

Welcome to Calculus! Here are my supplemental notes for one-variable Calculus, giving alternative ways to think about some things, practical advice, and sometimes more theoretical detail.

This does not cover everything that you need to know (although it covers a lot); you should also have the official course textbook, which is the 3rd Edition of *University Calculus: Early Transcendentals* by Hass et al published by Addison–Wesley (Pearson). There are also some references in these notes to that textbook.

Preliminaries

Before beginning this class, you should be familiar with the basic algebraic properties of real numbers.

By default, all of the numbers that we work with will be real numbers; most of Calculus applies just as well to complex numbers, but a complete understanding of Calculus in even one complex variable requires some ideas from multivariable Calculus, which these notes do not cover. In particular, if a is a negative number, then $\sqrt[n]{a}$ is undefined when n is an even integer and negative when n is an odd integer. More generally, if a is a negative number, then a^p is defined *only* if p is a rational number whose denominator in lowest terms is odd; in this case, a^p is positive if the numerator of p is even and negative if the numerator of p is odd. Note that $(a^2)^{1/2} = \sqrt{a^2} = |a|$, while $a^{2 \cdot 1/2} = a^1 = a$, which is different when a is negative, so the rule that $(a^x)^y = a^{xy}$ does not hold in general (although it does hold when a is a positive number).

Although 0^x is undefined whenever x is negative (because this amounts to dividing by zero), we need to define $0^0 = 1$ in order to make some formulas work correctly. Although the course textbook says that 0^0 is undefined, this contradicts some things that the textbook says about polynomials and power series. (Section 9.7 of the textbook, beginning with the definition of power series on page 523, is the first place where this is important; see also the discussion of power series in the full version of these notes.) It's possible to take a more nuanced approach where 0^x is 1 when x is a *constant* with the value 0 while 0^x is undefined when x is a *variable* with the value 0; however, this makes the meaning of 0^0 ambiguous without context, so I prefer to simply say that $0^0 = 1$. Nevertheless, this will require some care when it comes to rules for evaluating limits.

The main difference in my approach to Calculus from the textbook's is that I make more use of differentials. Calculus was originally developed using differentials, and many calculations are easier to do this way. Furthermore, differentials are often used in applications, especially (but not only) to physics. They fell out of fashion with mathematicians towards the end of the 19th century, when Calculus was first put on a rigorous logical foundation, because this foundation did not include differentials. However, that problem was solved by the middle of the 20th century, so there is no longer any reason to avoid them. You can do almost everything in the textbook's way if you want, but I encourage you to try using differentials.

Functions

Another difference between these notes and the textbook is that I will never be sloppy with function notation.

In an expression such as

$$y = f(x),$$

the variables x and y stand for real numbers, while the variable f stands for a function. (Usually this variable is actually a constant, because f always refers to the same function throughout the problem, although there can also be situations where the function itself is allowed to vary.) A function is not a number but rather a process for turning one number into another. When speaking of specific numbers, this is usually not a problem; for example, $f(2) = 4$ means that the function f is a process that (among other things) turns the number 2 into the number 4.

The statement that $f(x) = x^2$ is more ambiguous; in a context where the variable x already appears, this means that the function f is a process that turns the number x (whatever number that is) into the number x^2 . But in a context where x does not already have a meaning, this statement usually means that the function f is a process that turns *every* real number into its square, which is a complete description of the function. In this case, it is better to say something like

$$f(x) = x^2 \text{ for all } x,$$

and I will usually say something like this.

Another way to completely describe this function is to write

$$f = (x \mapsto x^2).$$

This is analogous to defining a set S as $S = \{x \mid x > 2\}$; in each case, you introduce a new *dummy variable* and then give an expression (to define a function) or an equation or inequality (or other statement, to define a set) that refers to that dummy variable. You can even do this without giving the function (or set) a name, by just referring to the function ($x \mapsto x^2$). Although the textbook does this with sets, it never does this with functions; so in order to avoid burdening you with additional notation to learn, I will not do this either. It can be very handy, however.

The real problem is when the same symbol is used both to refer to a function and to its output value, as in

$$A = A(x),$$

which you might see (for example) in a problem in which the area of some shape depends on something else. I will never do this! Either I will use A to refer to the area itself, or I will use A to refer to the function that indicates how this area depends on x , but I will not use the same symbol for both of these. If I need to refer to both of these, then I will use two different symbols. Most of the time, however, it's enough to have a symbol for the area itself and to leave the function unnamed. (The notation for evaluation described below on page 3 can help with this.)

When we cover derivatives later on, there are various symbols used for this concept; and when $y = f(x)$, then I will also write

$$\frac{dy}{dx} = f'(x).$$

(What this means is explained in the full version of these notes.) The textbook will sometimes write y' or df/dx in this situation, but I never will, and this is important to ensure that the ordinary rules of algebra continue to apply to such expressions. (For example, you can multiply both sides of the equation above by dx to get $dy = f'(x) dx$, which would be difficult to do correctly using the wrong symbols.) I will not count it against you if you are as sloppy as the textbook about this, because I don't think that it's fair to require more of you than the textbook writers can manage; however, if you get confused by your notation and make a mistake, then that will count against you! So I encourage you to follow my notation.

Variables

In Calculus, we study *variable* quantities, that is quantities whose values may vary (or change).

In Algebra, we often use the word 'variable' to refer to any quantity whose value we don't know, even if this value is fixed and never changes throughout the problem. In fact, the standard Algebra problem, solving an equation such as $2x + 3 = 5$, involves figuring out the value of the variable; so it had only one value all along, and we just had to figure out what it was. So if x is a variable in an Algebra problem, and at some point we decide that the value of x is 1, then this may well mean that x is 1 throughout the entire problem. (That's not always the case in Algebra, but it often is.)

In Calculus, we take the word 'variable' more seriously. If x is a variable in a Calculus problem, then x might be 1 at some point, but it will probably be 6 at some other point in the problem. (And more often than not, it will take all of the values in between 1 and 6 along the way, such as $1\frac{1}{2}$, π , and 5.789.) Furthermore, if x and y are two variables that appear in the same problem, then the value of y will usually change as the value of x changes. Calculus is primarily about exactly this sort of thing: *how* one quantity changes as another quantity changes.

In the simplest cases, it turns out that y is a function of x ; that is, there is a fixed function f such that $y = f(x)$ remains true as x and y vary. Calculus textbooks generally try to fit everything into this mould, but it doesn't always come out like this naturally. Often, you know that both x and y are changing, but it's not obvious that the value of x at some point is enough information to figure out the value of y at that point; yet when you write $y = f(x)$, you're assuming that it is enough information.

Most of the time, however, we can assume that there is some variable t , called the *independent variable*, such that every other variable in the problem is a function of t . That is, if x and y appear in the problem, then there are fixed functions f and g such that $x = f(t)$ and $y = g(t)$ throughout the problem. (Then x and y are called *dependent variables*, since their values depend on the values of t , through the functions f and g .) But this variable t might not show up directly! Calculus books will usually tell you (especially in word problems) that it's necessary to pick an independent variable, but it's enough to visualize the range of variation in the problem, and you can treat all of the variables on an equal footing. All the same, for the sake of formal definitions, we will assume that there is an independent variable t and that every other variable is a function of it, even though in practice we don't have to identify it. (Of course, you don't have to call the independent variable 't', but I usually will, just to have a consistent name.)

If we're not going to refer directly to t , then we're not going to refer directly to f and g either, so we need some way to refer to the values of these functions without referring to the functions themselves. Here is how we do it formally:

If $u = f(t)$, then $u|_{t=c} = f(c)$.

(This is called *evaluation notation*.) More generally, if P is some statement that is only true once, then P implies the statement $t = c$ for some value of c , so we can make sense of $u|_P$. Even if P is a statement that might not only be true once, as long as every possible value of $u|_P$ is the same, then we can still make sense of $u|_P$. Finally, even if there are different possible values of $u|_P$, then the value of $u|_P$ still varies, but at least it doesn't vary as much as u itself, since there are now fewer possibilities.

This all sounds very abstract (because it is), but the concrete application is straightforward; here are some examples:

$$\begin{aligned}x|_{x=5} &= 5, \\(2x + 3)|_{x=4} &= 2(4) + 3 = 11, \\(2x + 3y)|_{\substack{x=4, \\ y=5}} &= 2(4) + 3(5) = 23.\end{aligned}$$

Taking the last of these for example, there is no need to think about what t is when $x = 4$ and $y = 5$, and indeed without considering how x and y depend on this unspecified independent variable t , the value of t is impossible to know. Nevertheless, we know that no matter what t may be, if $x = 4$ and $y = 5$, then $2x + 3y$ is definitely $2(4) + 3(5) = 23$, and that is enough. So all that you have to do in practice is to plug and chug.

Sometimes (generally only in the middle of a problem or in something theoretical) you can't work out the value completely; for example,

$$(2x + 3y)|_{x=4} = 2(4) + 3(y|_{x=4}) = 8 + 3y|_{x=4}.$$

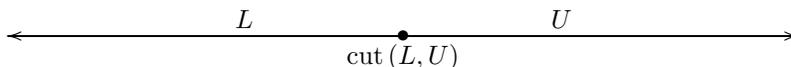
If we don't know anything more about the relationship between x and y , then we don't know the value of y when $x = 4$, so this is all that we can say in this example, but at least we were able to work out part of it.

Real numbers

In this course, we work with the real numbers, which are supposed to correspond to points on a number line. Ultimately, all of the properties of real numbers derive from intuitive geometric properties of points on a line. For example, the arithmetic operations of addition, subtraction, multiplication, and division can be defined in terms of changes of position and scale on the number line. The order relation between real numbers ($<$ and $>$) also derives from relative position on a line. (You have to specify the numerical values of at least two points, such as 0 and 1, in order to make a geometric line into a number line, but once you have those two points, then everything else follows.)

The most advanced of the fundamental properties of the number line is its *completeness*. There are many ways to express completeness, but my favourite is this:

If you pick out two nonempty regions of the number line, one on the left called L and one on the right called U , which don't overlap but otherwise cannot be extended further, then there is a single point between them, called $\text{cut}(L, U)$, the *cut* between L and U .



We can make this logically precise (in terms of the order relation on real numbers): Suppose that L and U are two sets of real numbers (making precise what regions of the number line are), with these properties:

- There is some $r \in L$ and some $s \in U$ (which is what it means for L and U to be nonempty);
- If $r \in L$ and $s \in U$, then $r < s$ (which is what it means for L to be on the left and U on the right without overlapping);
- If $r < s$, then $r \in L$ or $s \in U$ (which is what it means to say that L and U cannot be extended further).

(Note that 'or' in math, as here, normally includes the possibility of both.) Then there exists a real number $\text{cut}(L, U)$ with this property:

- If $r \in L$ and $s \in U$, then $r \leq \text{cut}(L, U) \leq s$ (which is what it means for $\text{cut}(L, U)$ to be between L and U).

A couple more important properties follow from what was said above:

- The number $\text{cut}(L, U)$ is the *only* real number between L and U ;
- If $r < \text{cut}(L, U) < s$, then $r \in L$ and $s \in U$.

The point of all this is to be able to prove that a real number exists. For example, in order to prove rigorously that every real number c has a cube root $\sqrt[3]{c}$ (and has anybody ever showed you why this is true or did you just take it on faith?), you first define L as $\{x \mid x^3 < c\}$ and U as $\{x \mid x^3 > c\}$, check that L and U have the necessary properties listed above, conclude that $\text{cut}(L, U)$ exists, and check (using the properties of $\text{cut}(L, U)$ listed above) that $\text{cut}(L, U)^3 = c$. Thus, this cut is the cube root $\sqrt[3]{c}$.

This method of proving that a real number exists is also practical, because it shows us how to approximate its value as closely as we like. For example, to approximate $\sqrt[3]{2}$ to 4 decimal places, you look at some nearby possibilities, such as 1.0001, 1.0002, 1.0003, ..., 1.9997, 1.9998, 1.9999. Somewhere in this list are two numbers right next to each other, one of which has a cube less than 2 (so it's in L) and one of which has a cube greater than 2 (so it's in U). Then we approximate $\sqrt[3]{2}$ to 4 decimal places by saying that it's in between these two numbers. (As it happens, these two numbers are 1.2599 and 1.2600; also, $1.25995^3 > 2$, so $\sqrt[3]{2}$ rounds to 1.2599.) There are more efficient ways to calculate cube roots, but this proof that they exist at least gives *one* way to calculate them, to start with.

More to come!