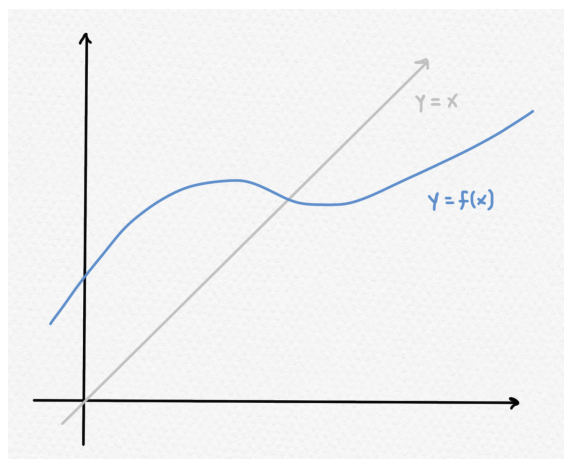# 3 Fixed-point iteration

Here are the student questions that I'm attempting to answer thoroughly with this write-up:

- When does fixed-point iteration converge/diverge at a certain fixed point?

- In the convergence criterion for FPI, what exactly does "sufficiently close initial guesses" mean?

- How do we know (rigorously) that FPI converges in certain cases? When can we guarantee convergence?

- How can we say whether $f'(\phi) \in [-1, 1]$ if we don't actually know the value of ϕ yet?

- What if FPI diverges for one of the fixed points of a certain function? Is it impossible to approximate?
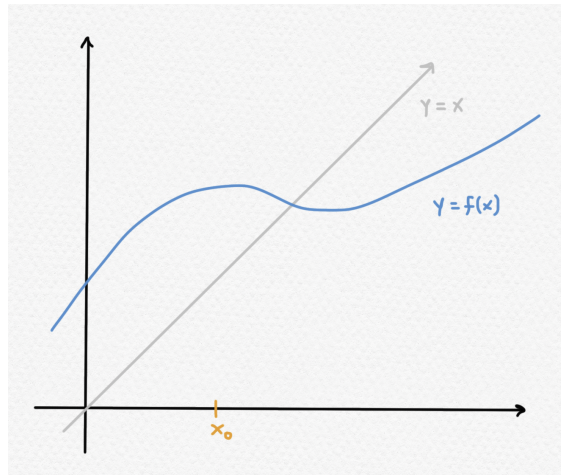
3.1 Cobweb plots and convergence

One useful way of visualizing fixed-point iteration is to use a **cobweb plot**. Given a function f with a fixed point ϕ, plot the graphs of $y = x$ and $y = f(x)$ on the same plane, like this:
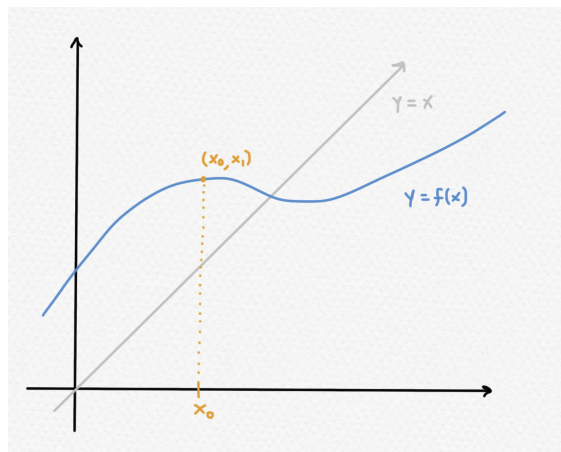

Mark the initial guess x_0 on the x-axis:

Find the point $(x_0, f(x_0))$ on the curve $y = f(x)$ corresponding to this x-value. Notice that since FPI is defined using the recurrence $x_{n+1} = f(x_n)$, this is the same as the point (x_0, x_1).


Now, draw a horizontal line segment connecting this point to the line $y = x$, obtaining the point (x_1, x_1), and then find the point $(x_1, f(x_1))$ by drawing another vertical line segment from this point to the corresponding point on $y = f(x)$ with the same x-value. Notice that this point is the same as (x_1, x_2).
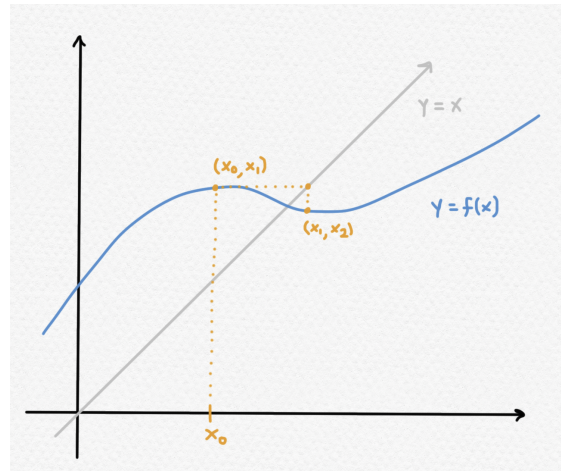
Now you can repeat this process starting with the point (x_1, x_2) rather than (x_0, x_1). By repeating this process, you will obtain a sequence of points $(x_0, x_1), (x_1, x_2), (x_2, x_3)$ and so on, and by examining how they appear on the graph, you can tell whether or not FPI converges with the given starting value. For instance, here's what a zoomed-in version of the following cobweb plot might look like after several iterations:
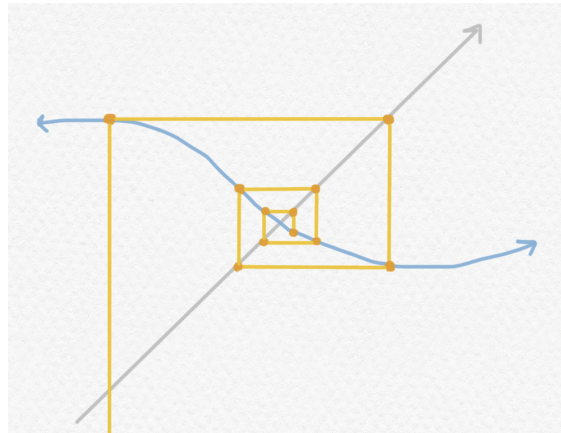
See how the estimates seem to be getting closer and closer to the fixed point? This is a case in which fixed-point iteration is *successful* at producing a sequence of increasingly accurate approximations. Here's an example with a different function where it does not turn out so well:
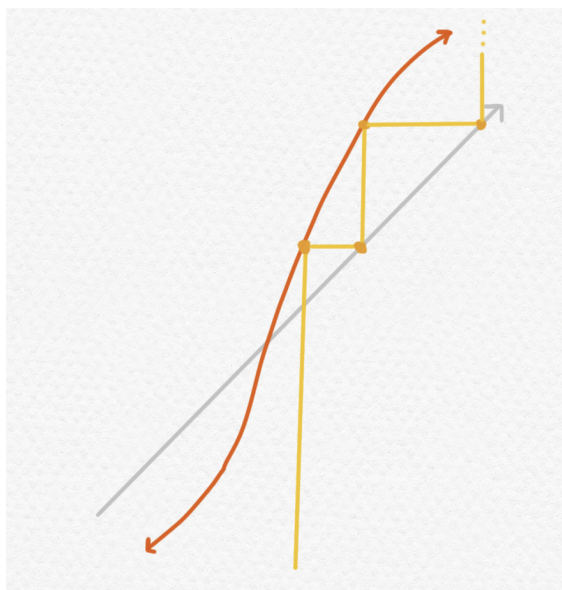
Notice how the approximations *move away* from the actual fixed point after a few iterations? If you try sketching some more examples by hand, you might notice that when the curve $y = f(x)$ crosses the line $y = x$ with a *very steep* slope, FPI fails to converge to the fixed point in question, but when it crosses $y = x$ with a *very shallow* slope, FPI converges to that fixed point, as long as the initial guess is close enough to start with. To be specific, we will have guaranteed convergence if the *slope of the function's curve has magnitude less than one at the fixed point*, or if $|f'(\phi)| < 1$, where ϕ is the fixed point.

By the way, there are a couple of very neat cobweb plotting tools online, like this one.

3.2 Rigorous proof of convergence

A rigorous proof that FPI converges requires you to use - you guessed it - Taylor's Theorem! Roughly, Taylor's Theorem tells us that "well-behaved" (twice-differentiable) functions behave approximately like linear functions in small neighborhoods of each of their points, and it's much easier to determine how FPI behaves on linear functions than on arbitrary functions.

Suppose we have a function $f : [a, b] \to \mathbb{R}$ that is twice-differentiable in the interval (a, b), and has a fixed point $\phi \in [a, b]$. By Taylor's Theorem, we may expand f into its Taylor Series centered at $x = \phi$, so that we have

$$f(x) = f(\phi) + (x - \phi)f'(\phi) + O((x - \phi)^2)$$

or, since ϕ is a fixed point,

$$f(x) = \phi + (x - \phi)f'(\phi) + O((x - \phi)^2)$$

In case you're confused by the big-O notation on the right side of the equation, you just need to know that this means that $f(x)$ is equal to $\phi + (x - \phi)f'(\phi)$ *plus some error term* which we don't know the exact value of, but which is guaranteed to be less than $C(x - \phi)^2$ for all x, for some constant $C > 0$. Notice than the error term of $C(x - \phi)^2$ shrinks to zero much faster than $(x - \phi)f'(\phi)$ as x gets closer and closer to ϕ, meaning that the relative error in this approximation gets better and better as x gets closer to ϕ. (Assuming, for now, that $f'(\phi)$ is $\neq 0$. In that case, things are a little more complicated.)

Now, recall that fixed-point iteration starts with some initial guess x_0 for the value of the fixed point, and calculates successive approximations as follows:

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

We may also define a sequence of numbers measuring the *error* between the FPI approximations x_n, and the actual value of ϕ:

$$e_n = |x_n - \phi|$$

If things go well (i.e. if FPI converges) then these error values should tend towards zero, meaning that our approximations are getting arbitrarily close to the actual value.

Now, plugging $x = x_n$ into the approximation that we obtained from Taylor's Theorem, we have the following:

$$f(x_n) = \phi + (x_n - \phi)f'(\phi) + O((x_n - \phi)^2)$$

However, since $x_{n+1} = f(x_n)$ by definition, we can rewrite this as follows:

$$x_{n+1} = \phi + (x_n - \phi)f'(\phi) + O((x_n - \phi)^2)$$

or, subtracting ϕ from both sides,

$$x_{n+1} - \phi = (x_n - \phi)f'(\phi) + O((x_n - \phi)^2)$$

If we take absolute values of both sides (and use the triangle inequality), we can rewrite this as

$$|x_{n+1} - \phi| = |x_n - \phi| \cdot |f'(\phi)| + O(|x_n - \phi|^2)$$

but this actually expresses a relationship between the errors of successive terms:

$$e_{n+1} = |f'(\phi)| \cdot e_n + O(e_n^2)$$

Using the definition of big-O notation, this gives us the following inequality, for some unspecified constant C:

$$e_{n+1} \leq |f'(\phi)| \cdot e_n + Ce_n^2$$

or

$$e_{n+1} \leq \left(|f'(\phi)| + Ce_n\right) \cdot e_n$$

Hence, the term $|f'(\phi)| + Ce_n$ tells us the maximum factor by which the error is multiplied after one iteration. This means that *if we guarantee that this quantity is less than one*, then *we will*

know that the error is decreasing. In particular, if we are able to choose an initial guess x_0 that is sufficiently close to ϕ such that the error e_0 is less than

$$\frac{1 - |f'(\phi)|}{C}$$

then we will know that $|f'(\phi)| + Ce_0 < 1$, meaning that $e_1 < e_0$. Then, by similar reasoning, we could show that $e_2 < e_1$, and $e_3 < e_2$, and so on. Hence, if our initial guess has an error satisfying

$$e_0 < \frac{1 - |f'(\phi)|}{C}$$

then the errors are *guaranteed to shrink to zero*, since they will be shrinking by a factor that keeps getting smaller and smaller at each iteration. When the errors converge to zero, of course, the approximations x_n must be converging to the true fixed point value ϕ.

Of course, this is hardly practical if we don't actually know the exact value of the constant C. Taylor's Theorem allows us to actually determine one possible value of this constant: it tells us that

$$C = \frac{|f''(\xi)|}{2}$$

for some $\xi \in [a, b]$. Of course, this still does not actually tell us the precise value of ξ - but it does allow us to give an *upper bound*. If we are able to determine the *maximum value* of $|f''(x)|$ on $[a, b]$ (let's call it M), then we can say that the constant C guaranteed by Taylor's Theorem satisfies

$$C \leq \frac{M}{2}$$

which means that if we conservatively choose our first estimate x_0 such that the error is less than

$$e_0 < \frac{1 - |f'(\phi)|}{M/2} \leq \frac{1 - |f'(\phi)|}{C}$$

then our estimates are also guaranteed to converge. Hence, a precise sufficient condition for convergence of FPI is the following:

$$e_0 < \frac{1 - |f'(\phi)|}{\frac{1}{2} \max_{x \in [a,b]} |f''(x)|}$$

This is not to say that other initial guesses x_0 will definitely *fail to converge* - there may well be other initial guesses that work. This is just one criterion that will *guarantee* convergence.

Notice, by the way, that if $|f'(\phi)| > 1$, the above criterion cannot be satisfied, since $1 - |f'(\theta)|$ will be negative, but the value e_0 can only be nonnegative. This corresponds to our observation that when $|f'(\phi)| > 1$, fixed-point iteration will diverge.

3.3 Convergence without knowing the fixed point

One question that might have occurred to you is the following: if we don't know the actual value of a fixed point ϕ of some function $f(x)$ (which we probably don't, if we're trying to approximate it) then how can we determine the magnitude of the derivative $|f'(\phi)|$ to determine whether or not FPI converges to that fixed point? Isn't it impossible to determine convergence without knowing the fixed point in the first place?

As a matter of fact, it is possible to determine convergence without knowing ϕ beforehand. This is because, to know that FPI converges, we don't need to know the *exact* value of $f'(\phi)$ - we just need to know whether it has magnitude less than 1. And if we know that ϕ lies within some interval $[a, b]$, and that $f'(x)$ is bounded between -1 and 1 for all x inside of that interval, then we can conclude that $|f'(\phi)| < 1$ without ever calculating its value explicitly.

To take an example, consider the following function:

$$f(x) = e^x + 2x$$

Without graphing this function, let's figure out how many fixed points it has, and whether or not FPI converges to those fixed points. First of all, a fixed point of f is a solution to the equation

$$e^x + 2x = x$$

or, by subtracting x from both sides,

$$e^x + x = 0$$

Notice that both e^x and x are increasing functions of x, meaning that $e^x + x$ is also a strictly increasing function of x. Hence, $e^x + x$ cannot have *more than one zero* - since, once it crosses the x-axis, it cannot "go back down" again. So now we just need to determine whether it has one zero or no zeros at all.

To show that it does in fact have a zero, consider the following two values:

$$f(0) - 0 = e^0 + 0 = 1$$

$$f(-1) - (-1) = e^{-1} - 1$$

We can see that $f(0) > 0$, and also that $f(-1) < -1$, since $1/e$ is less than 1. This means that $f(x) - x$ must change signs from negative to positive somewhere between $x = -1$ and $x = 0$. This means that it must have a zero somewhere between these two x-values, and therefore f must have a fixed point somewhere in the interval $(-1, 0)$. (Technically, this relies on the fact that f is a continuous function, and uses the **Intermediate Value Theorem**.)

So now we know that there's a unique fixed point between $x = -1$ and $x = 0$. How can we determine whether FPI converges to this fixed point? Let's consider the derivative of f:

$$f'(x) = e^x + 2$$

Recall that the function e^x is always positive, meaning that $f'(x)$ is always greater than 2. Hence, we can conclude that FPI will *not converge* to the fixed point of f, since we need $f'(\phi)$ to have magnitude less than 1 to confirm convergence.

3.4 Solving other equations using FPI

As it is presented, it seems that fixed-point iteration can only be used to solve equations of the following form:

$$f(x) = x$$

However, FPI can actually be used to solve a wide variety of other equations, provided that it converges. For instance, suppose you want to find a root of a function f, i.e. a solution to the equation

$$f(x) = 0 \tag{i}$$

Notice that this equation is equivalent to the following:

$$f(x) + x = x \tag{ii}$$

by adding x to both sides of the original equation. That is, every solution to equation (i) is a solution to (ii), and vice versa. If we define a second function called g by letting $g(x) = f(x) + x$, then the second equation states that

$$g(x) = x \tag{iii}$$

so that the equations (i), (ii) and (iii) all have exactly the same solutions. Notice that equation (iii) expresses that x is a *fixed point* of the function g! This means that every root of f is a fixed point of g, and vice versa. Since FPI is effective for calculating fixed points, we can also use it for calculating roots of an arbitrary function $f(x)$ by instead calculating fixed points of the transformed function $g(x) = f(x) + x$.

As a matter of fact, we can essentially solve any equation with FPI (again, ignoring convergence issues). If we have *any two functions* f_1 and f_2 and we want to solve the equation

$$f_1(x) = f_2(x)$$

we may rewrite this equation as follows:

$$f_1(x) - f_2(x) = 0$$

Hence, if we let $f(x) = f_1(x) - f_2(x)$, any solution to our equation will be a *root* of the function $f(x)$. And we have just seen how to use FPI to calculate roots!

3.5 Fixing divergent behavior

In the previous section, we ignored matters of convergence, but convergence of FPI is extremely important. After all, if FPI does not converge to a fixed point ϕ of some function f, then FPI will not help us approximate that fixed point. As we've seen, we need $|f'(\phi)| < 1$ for FPI to reliably converge to the fixed point ϕ. But what if we want to approximate a fixed point where $|f'(\phi)| > 1$? Is FPI completely useless in this case?

As it turns out, there is actually a creative way of using FPI to approximate a fixed point ϕ of a function f where $|f'(\phi)| > 1$. Consider once more the following equation:

$$f(x) = x \tag{i}$$

Let us choose some fixed constant $c \in \mathbb{R}$, and add the quantity cx to both sides of this equation. This yields the following:

$$f(x) + cx = (c + 1)x \tag{ii}$$

Note that this equation has *exactly the same solutions* as (i), so if we can approximate the solutions of this equation, then we can approximate the solutions of our original equation. So far so good. Now, let's divide both sides of this equation by $(c + 1)$, assuming that $c \neq -1$. We get the following:

$$\frac{f(x) + cx}{c + 1} = x \tag{iii}$$

Again, the solutions of (iii) are precisely the same as the solutions of (i). In fact, if we define a second function, call it g_c, to equal the left-hand side of this equation:

$$g_c(x) = \frac{f(x) + cx}{c + 1}$$

then we have that (iii) simply states $g_c(x) = x$. Thus, *the fixed points of f are exactly the same as the fixed points of g_c*. If we manage to approximate the fixed points of g_c, then we will have also obtained approximations for the fixed points of f.

If we can't apply fixed-point iteration to f successfully, then maybe it will work on g_c. To know whether FPI is viable for g_c, we need to know what $g_c'(\phi)$ looks like, and whether its magnitude is between 0 and 1. We have the following derivative from the definition of g_c:

$$g_c'(x) = \frac{f'(x) + c}{c + 1}$$

24

meaning that

$$g_c'(\phi) = \frac{f'(\phi) + c}{c + 1}$$

Now, recall that c can be *any constant we want* (other than $c = -1$). Thus, we can simply *choose a value* of c that makes $g_c'(\phi)$ have a value between -1 and 1, guaranteeing convergence.

Let's see an example of this in action by considering the fixed point of the following function that we looked at earlier:

$$f(x) = e^x + 2x$$

We already found that this function has exactly one fixed point ϕ with $f'(\phi) > 2$, meaning that FPI would not be viable for approximating this fixed point. Notice that we have

$$f'(x) = e^x + 2$$

so since this fixed point is negative, we can reason that the slope at this fixed point is no greater than $3 = e^0 + 2$. This means that if we choose $c = -3$, then the quantity

$$g_c'(\phi) = \frac{f'(\phi) + c}{c + 1}$$

is guaranteed to be between -1 and 0, which means that applying FPI to the function

$$g_{-3}(x) = \frac{f(x) - 3x}{-2} = \frac{x - e^x}{2}$$

would give us a sequence of convergent approximations to this fixed point of the desired function $f(x)$. Here's a plot of these two functions, with $f(x)$ in red, $g_c(x)$ in green, and the line $y = x$ in blue:
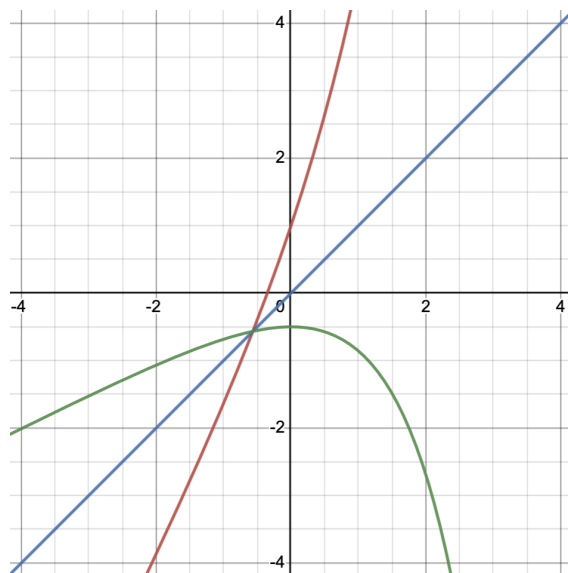
As an interesting aside, not only can we use this trick to get convergent approximations for a fixed point where FPI normally fails, but we can also use it to *accelerate the convergence* for fixed points that FPI converges to *slowly*. For instance, if we have a function $f(x)$ with a fixed point ϕ at which $|f'(\phi)| = 0.99$, FPI will converge to this fixed point for sufficiently close initial guesses, but it may do so extremely slowly, with the error decreasing by approximately 1% each iteration. However, by choosing a suitable value of c, we can find a function g_c with the same fixed point and a smaller value of $|g'_c(\phi)|$.