attracting 2-cycle.

3.3 Quadratic Sequence 3-Cycles

We will now look at the interesting point on the bifurcation diagram where everything settles and for a brief interval there is a 3-cycle. However, as the functions and number of values have grown we will not carry out such a detailed study into this area, but rather focus on computing the actual values at which it happens. Indeed once we have found the onset the 3 cycle we will not go any further.

Let's begin by eliminating the 1 cycles,

$$\begin{aligned} \frac{Q_3(x) - x}{Q(x) - x} &= q_3(x) \\ q_3(x) &= 0 \\ \lambda^6 x^6 + (-3\lambda^6 - \lambda^5)x^5 + (3\lambda^6 + 4\lambda^5 + \lambda^4)x^4 + & \\ (-\lambda^6 - 5\lambda^5 - 3\lambda^4 - \lambda^3)x^3 + (2\lambda^5 + 3\lambda^4 + 3\lambda^3 + \lambda^2)x^2 + & \\ (-\lambda^4 - 2\lambda^3 - 2\lambda^2 - \lambda)x + \lambda^2 + \lambda + 1 &= 0 \end{aligned}$$

This last sextic polynomial is very complicated and to try and solve it would be futile, however as we are only interested in where it returns real values it is sufficient to take its determinant. Before we do this let's first consider a similar idea about quartic polynomials.

Every quartic polynomial is conjugate to a quartic polynomial of the form $f(t) = t^4 + xt^2 + yt + z$ and each quartic of this form can be represented by a point (x, y, z) in parametric space. A polynomial will have a double root only if $f(t) = f'(t) = 0$. The set of points (x, y, z) which correspond to polynomials with a double root forms a boundary in parametric space between the points representing the quartics with zero, two, and four roots. We may now graph such a surface,

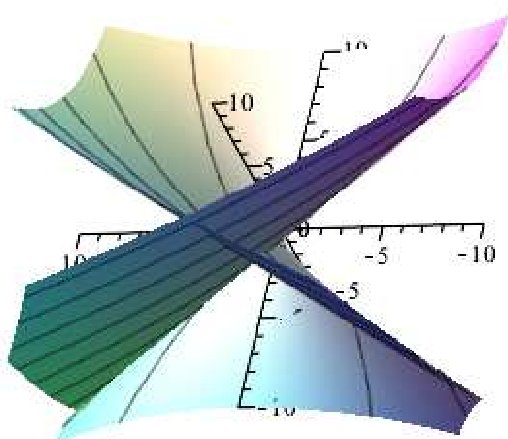

Figure 14: (Appendix F) Queue d'Aronde.

This boundary surface is known as the swallowtail. If we chose a point; above the surface we find that there are no roots to the polynomial; below and there are 2 roots; on it there 2 single roots and 1 double root; on the cusp there is 1 triple root and 1 single root; inside there are 4 roots; at self-intersection points there are 2 double roots; and where the curve intersects itself completely is a quadruple root. Notice that at the origin in parametric space the corresponding polynomial is $f(t) = t^4$, and its root is the quadruple root. This point is the only point corresponding to a polynomial with a quadruple root, because it is the only point for which $f(t) = f'(t) = f''(t) = f'''(t) = 0$ which is the necessary condition. (For a more detailed study on this surface see [2]).

This graph gives us a very visual understanding of how the polynomials behave, at least were we are interested in them, and we could theoretically do the same thing for sextic polynomials as all sextic polynomials are conjugate to $f(t) = t^6 + xt^4 + yt^3 + zt^4 + wt + v$. (In fact to remove the x^{n-1} term of a polynomial with degree n we must simply use the substitution $x = t + k$ where k is an appropriate constant constant.) However, we would have to graph it in hyperspace as we have (x, y, z, w, v) to consider. Therefore, to continue with the polynomial $q_3(x) = 0$ let us just take the discriminant,

$$\Delta = \frac{(\lambda^2 - 5\lambda + 7)^2(\lambda^2 - 2\lambda - 7)^3(1 + \lambda + \lambda^2)}{\lambda^{30}}$$

We see in this factorised form of the discriminant that if we want $\Delta \geq 0$, then we must just consider middle term in the numerator as it is the only one which can make the Δ negative. So the earliest point at which we get it to be non-negative is,

$$\begin{aligned} \lambda^2 - 2\lambda - 7 &= 0 \\ \lambda &= 1 + 2\sqrt{2} \approx 3.828427124 \end{aligned}$$

This is the onset of the 3-cycle and it does match the observation in our diagram from the last section. (The root $\lambda = 1 - 2\sqrt{2}$ is negative and so has been ignored). To prove that the sequence converges to the attracting 3-cycle is a longer, but not harder, analysis than for the 2-cycle and indeed the 4-cycle which is left up to the reader. The reason for this computation is the aesthetics involved: the point in the bifurcation diagram, the number $1 + 2\sqrt{2}$, the curious surface and the link of sequences to many other fields of mathematics.

3.4 Conclusion

We have so far seen that though quadratic functions, especially the logistic map, may seem to be very simple and well behaved are in fact very complicated. Indeed the entire investigation has been much more sophisticated than for the exponential function and required many more techniques and ideas to be completed. u We have seen that the population of fish in a pond can be a very strange thing, it may stabilise and get a fixed population or it may oscillate such that one month there are α fish and the other there are γ number of fish. However though this analogy is a good tool to introduce the idea of the logistic map, the function and the sequence in and of themselves are much more interesting. Indeed the relationships we have seen with bifurcations, fold bifurcations, and many other fields from this simple idea has been astounding. The idea of fish was in fact just a front, the truth is that there is a great amount of wealth to be gained when looking at the negative values of λ as a similar bifurcation process happens when we move to the left.

Indeed a much more complete analysis would answer the question 'why' which we have by no means addressed. To do this requires looking in the complex plane where a great deal of structure can be seen, though the analysis is much more difficult.

There have been many things left up to the reader to complete, this is a unfortunate consequence of timing constraints when writing the article. These tasks are by no means trivial, though they can be very similar to the ideas presented. Indeed to get a much better understanding of how the analysis of 2-cycles was completed it is advised to do one for the 4-cycle individually as this gives an insight into some of the more subtle ideas in section 3.2 and in general about the function.

There exists many more general theorems about convergent to 2p-cycles and such than was presented here and so with the wealth of analysis left to do on the logistic maple we may call this investigation a *'short introduction'* and hope it has interested the reader sufficiently to pursue the ideas further.

4 Newton-Raphson Method

The Newton-Raphson method is a classic and powerful way of solving equations of the form $f(x) = 0$ where f is some differentiable function. The principle lies in approximating the functions to a tangent and sliding x down the line until it reaches the x-axis where the approximate root lies. From this approximate root we can then repeat the method and the next value will be closer to the root than the former, an example of an iteration looks like,

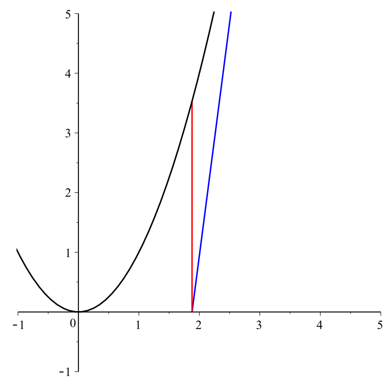

Figure 15: (Appendix G) The blue is the tangent, the red line is the new value.

The formula for the Newton-Raphson method of iteration is based on constructing the tangent, such that we have the linear approximation of a function at $(x_0, f(x_0))$ is

$$y - f(x_0) = f'(x_0)(x - x_0)$$

Since we wish solve $f(x) = 0$ it is obvious that we now take $y = 0$ and receive the value $x = x_1$ by rearranging the above,

$$x_1 = x_0 - \frac{f(x_0)}{f'(x_0)}$$

Using this new value x_1 we may now apply the method repeatedly such that we have the recursive formula,

$$x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$$

Naturally this formula is only valid when we are not a stationary points as this would make the tangent parallel to the x-axis. From the above definition and geometric interpolation of the formula it is not hard to see that we can have good and bad starting values, that is to say the sequence might converge to a root or diverge. Furthermore, if there are two roots and only the positive is wanted, then if the staring value is chosen badly the sequence might converge to the negative one.

Suppose that $x_n \to \alpha$ as $n \to \infty$, we must not be fooled into thinking that in this case $f(\alpha) = 0$ and this is a problem as we have not solved the equation. To demonstrate the deceitful kind of behaviour the sequence can have suppose that we have the function

$$f(x) = \begin{cases} 1 - 2x\sin(\frac{1}{x}) & \text{for } x \neq 0 \\ 1 & \text{for } x = 0 \end{cases}$$

This function is continuous everywhere, but only differentiable when $x \neq 0$, so we may use Newton-Raphson to approximate it. Take $x_0 = \frac{1}{2\pi}$, so that $f(x_0) = 1$, it is then not hard to see that we get the sequence,

$$\begin{aligned} x_0 &= \frac{1}{2\pi} \\ x_1 &= \frac{1}{4\pi} \\ &\vdots \\ x_n &= \frac{1}{2^{n-1}\pi} \end{aligned}$$

Thus it is definitely true that the sequence converges, that is to say we have as $n \to \infty$ the sequence $x_n \to 0$, however we know that $f(0) \neq 0$ and so we have not reached the root. It should be noted that this isn't because there are no roots, as can be seen in the graph below.

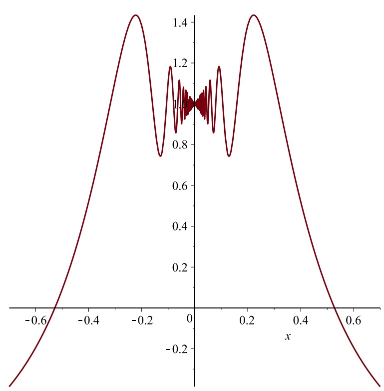

Figure 16: The function $f(x)$ which can deceive Newton Raphson.

In fact it is not hard so show that Newton-Raphson method only works as a numeric way of solving equations when certain conditions are met. Suppose that $x_n \to \alpha$ as $n \to \infty$ and for all x we have $|f'(x)| < M$ where M is some finite number. Then using the linear approximation which the Newton-Raphson method comes from,

$$(x_{n+1} - x_n)f'(x_n) = f(x_n)$$

since we know that $x_n \to \alpha$ we know that the left hand side disappears because the factor $(x_{n+1} - x_n) \to 0$. Thus we end up with the equation $\lim_{n\to\infty} f(x_n) \to 0$. This means that α is a root of the equation $f(x)$ as we have shown that $f(\alpha) = 0$.

The issue we had with the deceitful example is that the derivative is unbounded, in fact it is not hard to show that $f'(x_n) = 2^{n+2}\pi \to \infty$ as $n \to infty$, meaning the condign of M was not satisfied. Naturally we must now ask when does the sequence actually converge to a root. That is to say suppose for some α we have a function $f(\alpha) = 0$, we must deduce the conditions on x_0, f and α such that $x_n \to \alpha$ as $n \to infty$. The below theorem answers this exact problem.

Theorem 4.1. *Newton Raphson Converges to Root*
Suppose that a function $f, f'andf''$ are all continuous for x near α where $f(\alpha) = 0$ and $f'(\alpha) \neq 0$. Then provided that x_0 is sufficiently closet to α we get that α is the limit of the sequence when using the Newton-Raphson method of iteration.

Proof. Let the function $N(x)$ be the Newton-Raphson formula such that we have a sequence $x_{n+1} = N(x_n)$ with $N(\alpha) = \alpha$ for some value,

$$N(x) = x - \frac{f(x)}{f'(x)} \tag{13}$$

Since $f, f'andf''$ are continuous and $f'(\alpha) \neq 0$ there exists a c such that $|f'(x)| \geq c$ and $|f(x)f''(x)| \leq \frac{1}{2}c^2$ for all x near α. Furthermore, we see that

$$\begin{aligned} N'(x) &= \frac{f(x)f''(x)}{f'(x)^2} \\ \Leftrightarrow |N'(x)| &\leq \frac{1}{2} \end{aligned}$$

Applying the mean value theorem to N we see that

$$\begin{aligned} x_{n+1} - \alpha &= N(x_n) - N(\alpha) \\ &= (x_n - \alpha)N'(x) \end{aligned}$$

where x is some number between α and x_n. From this we can see that if x_0 is close enough to α we have,

$$\begin{aligned} |x_1 - \alpha| &\leq \frac{1}{2}|x_0 - \alpha| \\ |x_2 - \alpha| &\leq \frac{1}{4}|x_0 - \alpha| \\ &\vdots \\ |x_n - \alpha| &\leq \frac{1}{2^n}|x_0 - \alpha| \end{aligned}$$

Therefore $|x_n - \alpha| \to 0$ as $n \to \infty$, because $\frac{1}{2^n} \to 0$, which means that $x_n \to \alpha$. QED. □

We may cone again see that the deceitful function does not satisfy the conditions of this theorem because the first derivative is not defined at 0 and the second does not exist.

However, in truth the deceitful function is much more complicated than what we are interested in, our focus throughout this investigation will largely be on polynomials.(This is a real benefit as all polynomials are holomorphic) Further it may seem as though the Newton-Raphson function is very different from the other sequences we have studied, because it is based on another function, this is not the case. The Möbius function too was based on other functions, two linear functions, and so we should keep in mind that everything we have found out that is generally true about sequences we know is true here and vice versa. In fact we can think of the Möbius function as a special case of the Newton

Raphson sequence as we can solve the differing equation,

$$\begin{aligned}
\frac{ax+b}{cx+d} &= x - \frac{f(x)}{f'(x)} \\
\frac{f'(x)}{f(x)} &= \frac{cx+d}{cx^2+(a+d)x+b} \\
\ln(f(x)) &= \ln(cx^2+ax+dx+b)\sqrt{-a^2-2ad+4bc-d^2} - \\
&\quad 2\arctan\left(\frac{2cx+a+d}{\sqrt{-a^2-2ad+4bc-d^2}}\right)a + \\
&\quad \frac{2\arctan\left(\frac{(2cx+a+d)}{\sqrt{(-a^2-2ad+4bc-d^2)}}d\right)}{\sqrt{-a^2-2ad+4bc-d^2}} \\
&\quad +k
\end{aligned}$$

In fact many iterative sequences can be defined as special cases of the Newton-Raphson method provided we can solve the differential equation, however as we have just seen above it is usually a terrible way of thinking about the problem.

All the sequences we have seen so far we have studied only in terms of real numbers with the exception of a short discussion on Möbius transformation, all the while suggesting that there is a lot to be gained from looking in the complex plane. The Newton-Raphson iteration formula demands far too much to be complex to ignore it, therefore we will no longer limit ourselves to real numbers but now look at the properties of the complex Newton-Raphson sequence.

4.1 Complex Newton-Raphson

When sequences are brought into the complex plane some very peculiar things can happen, we see a great deal of structure in how they converge, but we also see that the plane is split up in a more complicated way than for real vales. In fact what we see is that there are regions which converge to one value and others that converge to other values and these regions can have very complicated structures. So to begin the investigation it is best to look at a simple special case first,

$$f(z) = z^2 + 1$$

This function is the most basic one we could have chosen, all the numbers used to define it are real but its roots are complex and that is what we want. Applying the Newton-Raphson formula we see that we get the equation,

$$\begin{aligned}
N(z) &= z - \frac{f(z)}{f(z)} \\
&= z - \frac{z^2+1}{2z} \\
&= \frac{z^2-1}{2z} \\
z_{n+1} &= \frac{z_n^2-1}{2z_n}
\end{aligned}$$

We wish to deduce which root of $f(z) = 0$, if any, the sequence z_n approaches as we iterate infinitely. We begin by redefining the transformation $T(z)$ we saw in section 1.4 to suit our new function,

$$T(z) = \frac{z-i}{z+i}$$

We see now that we can use this map to 'control' the Newton-Raphson by the following result. For clarity we will take $f_n(x)$ to be the n^{th} iteration of a function and $f^n(x)$ and $f(x)^n$ to be the n^{th} power of a function,

$$\begin{aligned} T(N(z)) &= T^2(z) \\ T(N_2(z)) &= T^2(N(z)) = T(z)^4 \\ &\vdots \\ T(N_n(z)) &= T(z)^{2^n} \end{aligned} \tag{14}$$

This is easy to show by induction. Something else we can see about $T(z)$ is that $|T(z)| = 1$ if and only if z is equidistant from the two roots. This comes from the fact that modulus really means distance of two complex numbers in the plane. Let us assume that z_0 is not equidistant from the roots therefore either $|T(z)| < 1$ or $|T(z)| > 1$.

Suppose that the former is true so that we have $|T(z)| < 1$, this means that $|T(z)^k| \to 0$ as $k \to \infty$ and so the right hand side of (14) disappears. Recalling that as $k \to \infty$ we have $n \to \infty$ we see that the left hand side disappears also, which means that that $T(z) = 0$. This happens if and only if $z = i$ which means that $N_n(z) \to i$ as $n \to \infty$. Hence, the Newton-Raphson iteration formula has converged to a root.

Suppose now that $|T(z)| > 1$, then we see that $T(z)^k \to \infty$ as $k \to \infty$. This can only be true if $z \to -i$ and so from 14 we see that this means that $N_n(z) \to -i$ as $k \to \infty$ where n and k have the same relationship as before.

So we now know that whenever z is not equidistant to the roots it converges to one of them. To see a visual representation of this we will colour the plane according to which root it converges, blue for i and red for $-i$ which is the image on the top of the next page. This graph demonstrates the *basins of attraction* of the function which to say it shows what ranges of values go where upon iteration. We can see three distinct sets of numbers in the plane,

$$\begin{aligned} R_1 &= \{z : N^k(z) \to i \text{ as } k \to \infty\} \\ R_2 &= \{z : N^k(z) \to -i \text{ as } k \to \infty\} \\ L &= \{z : N^k(z) \not\to i \text{ or } -i \text{ as } k \to \infty\} \end{aligned}$$

We have already deduced the condition for the first two sets which is that the starting value has to be closer to one root than the other. However, the set L has the property that the starting value is equidistant from both roots and it has the curious property of bisecting the plane. We can also deduce that if $x_0 \in L$, then $x_n \in L$ for all n by looking at T and its relationship to the Newton-Raphson map earlier. In fact what we can see in this special case is that the set L is the set of real numbers, which means that should you chose a real number as a starting value it will never converge to either root. Furthermore, it is quite clear that whenever the roots are complex conjugates the set L is the real numbers.

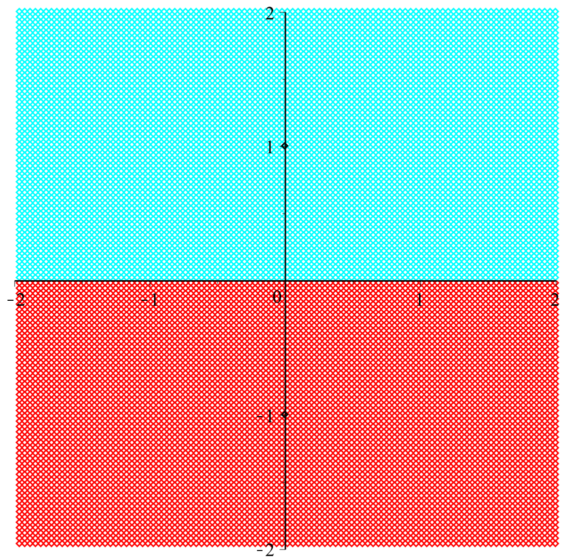

Figure 17: (Appendix H) The Basin of Attraction for $f(z) = z^2 + 1 = 0$.

Now we can ask a much more interesting question, is the set L always a line for quadratic functions? and what values, if any, does it always go through? To answer these questions we must leave the specific case and ask what happens for the general quadratic function $f(z) = z^2 + az + b$ when we iterate using the Newton-Raphson method. Note that by factorising or otherwise it is trivial to prove that for the two roots $\alpha, \beta \in \mathbb{C}$ we have that $a = -(\alpha + \beta)$ and $b = \alpha\beta$. Hence we can rewrite the functions N and T in the following general way,

$$\begin{aligned} N(z) &= \frac{z^2 - b}{2z + b} \\ T(z) &= \frac{z - \alpha}{z - \beta} \end{aligned}$$

Suppose then that $|T(z)| = 1$, as this is an equivalent condition for L, which means that $|z - \alpha| = |z - \beta|$ from which it obvious that we are on a straight line. There is no curve with the property that it is equidistant from two fixed points at all times. Furthermore, we know that the line must bisect the plane between α and β, because it is equidistant, so we may deduce two values in runs through,

$$\begin{aligned} l_1 = \frac{\alpha + \beta}{2} &= -\frac{a}{2} \\ l_2 &= -\frac{a}{2} + i\frac{1}{2}\sqrt{a^2 - 4b} \end{aligned}$$

The first value comes from taking the average of the roots as this is equidistant from both, the second comes from then rotating the line which connects one of the roots and l_1 by $\frac{\pi}{2}$ radians to get another value on L because it is perpendicular to the connecting line. Now we have two values which define the line and so we know what the set L is.

We must now observe that because the formula contains a fraction with variable denominator we may run into undefined values. Indeed when we chose $z = -\frac{1}{2}a$ or $z = \beta$ we see that we either have $N(z)$ or $T(\beta)$ having a denominator of zero. We ran into a similar issue with the Möbius transformation and the way we solved that is by saying we were in the field of $\mathbb{R} \cup \infty$, what we really mean by this is that we live on the 'Riemann Sphere' where infinity is a well-defined number.

There is one case we have not looked at yet which is what happens when we have repeated roots of $f(z) = 0$ such that $\alpha = \beta$. In this case it easy to see that there is no line L, in fact in this case every single starting value will converge to the root. This is actually very simple to prove, we just redefine the transforming function such that $T(z) = z - \alpha$ and get the resulting identity $T(N_n(z)) = \frac{1}{2^n}T(z) \to 0$ as $n \to \infty \Rightarrow N_n(z)$ showing that it converges. We now know everything about applying the Newton-Raphson formula to

quadratic equations. We know that when there are two roots that are complex conjugates no real starting values will ever let the method converge, we know that there is always a line L which bisect the plane for convergence to the α and β and we know two values which it goes through. Furthermore we know that for any quadratic function that has only one repeated root the Newton-Raphson method will always converge to the root.

4.2 Conclusions on the Newton-Raphson method

We have already stated everything we have learnt about the complex quadratic cases. However, there is much more work to be done on the sequence, for example asking what happens when we move into cubic equations is already a very good questions. As it happens the Newton-Raphson map for the simple cubic equation $f(z) = z^3 - 1 = 0$ (with roots $1, e^{\frac{2\pi i}{3}}, e^{-\frac{2\pi i}{3}}$) is

$$N(z) = \frac{2z^3 + 1}{3z^2}$$

This map cannot be reduced to anything simple using a transforming function and the basins of attraction for it are extremely complicated. There are three regions corresponding to each of the solutions and they all have the same boundary, which miscalled a Julia set of N. It is extremely bizarre to think of three regions in a plane sharing one boundary. Indeed the Julia set behaves like the line L where any value in it will never converge to a root when we iterate the method, but the reason we have chosen not to do any analysis on the cubic or higher degree polynomials is that the set is a fractal. It has infinite complexity and magnifying it arbitrarily gives no more structure. Therefore the analysis would have been far more difficult than anything else we have seen.

We will not progress any further with this sequences, instead on the next page are different fractal images demonstrating basins of attraction of various functions showing how beautiful this sequence can be. (The code used to generate all of these images can be found in appendix H).

Julia Sets for various Polynomials under Newton-Raphson Iteration

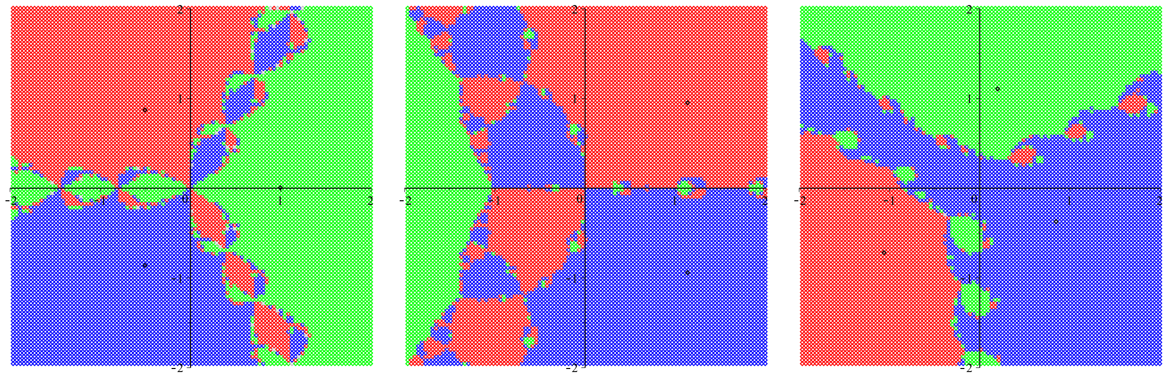

Figure 18: (A) $f(z) = z^3 - 1$ (B) $f(z) = z^3 - 3z + 5$ (C) $f(z) = 3z^3 - 2iz + 4i$

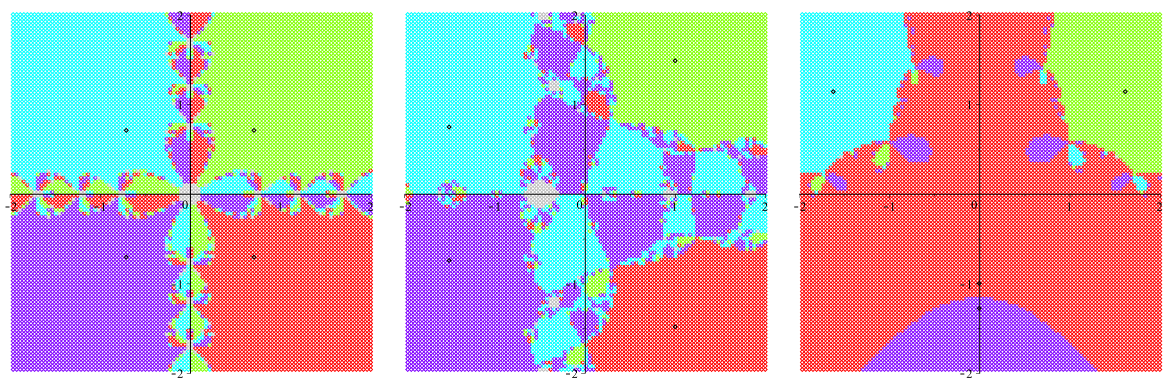

Figure 19: (A) $f(z) = z^4 + 1$ (B) $f(z) = z^4 + z^3 + 4z + 9$ (C) $f(z) = z^4 + 6iz + 5$

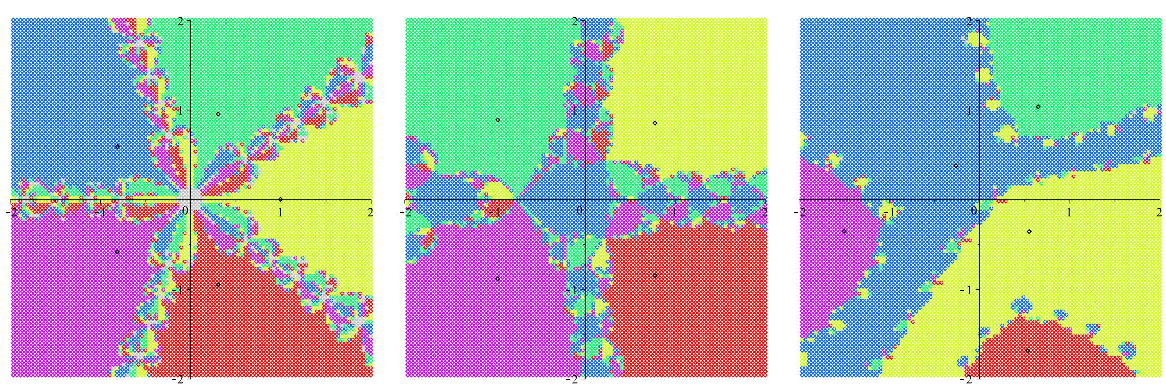

Figure 20: (A) $f(z) = z^5 - 1$ (B) $f(z) = z^5 + 3z^4 + z^3 + 2z + 6$ (C) $f(z) = z^5 + iz^4 + 3z^2 - z + i$

References

[1] Aneurin Jackson; 'An introduction and investigation into number theory.'; 07/2014.

[2] Eleanor Bell; 'Envelopes of lines and curves and surfaces of constant width'; 07/2014

[3] J. W. Bruce, P. J. Giblin, P. J. Rippon; 'MICROCOMPUTERS AND MATHEMATICS'; ISBN 0 521 37515 0.

[4] Ke Chen, Peter Giblin, Alan Irving; 'MATLAB'; ISBN 0 521 63078 9.

[5] http://www.wolframalpha.com/input/?i=logistic+map+r

[6] http://mathworld.wolfram.com/PolynomialDiscriminant.html

[7] http://mathworld.wolfram.com/NewtonsMethod.html

[8] http://physics.ucsc.edu/~peter/242/logistic.pdf

[9] http://mathonweb.com/help_ ebook/html/expoapps.htm


```
Note: The following codes have been presented as they look on screen,
this means  a simplification of the syntax. For anyone willing to reproduce
the codes it must be noted that all procedures are written
on single lines or in command boxes.
```

A Exponential Function Convergence when Changing the Parameter

```
#Maple 16
restart
with(plots): with(plottools): with(ColorTools):
f:=(l,x) ->  l^x:
P := Array([seq(1 .. 150*10)]):
sets := Array([seq(1 .. 150)]):

makeSets := proc(l::float)
local i, L, x;
L := [seq(1 .. 10)];
x := 0.0;
for i to 8 do
x := f(l, f(l, x))
end do;
for i to 10 do
x := f(l, x);
L[i] := point([l, x])
end do;
return L end proc

for l from 0.1e-1 by 0.1e-1 to 1.5 do
sets[floor(100*l)] := makeSets(l)
end do

c := 1:
for s in sets do
for i in s do
P[c] := i;
c := c+1
end do
end do;

display(seq(P[i], i = 1 .. 150*10), view = [0 .. 1.5, 0 .. 3])
```

B Exponential Sequence Cobweb Diagram

```
#Maple 16
restart
f := (l, x) -> l^x;
S := [seq(1 .. 41)];
L := plot(y, y, color = black);

Cobweb := proc (l::integer)
local i, L, x, Mh, Mv, r;
r := (1/20)*l;
x := 8;
L := [seq(1 .. 100)];
Mh := [seq(1 .. 98)];
Mv := [seq(1 .. 98)];
for i to 50 do
L[i] := x;
x := f(r, x)
end do;
for i from 2 to nops(L)-1 do
x := f(r, x);
Mh[i-1] := line([L[i-1], L[i]], [L[i], L[i]], color = "blue");
Mv[i-1] := line([L[i], L[i]], [L[i], L[i+1]], color = "red")
end do;
return [seq(Mh[i], i = 1 .. nops(Mh)), seq(Mv[i], i = 1 .. nops(Mv))]
end proc;

j := 1;
for i to 30 do
S[j] := Cobweb(i);
j := j+1
end do;

Web := proc (t::float)
local s, p;
s := S[floor(t)];
p := PLOT(seq(s[i], i = 1 .. nops(s)))
end proc;

C:=animate(Web, [r], r=1..3 0):
F := animate(plot, [f(r, y), y = 0 .. 10], r = 0 .. 1.5);
display([C, F, L], view = [0 .. 10, 0 .. 10]);
```

C Bifurcation Diagram

```
#Maple 16
restart
with(plots); with(plottools);
f := (l, x) -> l*x*(1-x);
P := Array([seq(1 .. 401*50)]); sets := Array([seq(1 .. 401)]);

makeSets := proc (l::float)
local i, L, x;
L := [seq(1 .. 50)];
x := .5;
for i to 1000 do
x := f(l, x)
end do;
for i to 50 do
x := f(l, x);
L[i] := point([l, x])
end do;
return L
end proc;

for l from 0. by 0.01 to 4 do
sets[floor(100*l+1)] := makeSets(l)
end do;

c := 1;
for s in sets do for i in s do
P[c] := i;
c := c+1
end do;
end do;

display(seq(P[i], i = floor(50*(2.8*100)) .. 401*50), view = [2.8 .. 4, 0 .. 1]);
```

D Quadratic Cobweb Diagram

```
#Maple 16
restart
with(plots); with(plottools);
x := 0.01; l := 2.9;
f := (x) -> l*x*(1-x);
L := [x, seq(1 .. 30)];
Mh := [seq(1 .. 29)];
Mv := [seq(1 .. 29)];

for i to 31 do
L[i] := x;
x := f(x)
end do;

for i from 2 to nops(L)-1 do
Mh[i-1] := line([L[i-1], L[i]], [L[i], L[i]], color = "blue");
Mv[i-1] := line([L[i], L[i]], [L[i], L[i+1]], color = "red");
end do;

display(seq(Mh[i], i = 1 .. nops(Mh)),
        seq(Mv[i], i = 1 .. nops(Mv)),
        plot(y, y, color = black), plot(f(y),
        y, color = green), view = [0 .. 1, 0 .. 1]);
```

E Basin of Quadratic 2-Cycle

```
#Maple 16
restart
with(plots); with(Optimization);

q2 := (x, l) -> l^2*x*(1-x)*(l*x^2-l*x+1);
r1 := (l) -> ((1/2)*l+1/2+(1/2)*sqrt(l^2-2*l-3))/l;
r2 := (l) -> ((1/2)*l+1/2-(1/2)*sqrt(l^2-2*l-3))/l;

obj := proc (x, l) -> abs(x-r1(l));
monster1 := (x, l) -> (q2(x, l)-r1(l))/(x-r1(l));
monster2 := (x, l) -> (q2(x, l)-r2(l))/(x-r2(l));

cnsts1 := [monster1(x, l) < 1, monster1(x, l) > -1];
cnsts2 := [monster2(x, l) < 1, monster2(x, l) > -1];

p1 := inequal([cnsts1, cnsts2], x = 0 .. 1, l = 3 .. 1+6^.5,
              optionsexcluded = (color = black),
              optionsfeasible = [[color = green], [color = blue]]);
p2 := contourplot(obj(x, l), x = 0 .. 1, l = 3 .. 1+6^(1/2));
display(p1, p3);
```

F Swallowtail Diagram

```
restart
with(plots); with(plottools);

y := (t, x) -> -4*t^3-2*x*t;
z := (t, x) -> -t^4-x*t^2-(-4*t^3-2*x*t)*t;

plot3d([x, y(t, x), z(t, x)],
       t = -10 .. 10, x = -10 .. 10,
       view = [-10 .. 10, -10 .. 10, -10 .. 10],
       grid = [200, 200]);
```

G Newton-Raphson Cobweb Diagram

```
#Maple 16
restart
with(plots); with(plottools);

a := 1/20; b := -2;
f := (x) -> x^2+a*x+b;
N := (x) -> x-(x^2+a*x+b)/(2*x+a);

newtonPoints := proc (x)
local i, L, z;
L := [seq(1 .. 11)]; z := x;
for i to 10 do
L[i] := z; z := N(z)
end do;
return L
end proc;

makeLines := proc (L::list)
local i, Mv, Mt, x1, x2, M;
Mt := [seq(1 .. nops(L)-1)]; Mv := [seq(1 .. nops(L)-1)];
M := [seq(1 .. nops(Mv)+nops(Mt))];
x1 := 0; x2 := 0;
for i from 2 to nops(L)-1 do
x1 := L[i-1]; x2 := L[i];
Mt[i-1] := line([x1, f(x1)], [x2, 0], color = blue);
Mv[i-1] := line([x2, 0], [x2, f(x2)], color = red)
end do;
for i to nops(L)-1 do
M[2*i] := Mt[i]; M[2*i-1] := Mv[i]
end do;
return M
end proc;

P := newtonPoints(6);
L := makeLines(P);

Con := proc (t::float)
local p;
p := PLOT(L[floor(2*t)], L[floor(2*t-1)])
end proc;

C := animate(Con, [r], r = 1 .. (1/2)*nops(L), frames = (1/2)*nops(L));
F := plot(f(x), x, color = black);

display([F, C], view = [-1 .. 20, -1 .. 20]);
```

H Newton-Raphson Basin of Attraction

```
#Maple 16
restart
with(plots); with(plottools); with(ColorTools);
f := (x) -> x^4-2*x-I;
r := 101;
R := [solve(f(z), z)];
N := (x) -> x-f(x)/(D(f))(x);
P := seq(point([Re(R[i]), Im(R[i])]), i = 1 .. nops(R));

for i from 0 to r do
for j from 0 to r do
zx := -2+4*i/r;
zy := -2+4*j/r;
z := zx+I*zy;
for k from 0 to 20 do
z := evalf(N(z))
end do;
t := 1;
for k to nops(R) do
if `and`(abs(z-evalf(R[k])) < 0.01, t = 1) then
p[i, j] := pointplot([zx, zy], color = Color("HSV", [k/nops(R), 1, 1]));
t := 0
end if;
end do;
if t = 1 then
p[i, j] := pointplot([zx, zy], color = grey)
end if;
end do;
end do;

display([seq(seq(p[i, j], i = 1 .. r), j = 1 .. r), P], view = [-2 .. 2, -2 .. 2]);
```

I Newton Raphson Route to Route

```
"This is a script which was not included in the investigation, but is still
of some interest. It shows how consecutive iterations of the Newton-Raphson
formula in the complex plane gets closer  to the root, effectively showing
the path the sequence shows. When running it pay express attention to the
fact that it curves slightly and that usually 5 iterations is sufficient
to be close enough to a root not to be able to see any real difference by eye.
This last point comes from the fact that the method has a quadratic
rate of convergence."

#Maple 16
restart
with(plots); with(plottools); with(ColorTools);

f := (z) -> z^3-I*z;
R := [solve(f(x), x)];

for i to nops(R) do
R[i] := point([Re(R[i]), Im(R[i])])
end do;

N := (z) -> z-f(z)/(D(f))(z);

newtonPoints := proc (z)
local i, L, x;
x := z;
L := [seq(1 .. 11)];
for i to nops(L) do
L[i] := x;
x := N(x)
end do;
return L
end proc:

joinPoints := proc (L::list)
local i, M, z1, z2;
M := [seq(1 .. nops(L)-1)];
z1 := 0;
z2 := 0;
for i from 2 to nops(L) do
z1 := L[i-1];
z2 := L[i];
M[i-1] := line([Re(z1), Im(z1)], [Re(z2), Im(z2)],
               color = ColorTools:-Color([.7/(i-1)^.5, .6/(i-1)^.5, .8/(i-1)^.5]))
end do;
return M;
```

```
end proc:

P := newtonPoints(3+3*I):
L := joinPoints(P):

Connect := proc (t::float)
local p;
p := PLOT(seq(L[floor(i)], i = 1 .. floor(t)))
end proc:

Walk := animate(Connect, [r], r = 1 .. nops(L), frames = nops(L)):
r := PLOT(op(R)):

display([Walk, r], view = [-5 .. 5, -5 .. 5]):
```