ALDAR ACADEMIES
AL BATEEN

IB Mathematics HL 12
A Review of the Derivative

1 The Derivative at a Point

Consider a function f, and a point a in the domain of f.[1] If we were to graph $y = f(x)$, we would find that the point $(a, f(a))$ lies on the graph. Similarly, for any b in the domain of f, we could produce another point on the graph, $(b, f(b))$. We would then be able to construct the line passing through $(a, f(a))$ and $(b, f(b))$, after which we could examine the gradient of that line (a line passing through (at least) two points on the graph of a function is called a *secant* line).

This is essentially the idea that leads us to the definition of the derivative at a point a: we consider how the gradient of the secant line would change as the value of b approaches the value of a. The gradient of the secant line would be given by

$$\frac{\Delta y}{\Delta x} = \frac{f(b) - f(a)}{b - a}$$

and b approaching a is of course one way of talking about a *limit*, leading us to consider

$$\lim_{b \to a} \frac{f(b) - f(a)}{b - a} \tag{1}$$

Equation (1) is one definition of the derivative of f at a point a.

In our lessons, we've used an alternative, equivalent definition. Instead of using b to represent a value in the domain of f, we've used $a + h$ instead. Our second point on the graph of the function then has coordinates $(a + h, f(a + h))$, and so the secant line would have gradient

$$\frac{\Delta y}{\Delta x} = \frac{f(a+h) - f(a)}{(a+h) - a}, \text{ or more simply,}$$
$$\frac{\Delta y}{\Delta x} = \frac{f(a+h) - f(a)}{h}$$

The fraction appearing on the right above is sometimes referred to as the *Newton quotient*. Now, instead of letting b approach a, we can let h approach 0 to get the same effect. So, we're left to consider

$$\lim_{h \to 0} \frac{f(a+h) - f(a)}{h} \tag{2}$$

Equation (2) is the definition of the derivative of f at a point a that we used in class.

[1] You may find that it is helpful to consider a particular function and a particular value for a, and I'd recommend thinking of $f(x) = x^2$ and $a = 1$ if you get stuck in any part of the discussion that follows.

If we make use of the Lagrange notation for the derivative of f, then the definition of the derivative of f at a point a is given by

$$f'(a) = \lim_{h\to 0} \frac{f(a+h)-f(a)}{h} \tag{3}$$

Note that, for a particular function f and a particular value a, the derivative of f at a will be a *numerical value*. In class we saw that, for $f(x) = x^2$ and $a = 1$, we get $f'(1) = 2$, and similarly we found that $f'(2) = 4$ and $f'(3) = 6$.

1.1 Questions

1. Use the definition of the derivative to find the derivative of $f(x) = -x^2 + 4x$ at $x = 5$.
2. Consider the function $f(x) = 2x$.
   (a) Use the definition of the derivative to show that the derivative of this function at $x = 3$ is 2.
   (b) Prove that $f'(a) = 2$ for any $a \in \mathbb{R}$.

2 *The* Derivative

After calculating the derivative of a simple function like $f(x) = x^2$ for various values of a, you may start to observe a pattern that would allow you to predict the value of the derivative without calculating the relevant limit. If you were interested in finding the value of $f'(a)$ for several different values of a, you might instead hope to find a *function* that would, given the value of a, allow you to calculate the value of $f'(a)$ without your needing to evaluate a limit. This leads to the notion of *the* derivative, which is exactly the function that gives you those values directly.

Definition (The Derivative)**.** Given a differentiable function f, the derivative of f, represented by f', is given by

$$f'(x) = \lim_{h\to 0} \frac{f(x+h)-f(x)}{h}$$

A function that has a derivative at all points in its domain is called *differentiable*. While we won't pause to consider such functions now, some functions are not differentiable: they may fail to be differentiable at certain points in their domain, or they may be nowhere differentiable.

2.1 Derivatives of Polynomials

When we need to determine the derivative of a given function, we would usually like to avoid having to determine the derivative by using the definition (and so, explicitly considering a limit). Instead, we'd like to develop techniques for determining the derivative by easier methods, when possible. To this end, some of you will already be familiar with the following result.

Theorem (The Derivative of $f(x) = ax^n$)**.** *Given a function of the form $f(x) = ax^n$ for some constant $a \in \mathbb{R}$, and $n \in \mathbb{Z}^+$, we have*

$$f'(x) = anx^{n-1}$$

Of course, this result alone doesn't allow us to find the derivative of *any* polynomial, but only those like $f(x) = 3x^2$, with derivative $f'(x) = 6x$, and $f(x) = \frac{2}{5}x^5$, with derivative $f'(x) = 2x^4$. Fortunately, we can easily extend this using the following result (which is rather simple to prove using the definition of the derivative and the relevant properties of limits).

Theorem (The Additive Property of Derivatives)**.** *Given differentiable functions f and g, the function h defined by $h(x) = f(x) + g(x)$ is such that*

$$h'(x) = f'(x) + g'(x)$$

In other words, *the derivative of a sum is the sum of the derivatives.*

Now, using this result, we *can* find the derivative of any polynomial.[2] For example, if $f(x) = 5x^4 - 2x^3 + 7x - 3$, then $f'(x) = 20x^3 - 6x^2 + 7$.

In fact, we have another powerful result at our disposal, as the earlier result concerning the derivative of functions of the form $f(x) = ax^n$ was unnecessarily restricted: we can, in fact, allow n to be *any nonzero real number*. Consider, for example the function $f(x) = \sqrt{x}$. We can reason as follows.

$$\begin{aligned} f(x) &= \sqrt{x} \\ &= x^{\frac{1}{2}}, \text{ and then, taking } n = \frac{1}{2} \text{ in our earlier theorem, we get} \\ f'(x) &= \frac{1}{2}x^{-\frac{1}{2}}, \text{ which could be rewritten as} \\ &= \frac{1}{2\sqrt{x}} \end{aligned}$$

Note that (both for the derivative and the original function), the domain must be restricted to non-negative real numbers.

The result above involved a positive rational exponent, but negative exponents are also permitted. Thus, for example, if $f(x) = \frac{1}{3x}$, then $f'(x) = -\frac{1}{3x^2}$. Many of the important results we'll cover later in the course concern how we can find derivatives for increasingly complicated functions without having to return to the limit definition of the derivative.

2.2 Questions

1. Find the derivative of $f(x) = -2x^3$

   (a) using the technique just introduced for finding derivatives of polynomials.

[2] Strictly speaking we will also be using the result that the derivative of a constant is 0, another result that is quite straightforward to prove.

(b) using the definition of the derivative.

2. Use the method introduced in this section to show that, if $f(x) = \frac{1}{3x}$, then $f'(x) = -\frac{1}{3x^2}$.

3. Find the derivative of $f(x) = x^4 - \frac{1}{2x^2} + \sqrt{x}$.

3 So What Exactly *is* a Derivative?

We can easily find derivatives for a wide variety of functions, but what exactly *is* a derivative? The short (but not very informative) reply is that the derivative of a function is another function, namely, the one given by taking the limit of the Newton quotient (as defined at the beginning of Section 2). Of course, this probably isn't very helpful, and to get a better understanding of derivatives we'll need to consider some of the ways in which we can interpret the derivative.

Consider the gradient of a line between two points (x_1, y_1) and (x_2, y_2). You know that the gradient of the line passing through those two points is given by

$$\frac{y_2 - y_1}{x_2 - x_1}$$

That is, after all, the definition of the gradient. However, you also know that the gradient of a line gives information about its steepness, you know that the line given by $y = 5x - 2$ is steeper than the line given by $y = 2x + 14$, that lines with positive gradients are increasing, lines with negative gradients are decreasing, and you know that you can use any two distinct points on the same line to calculate the gradient and you'll get the same value. These results are obvious to you now, but probably aren't obvious to someone who has just been given the definition of the gradient. We are in a similar situation now with the derivative: we've seen the definition and we can calculate derivatives for a variety of different functions, but we haven't yet discussed what information the derivative provides.

3.1 Derivatives and Tangent Lines

Part of what makes the derivative so useful is the relationship between tangent lines and the derivative. A line that is *tangent* to a curve (in other words, a *tangent line*) is one that, roughly speaking, just touches the curve at a single point.[3]

The important connection between tangent lines and derivatives is expressed below.

Definition (Derivatives and Tangent Lines)**.** Given a differentiable function f and a value a in the domain of f, the gradient of the line tangent to the graph of $y = f(x)$, passing through the point $(a, f(a))$, is equal to $f'(a)$.

[3] This description of a tangent line is not entirely accurate, as you'll see when working through the questions at the end of this section.

In other words, *derivatives give gradients of tangent lines*.[4]

This connection between derivatives and tangent lines can be illustrated by considering the function $f(x) = x^2$. If we graph $y = f(x)$, we get a parabola. Since $f'(x) = 2x$, we can easily calculate that $f'(1) = 2$, $f'(2) = 4$, and $f'(-1) = -2$. What do these values tell us? Well, the points with coordinates $(1,1)$ (with x-coordinate 1), $(2,4)$ (with x-coordinate 2), and $(-1,1)$ (with x-coordinate -1), all lie on the parabola, and for each point there is a a line tangent to the parabola that passes through that point. Calculating the derivatives tells us that

- the tangent line passing through $(1,1)$ has gradient 2, since $f'(1) = 2$,
- the tangent line passing through $(2,4)$ has gradient 4, since $f'(2) = 4$, and
- the tangent line passing through $(-1,1)$ has gradient -2, since $f'(-1) = -2$.

Notice that the calculation of the gradient of the tangent line only involves the x-coordinate of the relevant point on the curve: we can determine, for example, that the tangent line to the curve $y = x^2$ that passes through the point with x-coordinate 2.5 is 5, without having to calculate the y-coordinate of the point on the curve.

With a bit more work, we can also find the *equations* of the tangent lines (and here we do need to find the y-coordinates of the points on the curve). If we are trying to find the equations of each line in the form $y = mx + c$, then we can use the gradient (the value of m) and the given point on the parabola to determine the value of c in each case. For example, the equation of the tangent to the curve $y = x^2$ passing through $(1,1)$ is $y = 2x - 1$.

3.1.1 Questions

1. Find the equation of the tangent line to the curve given by $y = x^2$ at the point on the curve with x-coordinate 2.5.
2. Consider the function $f(x) = x^3 - 2x^2 - x + 2$.
   (a) Find the equation of the tangent line at the point on the curve $y = f(x)$ with $x = 0$.
   (b) Plot the graph of both the original function f and the tangent line from part (a) on the same set of axes.
   (c) Does the tangent line found in part (a) intersect the original curve at only one point?

[4] Note that this 'result' does not simply note a connection between tangent lines and derivatives, it actually gives the precise definition of a tangent line: the tangent line to the curve defined by $y = f(x)$ at a point $(a, f(a))$ is the the line passing through that point with gradient $f'(a)$. Thus, the gradient of the tangent line is *defined* to be equal to the value of the derivative. This allows us to avoid issues with the 'touches the curve at only one point' description of the tangent line—see question 2c in 3.1.1.

3.2 Derivatives and Rates of Change

Calculus is sometimes described as the study of *rates of change*, and the derivative at a point can also be interpreted as an *instantaneous rate of change*. To understand what that means, consider the following example.

A car has been driving for 30 minutes, and has covered a distance of 40 km. If you were asked to calculate the *average* speed of the car for this half-hour journey, you would calculate

$$\frac{40 \text{ km}}{0.5 \text{ h}} = 80 \text{ km/h}$$

Of course, over the 30 minute journey you might expect that the car was occasionally travelling faster than, and occasionally travelling slower than, 80 km/h. However, as far as the *average* speed is concerned, the calculation above is correct: you divide the total distance travelled by the total time taken to get the average speed over that interval—it doesn't matter whether the speed of the car was constant or not.

Let's now consider a similar scenario. Let's assume that the car *is* travelling at a constant speed of 80 km/h down a long, straight track. Along the track we've placed a marker, and we will start a stopwatch when the car passes that marker and then monitor the distance travelled as the car continues along the track. We could then plot a distance-time graph for this situation, as shown in Figure 1. Now consider: how could the speed of the car be determined from this graph? Given that *speed* is the *rate of change of distance with respect to time*, you should recognise that the speed of the car would be represented by *the gradient of the line*, and so calculating the gradient of the line in the distance-time graph would reveal the speed of the car.

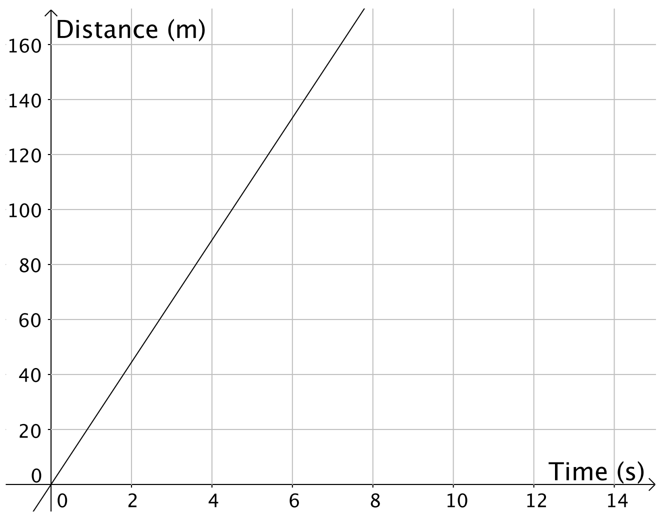

Figure 1: Distance-Time Graph for a Car Travelling at Constant Speed

3.2.1 Accelerated Motion

If cars always travelled at a constant speed we would be able to end our discussion here, but out in the 'real world' cars are constantly changing speed. A car that starts from rest has an initial speed of 0 km/h, but several seconds later that car could very well be travelling at a speed of 80 km/h. If the car is changing speed—if it is accelerating or decelerating—then the distance-time graph will not be a straight line, and so we will not be able to determine the speed of the car by finding a gradient as above.

Consider, then, a more realistic scenario, in which a car starts from rest and accelerates to 80 km/h in 10 seconds, with a distance time graph as shown in Figure 2.

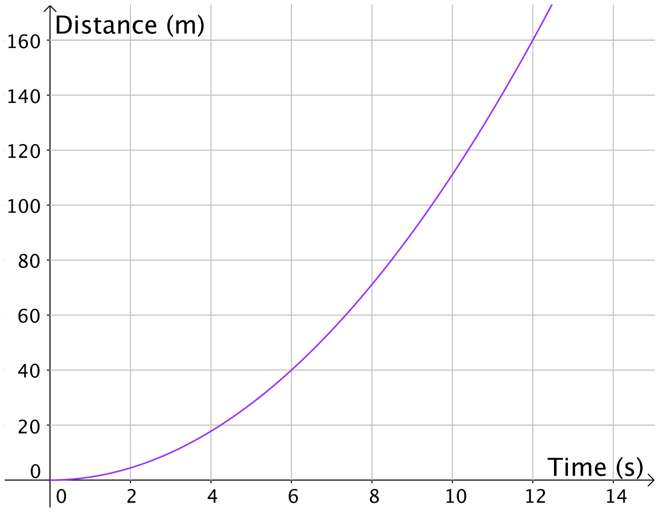

Figure 2: Distance-Time Graph for an Accelerating Car

Here, the function that gives the distance travelled by the car after x seconds is

$$f(x) = \left(\frac{10}{9}\right)x^2$$

Using this function, we can easily verify that the point $(6,40)$ lies on the curve. What was the car's average speed over the first 6 seconds? This is easily shown to be 24 km/h (since the car has travelled 40 m in 6 seconds), and corresponds to the gradient of the line shown in Figure 3.

Similarly, you can calculate that the car's average speed over the next 6 seconds is 72 km/h, which corresponds to the gradient of the line shown in Figure 4.

What none of these calculations provides, though, is the actual speed of the car at *exactly* 6 seconds—if you were to look at the car's speedometer 6 seconds after it started to move, what would you see? You might expect that it would be indicate a speed greater than 24 km/h, but less than 72 km/h. Given what we know about the distance-time graph for the car, is there some way we can calculate its actual speed at exactly 6 seconds?

Knowing what you now know about the derivative, you may suspect that it can be used to answer the question above. Since $f'(x) = \left(\frac{20}{9}\right)x$, we have $f'(6) = \frac{40}{3}$. As you

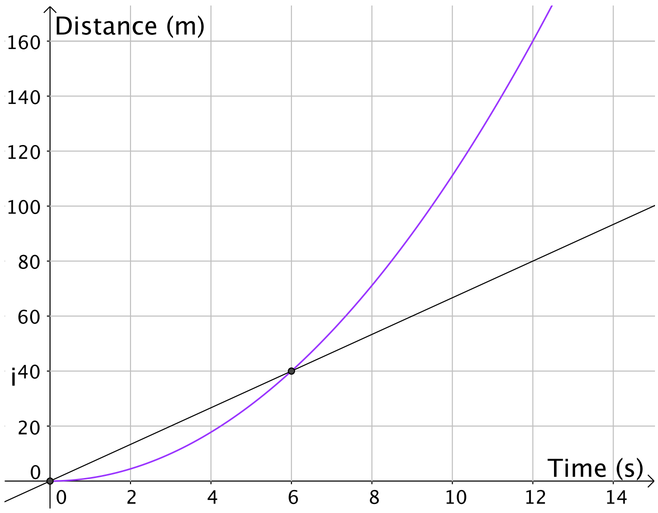

Figure 3: The average speed of the car in the first 6 seconds.

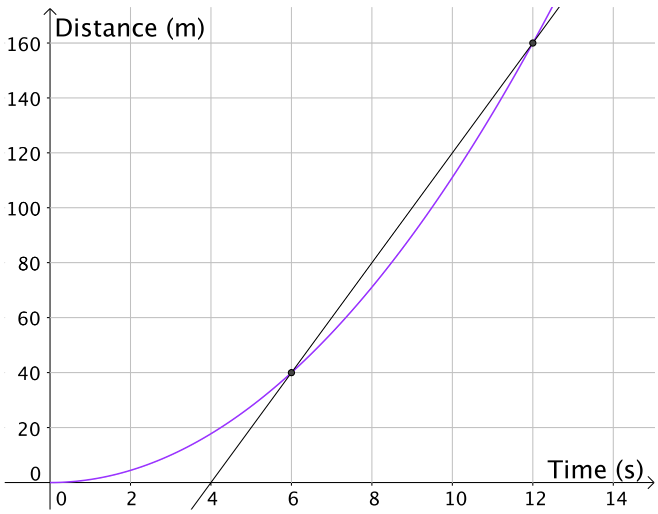

Figure 4: The average speed of the car in the next 6 seconds.

know, this value corresponds to the gradient of the tangent line passing through the point $(6, 40)$, as shown in Figure 5.

The gradient of the tangent line reflects the units of measurement used in the graph, and so the gradient of the tangent line is a value whose units are metres per second. Converting to km/h, we get that the gradient is equivalent to 48 km/h. If you were to look at the car's speedometer exactly 6 seconds after it began to accelerate, you would see that it was travelling at exactly 48 km/h. A moment later it would be travelling at a speed faster than 48 km/h, and a moment earlier it would have been travelling at a speed less than 48 km/h. Thus, since the car is accelerating, its speed is

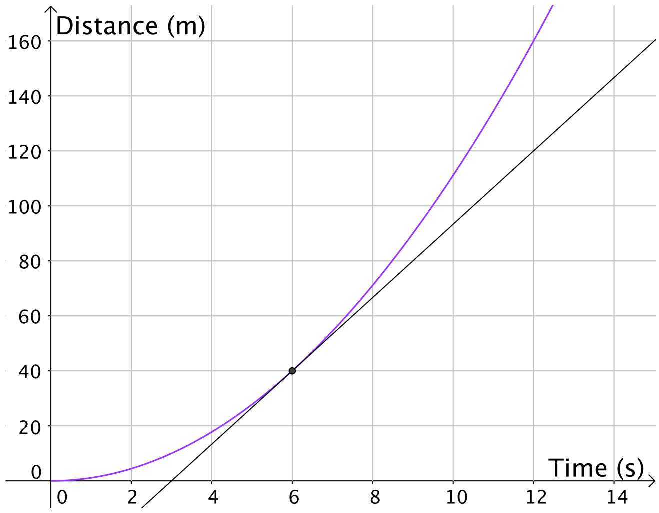

Figure 5: The speed of the car at 6 seconds.

exactly 48 km/h for only an instant, and it is in this sense that the derivative gives us an *instantaneous rate of change*.

3.2.2 Questions

All questions below concern the accelerating car from Section 3.2.1.

1. (a) Show that the car is travelling at 80 km/h 10 seconds after it begins to move.
   (b) Calculate the average speed of the car over the first 10 seconds.
   (c) Find the time at which the car is travelling at exactly the speed you calculated in question 1b.

2. (a) Find the derivative of the function

$$f'(x) = \left(\frac{20}{9}\right)x$$

   This new function is a called *second* derivative.
   (b) Explain why the units of the second derivative are ms^{-2}. What is the acceleration of the car?

4 Derivatives and Curve Sketching

It is sometimes the case in applications of calculus that we know more about the derivative of a function than we do about the original function. It is useful, then, to consider exactly what we can discover about a potentially unknown original function through the study of its derivative.

Here, we'll first start with the reverse situation, beginning with a known function and exploring the features of its derivative. Once we've identified how the features of the original function affect the features of its derivative, we'll study the derivative of a second, unknown function, to see what we can determine about the original function.

Again, a function may fail to have a derivative at some (or all) points in its domain, but we'll assume that all functions we consider have domain $\mathbb{R}$, and have a derivative defined for all values in $\mathbb{R}$ (unless explicitly indicated otherwise).

4.1 Intervals of Increase and Decrease

Assume that f is a function defined on the interval $[a,b]$ for $a,b \in \mathbb{R}$, f is said to be *strictly increasing* on the interval $[a,b]$ provided, for any $x,y \in [a,b]$

$$\text{if } x < y \text{ then } f(x) < f(y)$$

In other words, a function is strictly increasing over an interval if *the greater the input, the greater the output*.

Similarly, if f is a function defined on the interval $[a,b]$ for $a,b \in \mathbb{R}$, f is said to be *strictly decreasing* on the interval $[a,b]$ provided, for any $x,y \in [a,b]$

$$\text{if } x < y \text{ then } f(y) < f(x)$$

For example, the function $f(x) = x^2$ with domain $\mathbb{R}$ is increasing on $[1,4]$ and decreasing on $]-20,-3[$ (though the definitions involve closed intervals, that condition is not essential). Of course, it would be more informative to say the function is increasing on $[0,\infty[$ and decreasing on $]-\infty,0]$, as those are the largest intervals over which the two conditions are satisfied.[5]

Sometimes a weaker condition than *strict* increase/decrease is all that's required, which leads to the definitions below.

If f is a function defined on the interval $[a,b]$ for $a,b \in \mathbb{R}$, f is said to be *non-decreasing* on the interval $[a,b]$ provided, for any $x,y \in [a,b]$

$$\text{if } x < y \text{ then } f(x) \leq f(y)$$

In other words, a function is non-decreasing on an interval if it is either increasing or 'holding steady' as the input values increase.

Similarly, if f is a function defined on the interval $[a,b]$ for $a,b \in \mathbb{R}$, f is said to be *non-increasing* on the interval $[a,b]$ provided, for any $x,y \in [a,b]$

$$\text{if } x < y \text{ then } f(y) \leq f(x)$$

Unfortunately there are some alternative terminologies in use for these notions: sometimes what we have here called *strictly* increasing/decreasing is instead called increasing/decreasing, in which case 'non-decreasing' may be called *monotonically* increasing (and 'non-increasing' may be called *monotonically* decreasing). Even worse, sometimes a mixture of terminology is used. Consequently, speaking of an *increasing*

[5] Notice that 0 can be included in both intervals—think about why this is the case.

function, for example, could be ambiguous—do you mean *strictly* increasing, or *monotonically* increasing? The advantage of the terminology were are using here is that our terminology is unambiguous.

So, what does all of this have to do with derivatives? Consider the graph of the function $f(x) = x^3 - x^2 - x$, with domain $\mathbb{R}$, along with the tangent lines at the three zeros of this function, as shown in Figure 6.

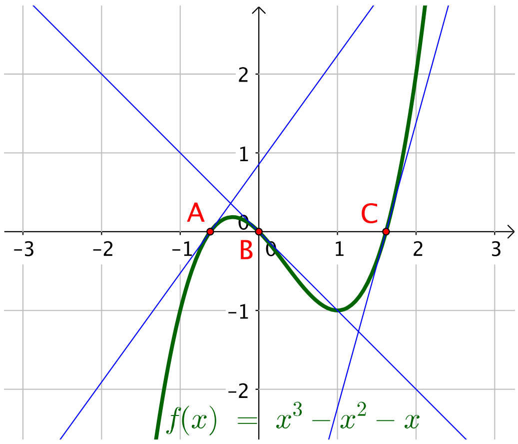

Figure 6: The graph of $f(x) = x^3 - x^2 - x$, with tangents at the zeros.

Notice that both A and C are located in intervals in which f is increasing, and B is located in an interval in which f is decreasing. How is this reflected in the gradients of the various tangent lines? In other words, how is this reflected in the value of *the derivative* at each point?

As you may now have realised, the *sign* of the derivative at a point helps to indicate where the function is increasing or decreasing.

- A point at which the derivative is *positive* has a tangent line with positive gradient, and so will lie in an interval in which the function is *increasing*.
- A point at which the derivative is *negative* will have a tangent line with negative gradient, and so will lie in an interval in which the function is *decreasing*.

Thus, *the sign of the derivative helps to identify the intervals of increase/decrease*: if f' is positive over some interval, then f is increasing over that interval, if f' is negative over some interval, then f is decreasing over that interval.

4.2 Stationary Points: Extrema

As you consider the values of the derivative of the function $f(x) = x^3 - x^2 - x$, you will recognise that there is a third option for the derivative: it is positive at certain values, negative at others, but it is also 0 at two points, as shown in Figure 7. This leads to the following definition.

Definition (Stationary Points)**.** Given a differentiable function f and a value a in the domain of f, a is a *stationary point* of f if $f'(a) = 0$.

As we know that the value of $f'(a)$ gives the gradient of the tangent line to the graph of $y = f(x)$ at the point with x-coordinate a, stationary points of a function f indicate where the tangent line is horizontal.

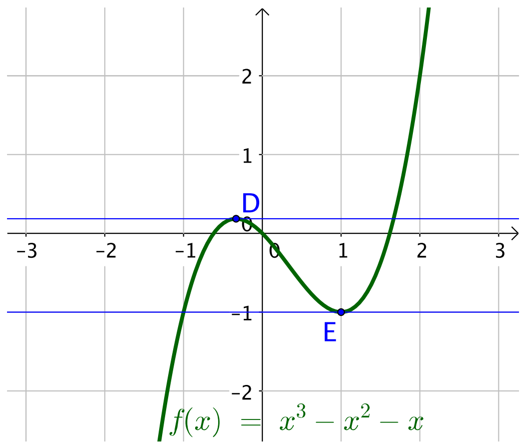

Figure 7: The graph of $f(x) = x^3 - x^2 - x$, with tangents of gradient zero shown.

The point E occurs at what is called a *local minimum* of the function f. It is a *local* minimum because -1 is *not* a minimum value of f (notice that f has no *global* minimum value on $\mathbb{R}$), but -1 *is* a minimum of the function in some intervals centred at 1, for example, in the interval $[0, 2]$. Similarly, f has a *local maximum* at D, though f has no (global) maximum value on its domain. Maxima and minima (whether local or global) are collectively called *extrema* (singular: *extremum*), although they may also be called *turning points*. Thus, for example, points D and E are turning points of f. This example should serve to suggest the following (correct) general principle: *the extrema of a differentiable function with domain $\mathbb{R}$ occur at stationary points.*[6]

We are now in a position to be able to easily calculate the (exact) coordinates of D. To do so, we'll (finally!) consider f' and its graph, as shown in Figure 8.

A quick derivation confirms that $f'(x) = 3x^2 - 2x - 1$, and this quadratic can be factored to give

$$f'(x) = (3x + 1)(x - 1)$$

[6]We need to be a bit careful in applying this result, since functions with domains restricted to an interval may have extrema at the endpoints of that interval. For example, the function $f(x) = x^2$ with domain $\mathbb{R}$ has a (global) minimum at $x = 0$ and no global maximum, but the function $f(x) = x^2$ with domain $[1, 4]$ has a global minimum at $x = 1$ and a global maximum at $x = 4$.

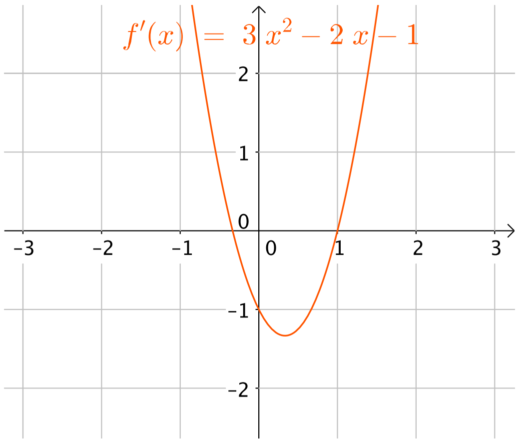

Figure 8: The graph of $f'(x) = 3x^2 - 2x - 1$.

Thus, we see that f' has two zeros, and so f has two stationary points: one at $x = -\frac{1}{3}$ and one at $x = 1$. From the graph of f (see Figure 7) we know that E has x-coordinate 1, and so $x = -\frac{1}{3}$ gives the x-coordinate of D. Finally,

$$\begin{aligned} f\left(-\frac{1}{3}\right) &= \left(-\frac{1}{3}\right)^3 - \left(-\frac{1}{3}\right)^2 - \left(-\frac{1}{3}\right) \\ &= \frac{5}{27} \end{aligned}$$

So, the coordinates of D are $\left(-\frac{1}{3}, \frac{5}{27}\right)$.

Putting all these pieces together, we can now identify a number of the features of f reflected in the graph of f' (Figure 8). First, we can see that f has two stationary points, since f' has two zeros. We already knew that the leftmost stationary point, at $x = -\frac{1}{3}$, was a local maximum, but we could also have determined this from the graph of f'. Since f' is *positive* to the left of $x = -\frac{1}{3}$, we know that f is *increasing* to the left of $x = -\frac{1}{3}$. Since f' is *negative* to the right of $x = -\frac{1}{3}$, we know that f is *decreasing* to the right of $x = -\frac{1}{3}$. Thus, f is increasing to the left of $x = -\frac{1}{3}$, stationary at $x = -\frac{1}{3}$, and decreasing to the right of $x = -\frac{1}{3}$, so the point with $x = -\frac{1}{3}$ on the graph of f must be a maximum. Similarly, we see from the graph of f' that f is decreasing to the left of $x = 1$, stationary at $x = 1$, and increasing to the right of $x = 1$, so the point with $x = 1$ on the graph of f must be a minimum.

The reasoning used here to identify minima and maxima can be applied in other situations to determine whether or not a stationary point is a maximum or a minimum. As shown in Figure 9, a stationary point that is a minimum will have a derivative that is negative immediately to the left of that point, and positive immediately to the right. Similarly, a stationary point that is a maximum will have a derivative that is positive immediately to the left of that point, and negative immediately to the right.

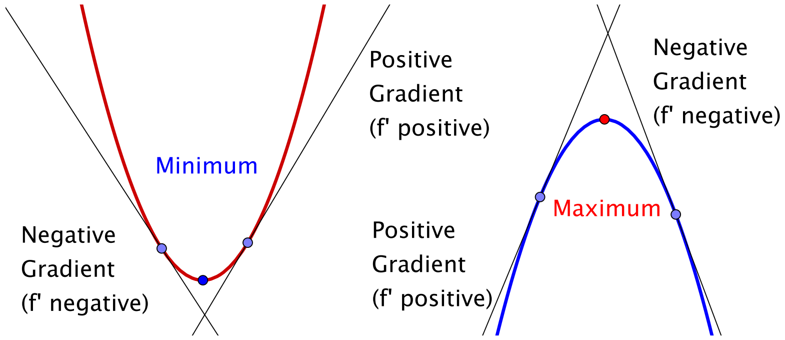

Figure 9: Tangents, Derivatives, and Extrema

Let's quickly look at a second example before we apply the methods we've just developed in a more complicated scenario. Consider the function $f(x) = x^2 + 2x$. Its derivative is $f'(x) = 2x + 2$, and so we can see immediately that f' has a zero, and so f has a stationary point, at $x = -1$.

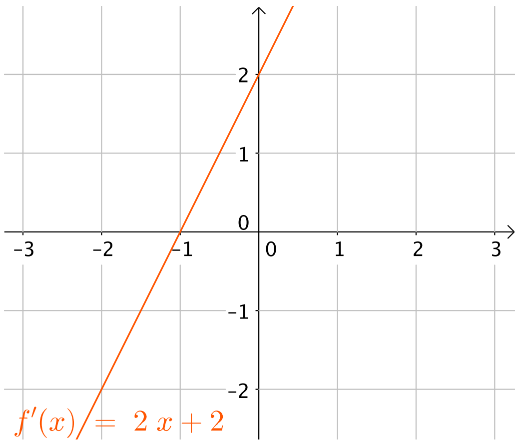

Figure 10: The graph of $f'(x) = 2x + 2$.

It is also clear that this stationary point is a minimum, since f' is negative to the left of the stationary point and positive to the right (which indicates that f is decreasing to the left of the stationary point, then increasing to the right).

Of course, you know quite a bit about quadratic functions, and so you were already able to recognise that f would be concave up, and so have a (global) minimum value.

4.2.1 Example

Let's now consider our final example. The unknown function f has derivative f', as shown in Figure 11.

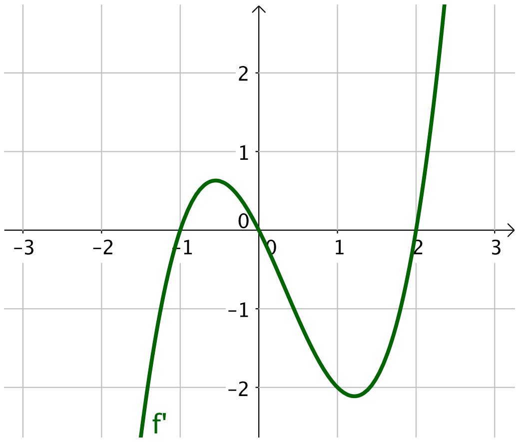

Figure 11: The graph of $y = f'(x)$.

What characteristics of f can we determine from the graph of f'? Well, clearly f has three stationary points, at $x = -1$, $x = 0$, and $x = 2$. It is also apparent that the stationary points at $x = -1$ and $x = 2$ are minima, and the stationary point at $x = 0$ is a maximum.

Given this information, there are a number of possible curves that will satisfy the given constraints, and so we can't expect our curve to match the original curve exactly (though it should have roughly the same shape). I suggest you complete your sketch of a possible graph of the original function f, then compare your sketch to that shown in Figure 13 on page 16.

4.3 Questions

1. Consider a quadratic function $f(x) = ax^2 + bx + c$, for some constants $a, b, c \in \mathbb{R}$, with $a \neq 0$.

   (a) Find an expression for $f'(x)$.

   (b) Find an expression for the x-coordinate of the stationary point of f.
   Hence, explain why the equation of the axis of symmetry for the graph of a quadratic function is $x = -\frac{b}{2a}$.

   (c) Use the derivative to explain how the sign of a determines whether the quadratic function f has a (global) minimum or a (global) maximum.

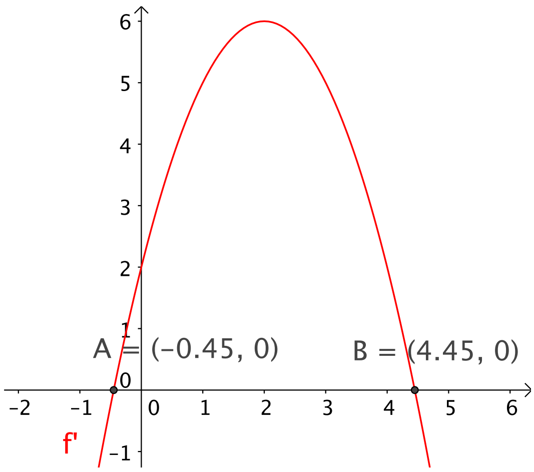

Figure 12: The graph of $y = f'(x)$.

2. Consider the graph of $y = f'(x)$, for a differentiable function f with domain $\mathbb{R}$, shown in Figure 12.
   (a) Identify the intervals of increase and decrease of f.
   (b) Sketch one possible graph of the original function f.

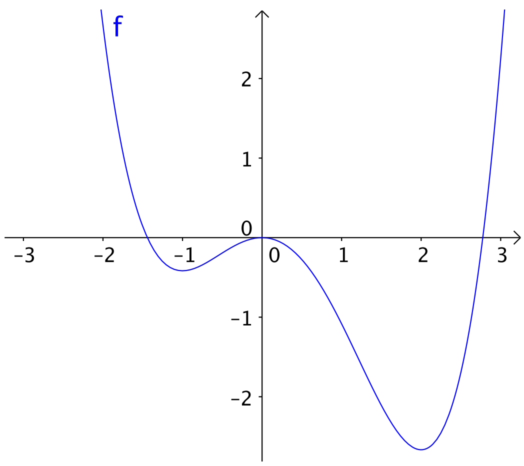

Figure 13: One possible graph of $y = f(x)$ from the example in Section 4.2.1.

5 Second Derivatives

5.1 Concavity

You now know that, for a given differentiable function f with domain $\mathbb{R}$, you could identify the intervals of increase and decrease from the graph of f': f will be increasing over intervals in which f' is positive, and decreasing over intervals when f' is negative. Of course, f' itself is a function, and so we could also study its intervals of increase or decrease. What do the intervals of increase or decrease for f' reveal about the original function f?

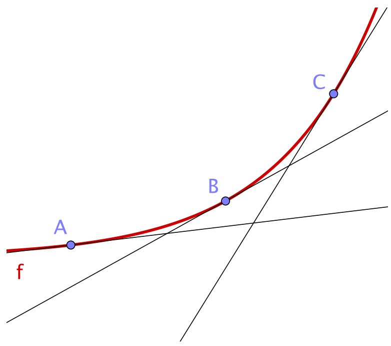

Figure 14: The graph of $y = f(x)$.

Consider the function f shown in Figure 14, with tangent lines indicated at A, B, and C. Consider the gradients of those tangent lines—as we move further to the right, the gradients increase, and so f' is an increasing function. Graphs of functions that bend upwards are said to be *concave up*, and so a function whose graph is concave up will have a derivative that is increasing.

Notice that f in Figure 14 is an increasing function, *and* its derivative is increasing (so f is an increasing function that is concave up). Is it possible for a *decreasing* function to be concave up? In other words, can a decreasing function have an increasing derivative? This question can be answered by considering the function g shown in Figure 15. While the function g is decreasing (and so its derivative is negative), the

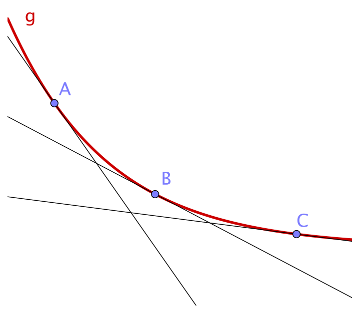

Figure 15: The graph of $y = g(x)$.

derivative of g is increasing. Thus, g is a decreasing function that is concave up.

A function whose derivative is decreasing is said to be *concave down*, and (as with functions that are concave up), a function that is concave down can either be increasing or decreasing. You should draw a quick sketch of a function that is increasing and concave down, and another of a function that is decreasing and concave down.

Let's consider the concavity of $f(x) = x^2$ with domain $\mathbb{R}$. Without further calculation, you should recognise that f is concave up over its entire domain.[7] Taking the derivative of f gives us $f'(x) = 2x$, which is clearly an increasing function. Thus, consideration of f' confirms that f is concave up, as expected. Notice, though, that we have another way to establish that f' is increasing—from Section 4.1, we know that a function is increasing when its derivative is positive, and so we can find *the derivative* of f' to check for intervals of increase or decrease. Thus, we want to find the *second* derivative[8] of f, which is symbolised in Lagrange notation as f''. The calculation here is simple, given $f'(x) = 2x$, we get

$$f''(x) = 2$$

Thus f'' is positive for all values in $\mathbb{R}$, and so f' is increasing throughout its domain, and thus f is concave up over $\mathbb{R}$.

Let's now consider another simple example, with $f(x) = -2x^2 + 3x - 4$ and domain $\mathbb{R}$. Here we get $f'(x) = -4x + 3$, and so $f''(x) = -4$. Clearly f'' is negative throughout its domain, and so f' is decreasing, which tells us that f is concave down over its entire domain.

5.1.1 Questions

1. Consider the function $f(x) = \sqrt{x}$ with domain $[0, \infty[$. Find the second derivative of f. **Hence**, show that f is concave down on its domain.

2. Show that $f(x) = \frac{1}{x}$, with domain $\mathbb{R}$, is concave down on $]0, \infty[$ and concave up on $]-\infty, 0[$.

5.2 Points of Inflexion

Quadratic functions with domain have rather unexciting behaviour when it comes to concavity: they're either concave down or concave up throughout their domain. It's not until we get to degree three polynomial functions that we our first examples of functions that *change concavity* over their domain, and a *point of inflexion* is the name given to a point at which a function changes concavity.

Let's study the function $f(x) = x^3 - x^2 - x$ (which you will recall had also been considered in Section 4.2). Quick calculations yield that $f'(x) = 3x^2 - 2x - 1$ and $f''(x) = 6x - 2$. The graphs of all three functions are shown in Figure 16.

It is clear that f'' has a zero at $x = \frac{1}{3}$, is negative on the interval $]-\infty, \frac{1}{3}[$, and is positive on the interval $]\frac{1}{3}, \infty[$. Thus, consideration of f'' confirms that f' is decreasing on $]-\infty, \frac{1}{3}[$ and increasing on $]\frac{1}{3}, \infty[$. This indicates that f is concave down on $]-\infty, \frac{1}{3}[$ and

[7] Just as a single function can be increasing in some intervals and decreasing on others, a function can also have intervals of different concavity.

[8] Recall question 2a in Section 3.2.2.

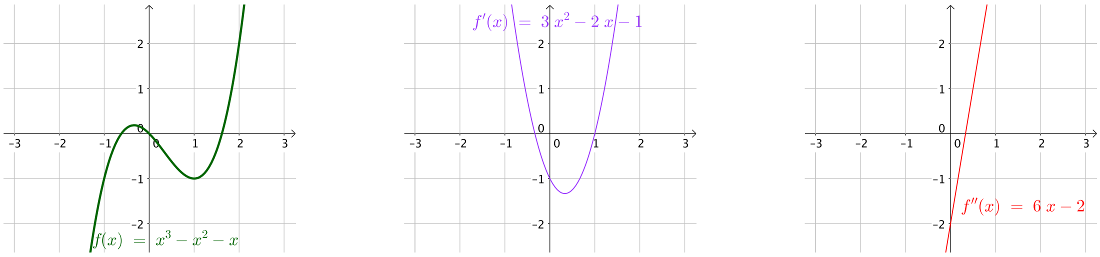

Figure 16: The graph of $f(x) = x^3 - x^2 - x$, along with its first and second derivatives.

concave up on $]\frac{1}{3}, \infty[$. Thus we see that f *changes concavity* at $x = \frac{1}{3}$, and so f has a point of inflexion. Substituting $\frac{1}{3}$ into the expression for f establishes that the coordinates of the point of inflexion are $\left(\frac{1}{3}, -\frac{11}{27}\right)$, and this point is indicated in Figure 17.

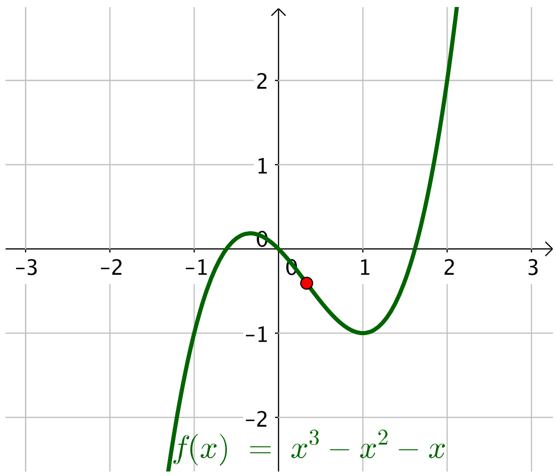

Figure 17: The graph of $f(x) = x^3 - x^2 - x$, showing its point of inflexion.

These considerations leads us to the following definition.

Definition (Point of Inflexion)**.** Given a twice-differentiable[9] function f and a point a in the domain of f, a is a point of inflexion of f provided $f''(a) = 0$ and f'' changes sign at a.

It is important to note that, given a twice-differentiable function f, knowing that $f''(a) = 0$ is *not sufficient* to guarantee that a is a point of inflexion—we need to be

[9]In other words, both the first and second derivatives of f exist.

certain that f'' changes signs at a as well. To see why, consider the function $f(x) = x^4$, with second derivative $f''(x) = 12x^2$. Clearly $f''(0) = 0$, but $(0,0)$ is *not* a point of inflexion. The function f'' does not change signs at 0, and f is concave up on all of $\mathbb{R}$.

5.2.1 Questions

1. Find the point of inflexion of the function $f(x) = -x^3 + 4x^2 - 7x + 2$, and so determine the intervals over which the function is concave up/concave down.

5.2.2 Stationary Points Revisited

Our consideration of $f(x) = x^4$ showed us that zeros of f'' don't necessarily yield points of inflexion—we need to be certain that f'' actually changes sign at the point in question as well. This raises an interesting question: do we need to worry about something similar happening when we use the zeros of the *first* derivative (i.e., the stationary points) to find extrema? We've seen that minima and maxima of differentiable functions with domain $\mathbb{R}$ occur at stationary points, but is it possible to have a stationary point that is *not* an extremum? Interestingly, the answer is *yes*.

Consider the function $f(x) = x^3$. As $f'(x) = x^2$, we immediately get that $f'(0) = 0$, and so 0 is a stationary point of f. However, (as the graph will confirm), $(0,0)$ is neither a maximum nor a minimum of f. What is happening at the point $(0,0)$? You should now be able to recognise that $(0,0)$ is a point of inflexion, which gives us our third (and final) option: given a twice-differential function f, if a is a stationary point of f then

- $(a, f(a))$ is a maximum, or
- $(a, f(a))$ is a minimum, or
- $(a, f(a))$ is a point of inflexion.

This (now completed) list will be explored further in your assignment as you develop the *Second Derivative Test*.

5.2.3 Questions

1. Consider the function $f(x) = x^3$. Plot the graph of f', and with reference to your graph, explain how you can tell that $(0,0)$ is not an extremum of f.

   [Hint: If 0 were, for example, a maximum of f, then f' would positive just to the right of 0 and negative just to the left. What do you actually observe on either side of $x = 0$ in the graph of f'?]