Extrema[1]

Let $f : D \to \mathbb{R}$ where $D \subseteq \mathbb{R}^n$.

$\overline{\mathbf{x}} \in D$ is a **local maximum** of f if and only if there exists an $\varepsilon > 0$ such that $f(\mathbf{x}) \leq f(\overline{\mathbf{x}})$ for all $\mathbf{x} \in B_\varepsilon(\overline{\mathbf{x}})$.

If $f(\mathbf{x}) \leq f(\overline{\mathbf{x}})$ for all $\mathbf{x} \in D$, then $\overline{\mathbf{x}}$ is a **global maximum**.

$\underline{\mathbf{x}} \in D$ is a **local minimum** of f if and only if there exists an $\varepsilon > 0$ such that $f(\mathbf{x}) \geq f(\underline{\mathbf{x}})$ for all $\mathbf{x} \in B_\varepsilon(\underline{\mathbf{x}})$.

If $f(\mathbf{x}) \geq f(\underline{\mathbf{x}})$ for all $\mathbf{x} \in D$, then $\underline{\mathbf{x}}$ is a **global minimum**.

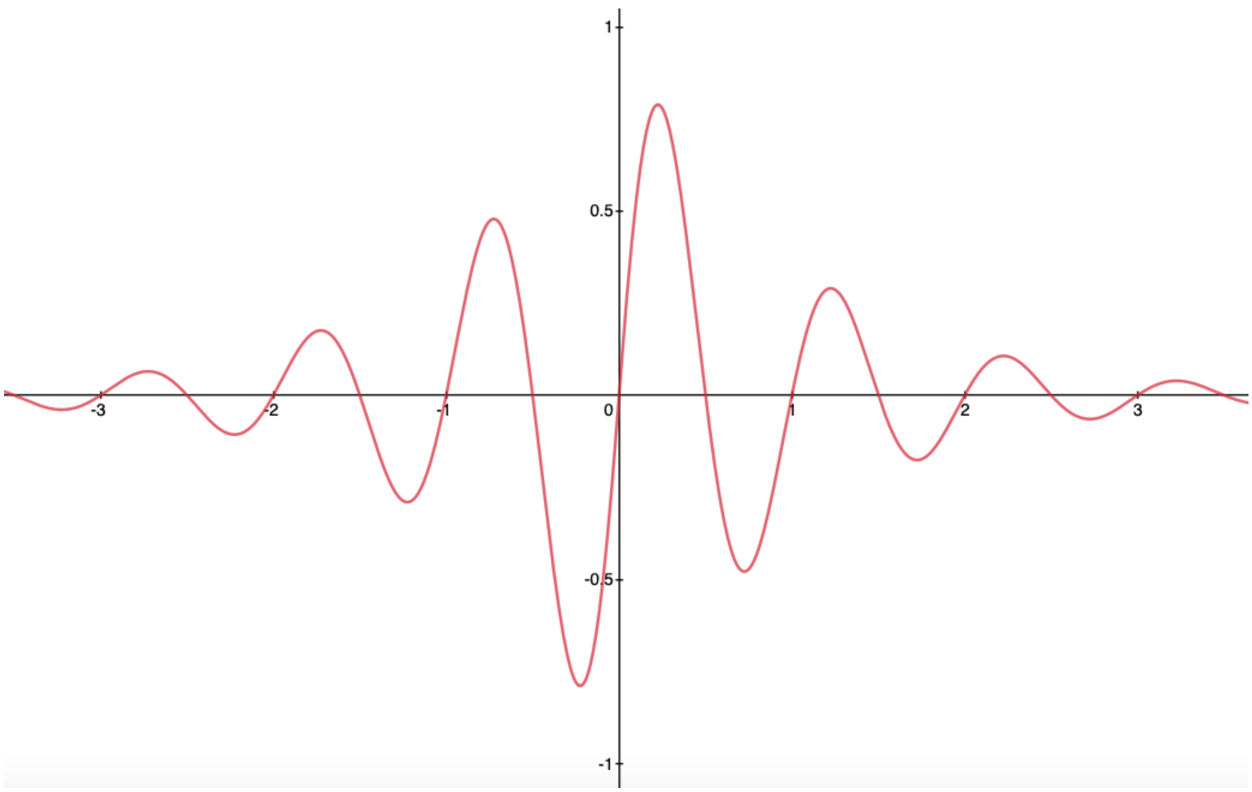

All global maxima (minima) are also local maxima (minima).

[1]Prepared by Sarah Robinson

Let $f(x)$ be a twice-continuously differentiable univariate function. Then $f(x)$ reaches a local, interior

- maximum at x^* iff $f'(x^*) = 0$ and $f''(x^*) \leq 0$
- minimum at $\tilde{x}$ iff $f'(\tilde{x}) = 0$ and $f''(\tilde{x}) \geq 0$

$f'(x) = 0$ is known as the **first-order condition**. It tells us that we are at an extrema of some kind (local maximum or minimum).

The sign of $f''(x)$ is the **second-order condition**. It tells us which kind of extremum we are at (whether we are at a local maximum or at a local minimum).

$$f''(x) \leq 0 \;\approx\; f \text{ is concave in that area } \approx\; x \text{ is a local maximum}$$

$$f''(x) \geq 0 \;\approx\; f \text{ is convex in that area } \approx\; x \text{ is a local minimum}$$

Example: $f(x) = \ln(x) - x$

We can extend to multivariate functions.

Recall that the equivalent of the first derivative is the gradient, and the equivalent of the second derivative is the Hessian. Also recall that Sarah hates remembering the rules associated with Hessians.

You can use the first-order condition with the gradient to identify potential extrema. Then you can use raw cunning and skill to sort out whether its a maximum, minimum, or saddle point. (Or look up Hessians on Wikipedia).

Let $f : D \to \mathbb{R}$ be a differentiable function, where $D \subseteq \mathbb{R}^n$. If $f(\mathbf{x})$ reaches a local extremum at $x^* \in \mathbb{R}^n$, then

$$\frac{\partial f(\mathbf{x})}{x_i} = 0 \quad \forall\, i = 1, \ldots, n$$

$$\nabla f(\mathbf{x}) = 0$$

Example. Find and classify extrema for the function

$$f(x, y) = 8x^3 + 2xy - 3x^2 + y^2 + 1$$

Finding potential extrema:

We have two potential extrema:

$$(0,0) \quad \text{and} \quad \left(\frac{1}{3}, -\frac{1}{3}\right)$$

And we have our first derivatives:

$$f_x(x,y) = 24x^2 + 2y - 6x \qquad f_y(x,y) = 2x + 2y$$

Let's look at the second derivatives for each variable:

$$f_{xx}(x,y) = 48x - 6 \qquad f_{yy}(x,y) = 2$$

Let's discuss the y situation first. For any value of x, $f_{yy} \geq 0$. For any "slice" of f holding x constant, f is convex in y.

Now let's look at the x situation:

$$f_{xx}(x,y) = 48x - 6 = 0$$

$$x = \frac{6}{48} = \frac{1}{8}$$

If $x \geq \frac{1}{8}$, then $f_{xx} \geq 0$ and f is convex in x.
If $x \leq \frac{1}{8}$, then $f_{xx} \leq 0$ and f is concave in x.

At the point $(\frac{1}{3}, \frac{-1}{3})$, f is convex in x and y. So this is a local minimum.

At the point $(0,0)$, f is concave in x and convex in y. So this is a saddle point (neither a minimum nor a maximum).

Illustration

What we typically really care about are global maxima or minima. This is easy when the function is (strictly) concave or convex everywhere.

Let $f : D \to \mathbb{R}$, where $D \subseteq \mathbb{R}^n$, be twice continuously differentiable and concave. Then the following statements are equivalent:

- $\nabla f(\mathbf{x}^*) = 0$
- f achieves a global maximum at $\mathbf{x}^*$

Further, if f is strictly concave, then x^* is the unique global maximizer, i.e., $f(\mathbf{x}^*) > f(\mathbf{x})$ for all $\mathbf{x} \in D$ such that $\mathbf{x} \neq \mathbf{x}^*$.

Let $f : D \to \mathbb{R}$, where $D \subseteq \mathbb{R}^n$, be twice continuously differentiable and convex. Then the following statements are equivalent:

- $\nabla f(\mathbf{x}^*) = 0$
- f achieves a global minimum at $\mathbf{x}^*$

Further, if f is strictly convex, then x^* is the unique global minimizer, i.e., $f(\mathbf{x}^*) < f(\mathbf{x})$ for all $\mathbf{x} \in D$.

(Minor note: we are requiring f to be twice continuously differentiable everywhere, so we are not talking about closed domains. For example, $f(x) = x$ is concave and convex. But if it's only defined on $D = [0, 1]$, then it's not differentiable everywhere. The global max is $x = 1$ but the derivative isn't defined).

Example:
Find the maximum for a profit function $\pi(K, L)$ given by

$$\pi(K, L) = p\big[\ln(K) + \ln(L)\big] - rK - wL$$

where p, r, and w are strictly positive parameters, while K and L are (strictly positive) choice variables.

Notation

Consider a general form of an unconstrained maximization problem:

$$\max_{\mathbf{x} \in D(\boldsymbol{\theta})} f(\mathbf{x}, \boldsymbol{\theta})$$

f is the **objective function** (function to maximize)

$\mathbf{x}$ is the **choice variable** (what we can move to maximize f)

D is the **choice set** (the set of options for the choice variable)

$\boldsymbol{\theta}$ is a **parameter** that may affect both the objective function and the choice set

The **solution set** is the set of all $\mathbf{x}$ that solve the maximization problem. If the solution is unique, then it's a value of $\mathbf{x}$. If the solution is not unique (i.e., there are multiple global maxima), then it's a set.

$$\mathbf{x}^*(\boldsymbol{\theta}) = \arg \max_{\mathbf{x} \in D(\boldsymbol{\theta})} f(\mathbf{x}, \boldsymbol{\theta})$$

The **value function** gives the value of the function at the solution for any parameter. If you put in the parameters, you get out the value of the function *at the optimum.*

$$V(\boldsymbol{\theta}) = f(\mathbf{x}^*, \boldsymbol{\theta})$$

Notice that even if there are multiple elements in the solution set, they each give the same value in the value function.

Example:
Find the maximum for a profit function $\pi(K, L)$ given by

$$\pi(K, L) = p\big[\ln(K) + \ln(L)\big] - rK - wL$$

where p, r, and w are strictly positive parameters, while K and L are (strictly positive) choice variables.

The objective function is $\pi(K, L, p, w, r)$

The choice variables are K and L (or a vector (K, L))

The choice set is $\{(K, L) \mid K > 0, L > 0\}$

The parameters are a vector (p, r, w) (here they affect the objective function but not the choice set)

We can write the optimization problem as:

$$\max_{K>0, L>0} \pi(K, L, p, r, w)$$

The solution set is:

$$K^*(p, w, r) = \frac{p}{r} \qquad\qquad L^*(p, w, r) = \frac{p}{w}$$

The value function is:

$$V(p, w, r) = \pi(K^*, L^*, p, w, r)$$

We discussed previously the Extreme Value Theorem, which says that there must be a global maximum and minimum (though not necessarily unique) for any continuous, real-valued function with a non-empty, compact domain.

In some cases, we might want to discuss a maximum and minimum where we aren't sure that they exist because the domain is not closed.

Then we can use the concept of the infimum and supremum $\approx$ the minimum and maximum in the limit.

$$\min\{[0,1]\} = 0 \qquad \inf\{(0,1)\} = 0$$

$$\max\{[0,1]\} = 1 \qquad \sup\{(0,1)\} = 1$$

If a maximum exists, it is also a supremum.