

Lecture 11: Extrema

Nathan Pflueger

2 October 2013

1 Introduction

In this lecture we begin to consider the notion of extrema of functions on chosen intervals. This discussion will continue in the lectures titled “optimization.” The basic idea is that the functions constructed in models (e.g. economic models, or models for how a physical system will respond to stimulus) often measure something desirable (profit, or productivity), or something undesirable (like expenses, or waste); a critical second stage in modeling is to identify the high points (for desirable quantities) or low points (for undesirable quantities).

For today, we discuss the basic techniques to find the maximum and minimum of a function, and demonstrate these on the functions that we currently know how to differentiate.

The reference for today is Stewart section §4.2 and §4.3.

2 Key phrases

There is a lot of vocabulary in this lecture. None of it is too complicated, but it can be confusing at first. For your reference, here is a list of the key phrases, so that you can make sure you have covered all the bases when you study.

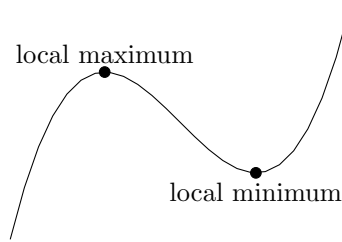
- Local maximum/minimum.
- Critical number.
- The first derivative test.
- Absolute maximum/minimum.
- The closed interval method.

3 Local extrema

The word “extremum” (plural “extrema”) simply means “maximum or minimum.” So this lecture is about locating points that are extremal in one sense or another: the largest or the smallest in some domain. First, we consider so-called *local* extrema.

I’ll start by stating the main theoretical facts, and then illustrate them all with examples. I suggest that you read the examples in parallel with the general statements.

Definition 3.1. A *local maximum* is a point $(c, f(c))$ on the graph of a function $f(x)$ such that $f(c) \geq f(x)$ for all x in *some* interval around c . Similarly, a *local minimum* is a point $(c, f(c))$ such that $f(c) \leq f(x)$ for all x in some interval around c .



Note: remember that a local maximum does *not* need to be where the function takes its largest possible value. It is just the largest in some region around itself.

A basic observation (apparently first described by Fermat) is that a local extremum (max or min) can only occur when the graph either lies perfectly flat, or has some sort of sharp corner. Therefore make the following definition.

Definition 3.2. A *critical number* is a number c such that the **either** $f'(c) = 0$ **or** $f'(c)$ is undefined.

Two basic facts govern how calculus can be used to locate local extrema: Fermat's theorem and the first derivative test.

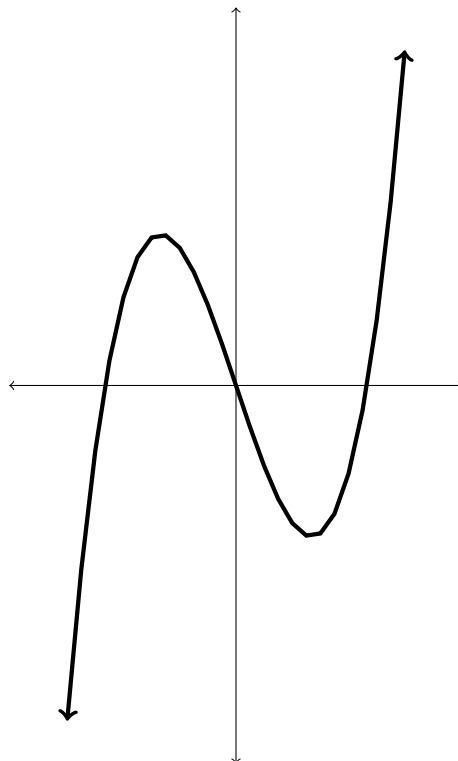
Theorem 3.3 (Fermat's theorem). *If $(c, f(c))$ is a local extremum of $f(x)$, then c is a critical number.*

Theorem 3.4 (First derivative test). *Suppose that c is a critical number.*

- *If $f'(x)$ changes sign from **positive to negative** around c (positive on the left, negative on the right), then $(c, f(c))$ is a **local maximum**.*
- *If $f'(x)$ changes sign from **negative to positive** around c (negative on the left, positive on the right), then $(c, f(c))$ is a **local minimum**.*
- *If $f'(x)$ has the **same sign** on both sides of $x = c$, then there is **no local extremum** at $x = c$.*

The following examples illustrates all the facts and terms mentioned above.

Example 3.5. Let $f(x) = x^3 - 3x$. Find all critical numbers, and identify all local extrema.



Solution. First, to find the critical numbers, use the power rule to obtain the derivative $f'(x) = 3x^2 - 3$. Since $f(x)$ is a polynomial, the derivative is defined everywhere, so the only critical numbers are those values where $f'(c) = 0$. Now, $f'(x) = 0 \Leftrightarrow 3x^2 - 3 = 0$. Solving this equation gives $x = \pm 1$. So the critical numbers are -1 and 1 .

By Fermat's theorem, the only possible local extrema are $(-1, f(-1))$ and $(1, f(1))$, i.e. $(-1, 2)$ and $(1, -2)$.

Use the first derivative test to identify which is which (of course, in this case it is obvious from the picture). The derivative $f'(x) = 3(x^2 - 1)$ is positive when $x^2 > 1$ and negative when $x^2 < 1$. So it is positive for $x < -1$, negative for $-1 < x < 1$, and positive again for $x > 1$. Therefore the derivative changes sign from $+$ to $-$ at $x = -1$; therefore $(-1, 2)$ is a local maximum. The derivative changes sign from $-$ to $+$ at $x = 1$; therefore $(1, -2)$ is a local minimum. Fermat's theorem shows that there can be no other local extrema.

Example 3.6. Let $f(x) = e^x - x$. Find all local extrema.

Solution. We can compute that $f'(x) = e^x - 1$. This is always defined; it is positive for $x > 0$ (since this is when $e^x > 1$) and negative for $x < 0$; $f'(x) = 0$ only when $x = 0$ (since $e^x - 1 = 0$ if and only if $e^x = 1$, i.e. $x = \ln 1 = 0$). So the only critical number is 0 , around which the derivative changes sign from negative to positive. So there is a local minimum at $(0, 1)$ and no other local extrema.

Example 3.7. Let $f(x) = x^3$. Find all critical numbers and local extrema.

Solution. By the power rule, $f'(x) = 3x^2$. This is always defined. $f'(x) = 0$ if and only if $3x^2 = 0$, i.e. $x = 0$. So the only critical number is 0 .

The derivative is never negative (since the square of a real number is never negative); in particular it does not change sign around $x = 0$. So by the third possibility in the first derivative test, 0 is neither a local maximum nor a local minimum. There are no other critical numbers, so by Fermat's theorem there are no local extrema at all.

Example 3.8. Let $f(x) = |x|$. Find all critical numbers and all local extrema.

Solution. Recall that $|x| = \begin{cases} x & x \geq 0 \\ -x & x < 0 \end{cases}$. Therefore it follows from this that $f'(x) = 1$ for all $x > 0$ and $f'(x) = -1$ for all $x < 0$. However, $f'(0)$ is not defined at all, since the function is not even continuous at $x = 0$.

So since $f'(x)$ is undefined at $x = 0$, this is a critical number. There are no other critical numbers since $f'(x)$ is never 0. The only critical number is 0.

The sign of $f'(x)$ changes from negative to positive around $x = 0$, so the first derivative test tells us that (0, 0) is a local minimum. There are no other critical numbers, hence no other local extrema.

4 Absolute extrema

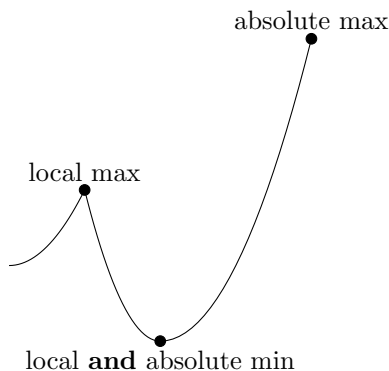
Now, we consider not local extrema, but what are called *absolute* extrema: those points on a graph that are greater (or less) than *all* other possible points. In physical situations, there is generally only some small range of possible values of x , so we will generally work on some chosen interval. Recall the standard notation for intervals.

Notation The notation $(-1, 1)$ denotes the interval of all x with $-1 < x < 1$ (in particular, not including the endpoints). This is called an open interval. The notation $[-1, 1]$ denotes the interval of all x with $-1 \leq x \leq 1$ (in particular, this includes the endpoints). This is called a closed interval. The notation $[-1, 1)$ denotes all x with $-1 \leq x < 1$ (including one endpoint but not the other), and similarly $(-1, 1]$ denotes all x with $-1 < x \leq 1$. These are called half-open intervals.

Definition 4.1. Given a function $f(x)$ and an interval I (open, closed, or half-open), we say that $(c, f(c))$ is an *absolute maximum* if $f(c) \geq f(x)$ for all other x in I . We say that it is an *absolute minimum* if $f(x) \leq f(c)$ for all other x in I .

Note. It is possible that there is a “tie” for the absolute maximum or absolute minimum. So the maximum and minimum do not have to be unique (though the maximum/minimum *value* is unique).

For example, see the following picture.

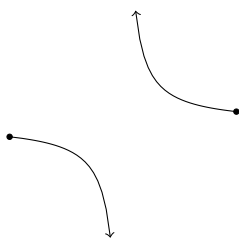


Absolute extrema are not always local extrema. An absolute extremum does not need to be a local extremum. As we see above, the absolute max is not labelled a local max as well. This is essentially a matter of terminology. The reason we don’t call this point a local max is that we only have information about the function on one side but not the other; the word “local” will be reserved for points in the inside of the interval. For our purposes, just think of a local extremum as only an extremum you can locate by finding a critical number of the function (where the derivative is zero or undefined).

A function does not always have absolute extrema on an interval. The following two examples illustrate two of the ways a function could fail to have absolute extrema. Then we will give a criterion for

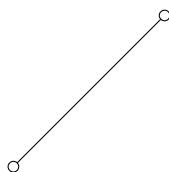
begin sure that it will have absolute extrema.

Example 4.2. Consider the function $f(x) = \frac{1}{x}$ on the interval $[-1, 1]$.



This function is **discontinuous** at $x = 0$. It also has **no absolute maximum or minimum**. This is because for any input where the function is defined, there is always another input that gives a larger (or smaller) value.

Example 4.3. Consider the function $f(x) = x$ on the open interval $(-1, 1)$.



Somewhat counterintuitively, this function also **has no absolute maximum or minimum**. This is counterintuitive because it really looks like the maximum is 1, so why isn't it? The issue here is fairly subtle: the problem is that the maximum ought to be at the endpoint, but the **interval is open**, meaning it doesn't contain the endpoints. So while $f(0.99) = 0.99$ is quite close to 1, you can still find a larger value of the function, such as $f(0.999)$. So there isn't single number x in $(-1, 1)$ such that $f(x)$ is the biggest value possible.

These examples show two things that can prevent a function from having a maximum on an interval. The following fact ensures that these are the *only* things that can go wrong: as long as the function is continuous, and the interval is closed, there is an absolute maximum and absolute minimum.

Theorem 4.4 (Extreme value theorem). *A **continuous** function $f(x)$ always has an absolute maximum and an absolute minimum on any **closed** interval $[a, b]$.*

Of course it's all well and good to know just that the absolute maximum exists, but how do we actually find it? It turns out that there's always a pretty short list of candidates: the absolute maximum is always either a local maximum on the inside of the interval (and therefore comes from a critical number), or else it occurs at an endpoint of the interval. The same holds for an absolute minimum. So there is a very straightforward method for finding the absolute extrema.

Closed interval method. Suppose $f(x)$ is a continuous function on a closed interval $[a, b]$.

1. Compute the derivative $f'(x)$ where it exists, and find all the critical numbers (that is, write down all the non-differentiable points, then solve $f'(x) = 0$ and add those solutions to the list).
2. Compute $f(c)$ for all critical values c (found in step 1) and also at the two endpoints $c = a$ and $c = b$ of the interval.
3. Find the maximum and minimum of the values computed in step 2. These are the absolute maximum and absolute minimum on this interval.

Here are a few examples of this method.

Example 4.5. Find the absolute extrema of $f(x) = x^{2/5}(x^{1/5} - 1)$ on $[0, 32]$.

Solution. First, compute the derivative with the power rule as follows.

$$\begin{aligned}f(x) &= x^{3/5} - x^{2/5} \\f'(x) &= \frac{3}{5}x^{-2/5} - \frac{2}{5}x^{-3/5}\end{aligned}$$

Note that this is defined everywhere except at $x = 0$, so 0 is one critical number. To find the others, solve $f'(x) = 0$ as follows.

$$\begin{aligned}\frac{3}{5}x^{-2/5} - \frac{2}{5}x^{-3/5} &= 0 \\ \frac{3}{5}x^{-2/5} &= \frac{2}{5}x^{-3/5} \\ \frac{x^{-2/5}}{x^{-3/5}} &= \frac{2/5}{3/5} \\ x^{1/5} &= \frac{2}{3} \\ x &= \left(\frac{2}{3}\right)^5\end{aligned}$$

So there are two critical numbers: $x = 0$ and $x = \left(\frac{2}{3}\right)^5$.

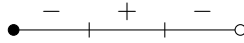
Now, the two endpoints are 0 and 32. So evaluate $f(x)$ on the critical values and the endpoints (note that 0 is both a critical number and an endpoint).

$$\begin{aligned}f(0) &= 0^{2/5}(0^{1/5} - 1) \\ &= 0 \\ f\left(\left(\frac{2}{3}\right)^5\right) &= \left(\left(\frac{2}{3}\right)^5\right)^{2/5} \left(\left(\left(\frac{2}{3}\right)^5\right)^{1/5} - 1\right) \\ &= \left(\frac{2}{3}\right)^2 \left(\frac{2}{3} - 1\right) \\ &= -\frac{4}{27} \\ f(32) &= 32^{2/5}(32^{1/5} - 1) \\ &= 4(2 - 1) \\ &= 4\end{aligned}$$

So the three values we compute in step 2 of the closed interval method are 0 , $-\frac{4}{27}$, and 4 . The maximum of these is the absolute maximum: $f(32) = 4$. The smallest is the absolute minimum: $f\left(\left(\frac{2}{3}\right)^5\right) = -4/27$.

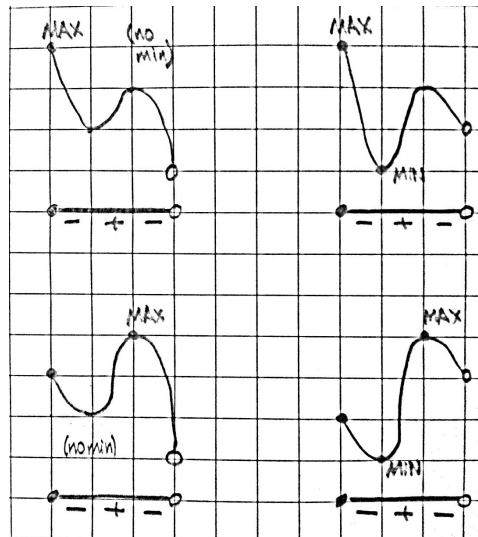
5 Sign diagrams

In this section, I just explore a bit exactly how much you can know about a function if you just know where it is increasing and decreasing. For example, suppose that you have a function $f(x)$ on $[1, 4]$, and all you know about it is that it has critical numbers 2 and 3, it is increasing on $[1, 2]$, decreasing on $[2, 3]$, and increasing on $[3, 4]$. This information can be denoted by a “sign diagram.”



Can we conclude from this whether $f(x)$ has absolute extrema, and if so where they must be?

To answer this, notice that we certainly know (by the first derivative test) that there is a local minimum at 2 and a local maximum at 3. Both of these could be absolute extrema, but they also might not. Shown below are some of the several possibilities.



In fact, we can reason that **there must be a global maximum** as follows: there is a global maximum on $[1, 3]$ (since this is a closed interval), and $f(x)$ is decreasing after 3, so it won't get any bigger. So in fact that absolute maximum always exists, and it must occur either at either 1 or 3. On the other hand, **we don't have enough information to tell if there's a global minimum..** Indeed, we can see in the pictures above that sometimes there is a global minimum at 2 (as in the pictures on the right), but sometimes there is no global minimum (as at the pictures on the left).