2: Limits and Continuity

We'll begin with a motivating idea—movement—since that is one of the driving ideas that forced the development of The Calculus.

2.1: Rates of Change and Limits

A Big Idea

Imagine for a moment that you are riding in a car with me, and the speedometer reads 113 kph. What does that mean?

Perhaps you would say "for every hour of travel, you'll go 113 km." In other words, it is the change in position (113 km) for a change in time (1 hour).

I only looked at the speedometer for a second—how can it know how far I'll go in an hour, based on a moment of time? Maybe there will be a roadblock ahead, or maybe I'll go crazy and exceed the speed limit…

You feel that you understand speed (or, maybe you did, up until a few minutes ago), but a little closer inspection will make you realize that you've been using this idea without really knowing what it was. Relax; it's okay. I'm here to help.

The problem is that there are two different ideas running around here: an *average speed* (measured over some interval of time), and *instantaneous speed*. It is precisely this idea of instantaneous speed that spurred The Calculus—finding a watertight explanation of what it is took over a hundred years of work, and changed mathematics from a necessary side skill for scientists into a discipline of its own.

A quick note: I'm going to use the terms speed and velocity interchangeably in this chapter. I know that they're different—it's not going to matter right now, and I don't want the words to get in the way.

Average speed can be defined as $\dfrac{\text{change in position}}{\text{change in time}} = \dfrac{p_{final} - p_{initial}}{t_{final} - t_{initial}}$. This formula works just fine, as long as some interval of time has elapsed. The word "instantaneous" should suggest that we're looking at reducing that interval of time to zero. Unfortunately, the formula above is not conducive to this endeavor—we can't divide by zero!

What we *can* do is investigate the behavior of the formula for time intervals that are *close* to zero. You did this before (should have, anyway) when you dealt with rational functions in PreCalculus. We're going to do it again, but this time we'll call the procedure…

Limits

We'll step away from speed for a bit…

Consider the function $y = \dfrac{x^2 - 1}{x + 1}$. Let's factor this to notice something interesting:

$y = \dfrac{x^2 - 1}{x + 1} = \dfrac{(x+1)(x-1)}{x+1}$. You might be tempted to cancel out that common factor, and declare that this function is really $y = x - 1$—but you'd be wrong. The domain of $y = x - 1$ is *all* real numbers; the domain of $y=\frac{x^2-1}{x+1}$ is all real numbers *except -1*. The graphs of these two functions are identical, except for one tiny little point.

Of course, you remember this as a *hole*, or a *removable discontinuity*. We're going to use this to talk about limits.

You can't plug $x=-1$ into $y=\frac{x^2-1}{x+1}$, but we can investigate what happens to the *y*-values as *x* gets close to -1. This is called a **limit**. Here's the notation: $\lim_{x\to -1}\frac{x^2-1}{x+1}$. This is read as "the limit, as *x* approaches negative one, of *x*-squared minus one over *x* plus one." The idea is that you are approaching the indicated value, but never quite reaching it.

If $x\neq -1$, then can you reduce $y=\frac{x^2-1}{x+1}$ to $y=x-1$?

Yes! As long as $x\neq -1$, the common factor $x+1$ isn't equal to zero. As long as $x\neq -1$, the domain of $y=\frac{x^2-1}{x+1}$ is the same as the domain of $y=x-1$. Thus, you *can* cancel that factor, and the two expressions are identical.

The notation: $\lim_{x\to -1}\frac{x^2-1}{x+1}=\lim_{x\to -1}\frac{(x+1)(x-1)}{x+1}=\lim_{x\to -1}(x-1)$.

Now we can ask "What happens to the values of $y=x-1$ as *x* approaches -1?" Remember: as long as $x\neq -1$, the graph of $y=x-1$ is going to behave identically to $y=\frac{x^2-1}{x+1}$.

Well, when $x=-1$, $y=x-1=(-1)-1=-2$.

Okay, okay! I said $x\neq -1$. But here's the deal: as long as plugging in that value you're approaching doesn't cause division by zero (or other things; more later), then *go ahead*. That's the beauty of limits: you avoid the problem value ($x=-1$) just long enough to eliminate it (*remove the discontinuity*), then you start using it.

Thus, we finish the limit: $\lim_{x\to -1}\frac{x^2-1}{x+1}=-2$.

When the values of *x* get really close (*arbitrarily close*) to negative one, the values of $y=\frac{x^2-1}{x+1}$ get close (but never equal) to -2.

Here are the technical details of what we just did: The limit of *f*(*x*) as *x* approaches *a* is *L* ($\lim_{x\to a}f(x)=L$) means that for every positive real value ε, there exists some positive real value δ, so that if $|x-a|<\delta$, then $|f(x)-L|<\varepsilon$.

Whew! What this is really saying is "if you tell me how close to *L* you want *f*(*x*) to get, then I can tell you how close to *a* *x* needs to be." This is a rigorous definition of what mathematicians mean when they say *arbitrarily close*—however close you want to get, I can tell you what you need to plug in to the function to get a function value that is that close.

Now relax—I just wanted you to see it. We're not going to mess with that anymore.

I would like to point out that finding limits graphically is usually a good idea. You'll see some examples of this in class, and farther along in these notes. Also, finding limits from (with) a table of values works.

Properties of Limits

The limit of a constant is that constant. If b is a real number, then the limit as x approaches any value a is $\lim_{x \to a} b = b$.

If a function is defined for $x = a$, then (usually) you can simply plug into the function to find the limit. $\lim_{x \to a} x = a$, $\lim_{x \to a} x^n = a^n$, $\lim_{x \to a} \sqrt{x} = \sqrt{a}$, etc. The main exceptions are piecewise-defined functions. We'll see some technical details about this later.

For the remaining properties, assume that $\lim_{x \to a} f(x) = L$, $\lim_{x \to a} g(x) = K$, and b is a real number.

Constants can be factored out: $\lim_{x \to a} \left[b \cdot f(x) \right] = b \cdot \lim_{x \to a} \left[f(x) \right] = b \cdot L$

Operations can be distributed: $\lim_{x \to a} \left[f(x) \pm g(x) \right] = \lim_{x \to a} f(x) \pm \lim_{x \to a} g(x) = L \pm K$,

$\lim_{x \to a} \left[f(x) \cdot g(x) \right] = \lim_{x \to a} f(x) \cdot \lim_{x \to a} g(x) = L \cdot K$, $\lim_{x \to a} \frac{f(x)}{g(x)} = \frac{\lim_{x \to a} f(x)}{\lim_{x \to a} g(x)} = \frac{L}{K}$ (provided that $K \neq 0$), and $\lim_{x \to a} \left[f(x) \right]^n = L^n$.

There is little need to memorize these—you'll find that the second paragraph above, plus some examples, will get you where you need to go.

Note: not every limit has an answer! If there is no way to eliminate division by zero, then the limit is some type of infinity (some people would say that the limit doesn't exist). There are others where everyone says that the limit simply doesn't exist…more on the weird ones later.

Examples

[1.] Find $\lim_{x \to 2} \frac{x^2 + 3x - 5}{x + 2}$.

Since letting $x = 2$ doesn't cause division by zero, I can just plug in:

$$\lim_{x \to 2} \frac{x^2 + 3x - 5}{x + 2} = \frac{(2)^2 + 3(2) - 5}{(2) + 2} = \frac{4 + 6 - 5}{4} = \frac{5}{4}.$$

[2.] Find $\lim_{x \to 1} \frac{x^3 - 1}{x - 1}$.

I can't let x be one…perhaps I'll try to make the problem go away by factoring (since that worked in the very first example above).

$$\lim_{x \to 1} \frac{x^3 - 1}{x - 1} = \lim_{x \to 1} \frac{(x - 1)(x^2 + x + 1)}{x - 1}$$

(*you do remember how to factor differences of cubes, don't you?*)

A-ha! Since I'm taking the limit as x approaches one, x is not equal to one, so I can cancel that pesky denominator!

$$\lim_{x\to1}\frac{(x-1)(x^2+x+1)}{x-1}=\lim_{x\to1}(x^2+x+1)$$

Now it is okay to let x be one.

$$\lim_{x\to1}(x^2+x+1)=1^2+1+1=3.$$

Let's look at the graph:

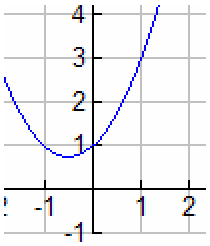

It looks like the graph actually passes through that point…however, closer inspection on your calculator will reveal the truth! This viewing window is approximately $[0.9,1.1][2.9,3.1]$.

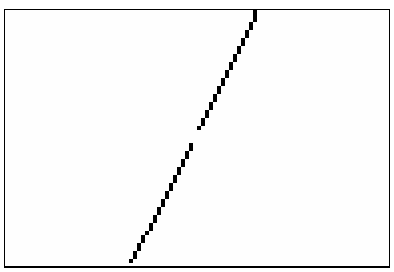

Finally, let's look at a table:

x	0.9	0.99	0.999	1.001	1.01	1.1
$\frac{x^3-1}{x-1}$	2.71	2.9701	2.997001	3.003001	3.0301	3.31

[3.] Find $\lim_{x\to0}\frac{g(3+x)-g(3)}{x}$ if $g(x)=x^2$.

Whew! Let's plug in, first. The problem is $\lim_{x\to0}\frac{(3+x)^2-(3)^2}{x}$. You can't let x be zero—let's FOIL that thing out and see what happens.

$$\lim_{x\to0}\frac{(3+x)^2-(3)^2}{x}=\lim_{x\to0}\frac{(3^2+6x+x^2)-(3^2)}{x}$$

I see something pretty! Cancel it!

$$\lim_{h\to0}\frac{(3^2+6x+x^2)-(3^2)}{x}=\lim_{x\to0}\frac{6x+x^2}{x}$$

Whoa! I see something else!

$$\lim_{x\to0}\frac{6x+x^2}{x}=\lim_{x\to0}\frac{x(6+x)}{x}=\lim_{x\to0}(6+x)$$

Now I can let x be zero. The limit is 6.

[4.] Find $\lim_{x\to 0}\frac{1}{x^2}$.

I can't let x be zero…and there isn't anything I can do algebraically. This must result in infinity—specifically, positive infinity (x^2 forces the fraction to always be positive). Of course, you could look at the graph:

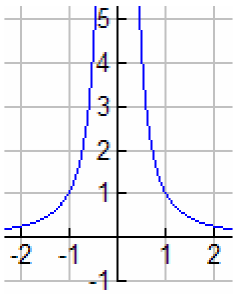

It's easy to tell that as x gets close to zero, the y-values are not getting close to any particular value.

Sided Limits

We have thus far considered limits from both sides of the point, but that's not the only way to do it! Consider the greatest integer function, $y=\lfloor x\rfloor$:

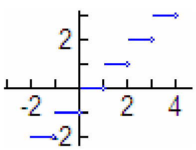

If you want to find the limit of this function as x approaches 2, you must consider both sides. For values of x that are lower than 2, the function limit is 1. However, for values of x that are larger than 2, the limit is 2.

If you are approaching from above (from the right), the notation is $\lim_{x\to a^+} f(x)$. If you are approaching from below (from the left), the notation is $\lim_{x\to a^-} f(x)$.

Technical detail: $\lim_{x\to a} f(x)$ exists only if $\lim_{x\to a^-} f(x)=\lim_{x\to a^+} f(x)$. In other words, you only get to talk about *the limit* if *the limit from the left equals the limit from the right*. So, since $\lim_{x\to 2^-}\lfloor x\rfloor=1\neq 2=\lim_{x\to 2^+}\lfloor x\rfloor$, $\lim_{x\to 2}\lfloor x\rfloor$ does not exist.

Typically, the sided limits *will* agree—the exceptions are few and far between. Typically, those exceptions are piecewise-defined functions, specifically designed to have "jumps" in the graph.

The Sandwich/Squeeze Theorem

Sometimes, we can't find limits directly, and must instead rely on stealth and cunning…one such method is the **Sandwich Theorem**. If you've got a recalcitrant function whose limit you desire (call it $\lim_{x\to a} f(x)$), and you can find two nice functions that bound the naughty function in the area where you're working the limit ($l(x)\le f(x)\le h(x)$ for some small interval surrounding

a), *and* those functions *just happen* to have the same limit at a ($\lim_{x \to a} l(x) = \lim_{x \to a} h(x) = L$), *then* $f(x)$ gets "sandwiched" between $l(x)$ and $h(x)$: $\lim_{x \to a} f(x) = L$.

Whew! This is a bit technical, and we won't dwell on it…if you become a math major, then you'll see this again.

2.2: Limits Involving Infinity

Limits When Approaching Infinity

The preceding section concerned limits at real numbers…but that's not the only way to do it! You can also consider the limit of a function as *x* approaches infinity (or negative infinity).

As was the case with the preceding limits, you've done this before! The limits in §2.1 came from finding holes in rational functions. Limits at infinity come from…(*drumroll*)…the end behavior of rational functions (and polynomials).

As an example, consider $f(x) = \frac{x^2 + 2x + 1}{x^2 - 2}$. What is its end behavior?

That's correct—it has a horizontal asymptote at $y = 1$. In other words, on the left, it approaches one, and on the right it approaches one.

Here is the same information, with notation: $\lim_{x \to -\infty} f(x) = 1$ and $\lim_{k \to \infty} f(x) = 1$.

Limits Resulting in Infinity

Must the end behavior be a number? Of course not! Consider $g(x) = \frac{x^3 + 1}{x}$. What is its end behavior? What are its limits at positive and negative infinity?

$\lim_{x \to -\infty} g(x) = \infty$ and $\lim_{x \to \infty} g(x) = \infty$

Limiting Functions

We can be a bit more descriptive about the end behavior of a rational function…in particular, we can find **end behavior models** (with any luck, you did this in PreCalculus). Consider again $g(x) = \frac{x^3 + 1}{x}$. Let's rewrite this to see its end behavior—in particular, let's divide it out.

$\frac{x^3 + 1}{x} = \frac{x^3}{x} + \frac{1}{x} = x^2 + \frac{1}{x}$. As *x* gets really large (far from zero), the x^2 part dominates the expression, because the $\frac{1}{x}$ part gets really close to zero. Thus, the end behavior model for $g(x) = \frac{x^3 + 1}{x}$ is $y = x^2$.

Try another: $y=\dfrac{x^2+2x+1}{x-2}$. To divide this one, you can work with polynomial long division, or synthetic division. In either case, it works out to $x+4+\dfrac{9}{x-2}$ (that's a quotient of $x+4$ and a remainder of 9). When x wanders far from zero, the remainder part of that expression gets close to zero, leaving you with the limiting function $y=x+4$.

2.3: Continuity

Perhaps you've heard about continuity before…"a function is continuous if you can draw it without picking up your pencil," or something like that. That's not a terrible idea, but it's not very technical, either. Time to fix that…

Continuity at a Point

Definition: A function $f(x)$ is continuous at $x=a$ if $\lim\limits_{x\to a} f(x)=f(a)$. In other words, the limit comes out to exactly what we were expecting: the function value. This is why (in many cases) you can simply plug in the number to find the limit.

There are a couple of assumptions in this definition. First of all, $\lim\limits_{x\to a} f(x)$ has to exist. Remember, that means $\lim\limits_{x\to a^-} f(x)=\lim\limits_{x\to a^+} f(x)$. The greatest integer function is a fine example of a function where $\lim\limits_{x\to a^-} f(x)\neq\lim\limits_{x\to a^+} f(x)$.

Side Note: If you are only looking at $f(x)$ on a closed interval, and you want to know if it is continuous at an endpoint of that interval, then you only need to look at the limit on one side (inside that interval).

Second assumption: $f(a)$ must exist. For the function $f(x)=\dfrac{x^2-1}{x+1}$, $\lim\limits_{x\to -1} f(x)=-2$, but $f(1)$ does not exist.

There are different ways that a function can be discontinuous at a point: **removable** ($f(x)=\dfrac{x^2-1}{x+1}$ at $x=1$), **jump** ($y=\lfloor x\rfloor$ at $x=2$), **infinite** ($y=\dfrac{1}{x^2}$ at $x=0$), and **oscillating**. I haven't shown an example of that last one—try $y=\cos\left(\dfrac{1}{x}\right)$ around $x=0$. Here's the graph:

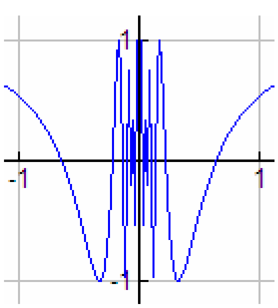

It just keeps wobbling back and forth between -1 and 1—in fact, the closer to zero you get, the worse the wobbling gets.

Continuity on an Interval

It's all well and good for a function to be continuous at a single point…just not very useful. It would be much nicer if the function was continuous over every value of some interval. When we say "$f(x)$ is continuous," what we (typically) mean is that $f(x)$ is continuous for all real numbers (even though that's not what that phrase technically means). When someone says that a function is continuous, make sure that you know what interval they mean!

Technical detail: the phrase "$f(x)$ is continuous" is supposed to mean that $f(x)$ is continuous on its domain. Some students are surprised by this, and point out that we determine the domain by eliminating discontinuities! Thus, they argue, every function is continuous on its domain.

Not so—and the fact even drove one man to insanity. Consider Cantor's Function: $C(x)=\begin{cases}1 & x \text{ is rational}\\ 0 & x \text{ is irrational}\end{cases}$. By its definition, the domain is all real numbers. This is a function—it passes the vertical line test. However, it is not continuous at any point, because there are an infinite number of jump discontinuities on any interval you care to investigate.

For extra credit brownie points, sketch the graph of Cantor's functions.

Combinations of Continuous Functions

Sums, differences, products, quotients and composites of continuous functions are continuous. 'Nuff said.

Intermediate Value Theorem

The nice thing about continuous functions is that they are predictable—they don't jump and twitch around; you know where they're going. The **Intermediate Value Theorem** says this: if a function $f(x)$ is continuous on some interval $[a,b]$, then $f(x)$ takes on *every* value between $f(a)$ and $f(b)$.

The immediate use/consequence of this: suppose I tell you that $f(x)$ is continuous, $f(2)=4$ and $f(5)=-1$. As a result, you immediately know where one of the zeros of $f(x)$ is located! By the Intermediate Value Theorem, since zero is between 4 and -1, there must be an *x*-value between 2 and 5 so that $f(x)=0$.

2.4: Rates of Change and Tangent Lines

Okay, enough about those technical details—let's get back to the motivating idea, speed (velocity; I know!). As you know from science class, speed is a change in position for a certain change in time. Thus, speed is the *rate of change in position*. Also, acceleration is the *rate of change in speed*. Keep this in mind.

Average Rates of Change

The speed that you know and love (or can calculate) is really *average speed*—an **average rate of change**. You use two fixed points to find it. Similarly, when you calculate acceleration, you use two fixed points and find an average rate of change of speed.

Here's a graphic example: I start at a very low speed as I get onto the onramp for I-26, then pick up speed until I approach the speed limit (which is about 114 kph for this example).

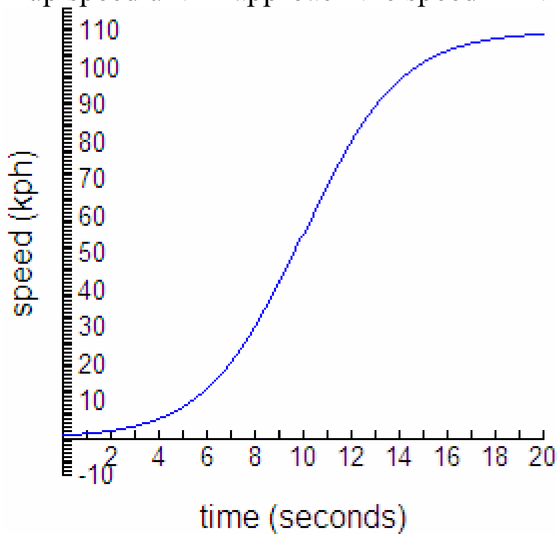

Since my speed is increasing, I am accelerating (despite the annoyed looks from the other drivers). I can find an average acceleration by taking two points from this graph and finding the slope between them. This is technically called the *slope of the secant line*.

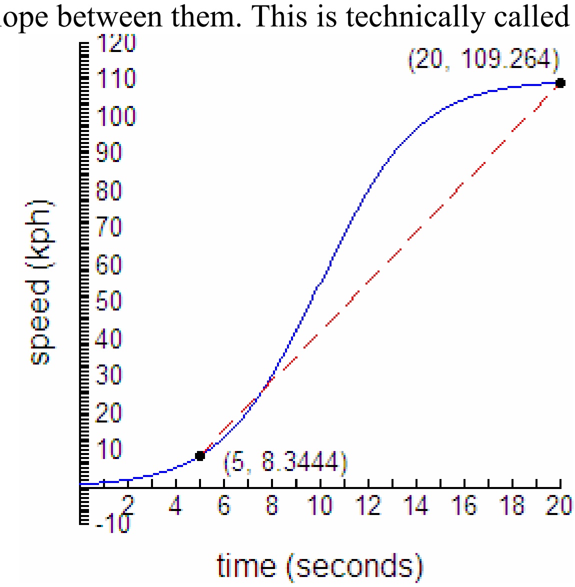

So, between $t = 5$ and $t = 20$, my speed increased from 8.344 kph to 109.264 kph. That's an average change of $\dfrac{109.264 - 8.34}{20 - 5} = \dfrac{100.924}{15} = 6.728$. For the interval between 5 and 20 seconds, my average acceleration was 6.728 kilometers per hour per second. Let's not fool with those units right now…leave them be.

Does that number (6.728) give a good approximation of my acceleration around time $t = 10$?

Instantaneous Rates of Change

I don't think so. Look at the graph; my speed looks like it is rising much faster than that secant line.

The problem is that the secant comes from a fairly large interval—there's some variation in that interval, especially around a time of 10 seconds. Maybe I should try a smaller interval around the 10 second mark.

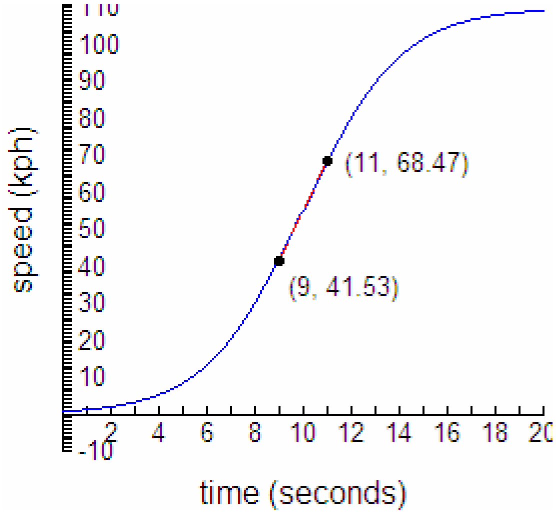

That's an average of 13.47—much higher. The secant line also appears to follow the curve better.

However, we can do even better. The smaller the interval, the better the secant line follows the curve—the better the line approximates the curve. Also, the smaller the interval, the closer we'll get to an **instantaneous rate of change**.

You can't actually have *no* time change in the formula, but you *can* take the limit as the time interval approaches zero.

Tangent Lines

As the distance between the two points on a curve gets small, the line connecting those points gets closer and closer to a tangent line. I know, I know—you thought tangents were just for circles. Not anymore. In fact, we're going to *define* the tangent line to a curve at a point as the **limit of the secant lines around that point**. Furthermore, we're going to define the slope of that tangent line to be the **slope of the curve at that point**.

The Slope of a Curve

Let's get technical.

We've got a curve—$y = f(x)$—and we want to find the slope of the curve (the slope of the line tangent to the curve) at the point $(a, f(a))$. In order to find the slope of the secant line, we'll need another point. Trust me on this: we want to let the *x*-coordinate of this other point to be $a + h$. That makes the other point on the curve (and on the secant line) $(a + h, f(a + h))$.

The slope between these two points is $\dfrac{\text{change in } y}{\text{change in } x} = \dfrac{f(a+h) - f(a)}{(a+h) - a} = \dfrac{f(a+h) - f(a)}{h}$.

Side Note: the term **difference quotient** refers to $\frac{f(a+h)-f(a)}{h}$.

I said that the tangent line is the limit of the secant lines as the two points get closer together. To do that, I can let the value *h* approach zero.

The slope of the tangent line is $\lim_{h\to 0}\frac{f(a+h)-f(a)}{h}$.

If I have a function that gives me my speed at any time *t*, then I can calculate my *instantaneous* acceleration at time *a* by finding the slope of tangent line at time *a*!

If I want to find my instantaneous speed at time *a*, I must find the slope of the tangent line for my position at time *a*!

The Normal to a Curve

Normals are perpendicular to tangents. The line normal to a curve at a point will be perpendicular to the line tangent to the curve at that point. This is especially useful in optics, where the angle of refraction depends on the angle that the light ray makes with the normal to the surface.

Examples

[5.] Find the average rate of change in $y = x^3$ on the interval $[1,3]$.

The point on the left is $(1,1)$, and the point on the right is $(3,27)$. The slope of the secant line gives the average rate of change: $\frac{27-1}{3-1} = \frac{26}{2} = 13$.

[6.] Find the slope of the line tangent to $y = x^2$ when $x = 3$.

Limit time, once again. The slope of the tangent line is $\lim_{h\to 0}\frac{(3+h)^2-(3)^2}{h}$. If you've been paying attention, we did that as an example earlier! The result was 6, so the slope of the line tangent to $y = x^2$ at $x = 3$ is 6.

[7.] Find the equation of the line tangent to $y = x^3$ when $x = 2$.

Careful—the result of this must be an equation of a line. I need two things for that: the slope, and a point on the line. Since this line is tangent to $y = x^3$, it will touch that curve. Since the *x*-value of interest is 2, that's the point where the curve and the tangent touch. What are the coordinates of that point? Specifically, what's the *y*-coordinate?

The point is $(2,8)$. Now for the slope of the tangent line: limits again!

The slope is $\lim_{h\to 0}\frac{(2+h)^3-(2)^3}{h}$. Let's FOIL:

$\lim_{h\to 0}\frac{(2+h)^3-(2)^3}{h} = \lim_{h\to 0}\frac{(8+12h+6h^2+h^3)-(8)}{h}$. A-ha! Cancel!

$\lim_{h\to 0}\frac{\left(8+12h+6h^2+h^3\right)-(8)}{h}=\lim_{h\to 0}\frac{12h+6h^2+h^3}{h}$. Cancel again!

$\lim_{h\to 0}\frac{12h+6h^2+h^3}{h}=\lim_{h\to 0}\frac{h\left(12+6h+h^2\right)}{h}=\lim_{h\to 0}\left(12+6h+h^2\right)$. Now that the coast is clear, plug in.

$\lim_{h\to 0}\left(12+6h+h^2\right)=12$. The slope of the tangent line is 4.

So the equation of the tangent line is $y-8=12(x-2)$.

That's a perfectly good equation—don't ruin it by simplifying! I'm serious—this would be accepted on the AP exam. If you try to simplify and fail, then you **lose points**. Don't simplify, except by request or necessity.

[8.] Find the equation of the line normal to $y=x^2+2x$ at $x=1$.

Let's first find the point of tangency: $(1,3)$.

Now, find the slope of the tangent line.

$$\lim_{h\to 0}\frac{\left[(1+h)^2+2(1+h)\right]-\left[(1)^2+2(1)\right]}{h}=$$

$$\lim_{h\to 0}\frac{\left[1+2h+h^2+2+2h\right]-\left[1+2\right]}{h}=\lim_{h\to 0}\frac{4h+h^2}{h}=\lim_{h\to 0}(4+h)=4$$

If the tangent line has slope 4, then the normal will have a slope of $-\frac{1}{4}$.

The equation of the normal is $y-3=\left(-\frac{1}{4}\right)(x-1)$.