

FOUNDATIONS OF OPTIMIZATION

Basics

- *Optimization problems*

- An optimization problem is

$$\text{minimise } f(\mathbf{x}) \text{ subject to } \mathbf{x} \in \mathcal{C}$$

f is the *objective* (real) \mathcal{C} is the constraint set/feasible set/search space.

- \mathbf{x}^* is an *optimal solution* (*global minimizer*) if and only if

$$f(\mathbf{x}^*) \leq f(\mathbf{x}) \quad \forall \mathbf{x} \in \mathcal{C}$$

- Maximizing $f(\mathbf{x})$ is equivalent to minimizing $-f(\mathbf{x})$.
- We consider problems in the following form

$$\begin{aligned} &\text{minimize } f(\mathbf{x}) \\ &\text{subject to } h_i(\mathbf{x}) = 0 \quad \forall 1 \leq i \leq m \\ &\quad \quad \quad g_i(\mathbf{x}) \leq 0 \quad \forall m \leq i \leq r \\ &\quad \quad \quad \mathbf{x} \in \mathbb{R}^n \end{aligned}$$

- We consider the following subsets of the problem
 - In *linear programming*, all functions are linear.
 - In *convex programming*, the f and g are convex, and the h are linear.
- If \mathcal{C} is the feasible set of a problem, a point $\mathbf{x} \in \mathcal{C}$ is a *local minimum* if there exists a neighborhood $N_r(\mathbf{x})$ such that $f(\mathbf{x}) \leq f(\mathbf{y}) \quad \forall \mathbf{y} \in \mathcal{C} \cap N_r(\mathbf{x})$. It is an *unconstrained local minimum* if $f(\mathbf{x}) \leq f(\mathbf{y}) \quad \forall \mathbf{y} \in N_r(\mathbf{x})$. (Strict equivalents exist).

- *Topology*

- An *open ball* around a point $\mathbf{x} \in \mathbb{R}^n$ with radius $r > 0$ is the set $N_r(\mathbf{x}) = \{\mathbf{y} \in \mathbb{R}^n : \|\mathbf{x} - \mathbf{y}\| < r\}$, where $\|\mathbf{x}\| = \sqrt{\sum x_i^2}$.
- A point $\mathbf{x} \in \mathcal{E} \subset \mathbb{R}^n$ is an *interior point* if there exists an open ball such that $N_r(\mathbf{x}) \subset \mathcal{E}$. A set $\mathcal{E} \subset \mathbb{R}^n$ is *open* if $\mathcal{E} = \text{int } \mathcal{E}$.

- A point $\mathbf{x} \in \mathcal{E} \subset \mathbb{R}^n$ is a *closure point* if, for every open ball $N_r(\mathbf{x})$, there exists $\mathbf{y} \in \mathcal{E}$ with $\mathbf{y} \in N_r(\mathbf{x})$. A set $\mathcal{E} \subset \mathbb{R}^n$ is *closed* if $\mathcal{E} = \text{cl } \mathcal{E}$.
- The set of reals is both closed and open.
- **Theorems:**
 - The union of open sets is open. The intersection of a *finite* number of open sets is open.
 - The intersection of closed sets is closed. The union of a *finite* number of closed sets is closed.

- **Analysis**

- A sequence of vectors $\{\mathbf{x}_n\} \subset \mathbb{R}^n$ converges to a limit $\mathbf{x} \in \mathbb{R}^n$ if $\lim_{k \rightarrow \infty} \|\mathbf{x} - \mathbf{x}_k\| = 0$, and we say that $\mathbf{x}_k \rightarrow \mathbf{x}$.
- A set $\mathcal{E} \subset \mathbb{R}^n$ is (sequentially) compact if, given a sequence $\{\mathbf{x}_k\} \subset \mathcal{E}$, there is a subsequence $\{\mathbf{x}_{k_i}\}$ converging to an element $\mathbf{x} \in \mathcal{E}$.
 - **Theorem** (Heine-Borel): A set $\mathcal{E} \subset \mathbb{R}^n$ is compact if and only if it is closed and bounded.
 - **Theorem:** A closed subset of a compact set is compact.
 - **Theorem:** Suppose $\{\mathcal{E}_n\}$ are a sequence of non-empty, compact sets that are nested (ie: $\mathcal{E}_{n+1} \subset \mathcal{E}_n$) – then their intersection is non-empty.
- A real-valued function f defined on a domain $\mathcal{X} \subset \mathbb{R}^n$ is continuous at the point $x \in \mathcal{X}$ if, for every sequence $\{\mathbf{x}_k\} \subset \mathcal{X}$ with $\mathbf{x}_k \rightarrow \mathbf{x}$, $\lim_{k \rightarrow \infty} f(\mathbf{x}_k) = f(\mathbf{x})$. f is *continuous* if it is continuous at all points in \mathcal{X} .
- A function f is *coercive* over a set $\mathcal{C} \subset \mathbb{R}^n$ if, for every sequence $\{\mathbf{x}_k\} \subset \mathcal{C}$ with $\|\mathbf{x}_k\| \rightarrow \infty$, we have $\lim_{k \rightarrow \infty} f(\mathbf{x}_k) = \infty$.
- The *inverse image* of the set $\mathcal{A} \subset \mathbb{R}$ is defined by $f^{-1}(\mathcal{A}) = \{\mathbf{x} \in \mathcal{X} : f(\mathbf{x}) \in \mathcal{A}\}$.
 - **Theorem:** If f is continuous and \mathcal{X} is open^{closed} and \mathcal{A} is open^{closed}, then $f^{-1}(\mathcal{A})$ is also open^{closed}. This is the standard way to prove that a set is open/closed.

- **Definition:** A function is *convex* if $f(\lambda \mathbf{x}_1 + (1 - \lambda)\mathbf{x}_2) \leq \lambda f(\mathbf{x}_1) + (1 - \lambda)f(\mathbf{x}_2)$. It is *strictly convex* if the inequality is strict for $\mathbf{x}_1 \neq \mathbf{x}_2$. We say f is convex over $\mathcal{X} = \text{dom } f$ if it is convex when restricted to \mathcal{X} . f is (strictly) concave if $-f$ is (strictly) convex.
- **Definition:** If f is convex with a convex domain \mathcal{X} , we define the *extended-value extension* $\tilde{f} : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ by

$$\tilde{f}(\mathbf{x}) = \begin{cases} f(\mathbf{x}) & \text{if } \mathbf{x} \in \mathcal{X} \\ \infty & \text{otherwise} \end{cases}$$

and we let

$$\text{dom } \tilde{f} = \{\mathbf{x} \in \mathbb{R}^n : \tilde{f}(\mathbf{x}) < \infty\}$$

An extended-value function is convex if

- Its domain is convex.
- The standard convexity property holds.
- Given $\mathcal{C} \subset \mathbb{R}^n$, the *indicator function* $I_{\mathcal{C}} : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ as

$$I_{\mathcal{C}}(\mathbf{x}) = \begin{cases} 0 & \text{if } \mathbf{x} \in \mathcal{C} \\ \infty & \text{otherwise} \end{cases}$$

If \mathcal{C} is a convex set, then $I_{\mathcal{C}}$ is a convex function.

- **Theorem:** If f is convex over a convex set $\mathcal{C} \subset \mathbb{R}^n$, then every sublevel set $\{\mathbf{x} \in \mathcal{C} : f(\mathbf{x}) \leq \gamma\}$ is a convex subset of \mathbb{R}^n . The converse is *not* true (eg: $\log x$ on $(0, \infty)$). However, we define...
- **...Definition:** A extended real valued function $f : \mathbb{R}^n \rightarrow \mathbb{R} \cup \{\infty\}$ is *quasiconvex* if, every one of its sublevel sets (ie: for every $\gamma \in \mathbb{R}$) is convex.

• **Calculus**

- A function $f : \mathcal{X} \rightarrow \mathbb{R}$ with $\mathcal{X} \subset \mathbb{R}^n$ is *differentiable* at $\mathbf{x} \in \text{int } \mathcal{X}$ if there exists a vector $\nabla f(\mathbf{x}) \in \mathbb{R}^n$, known as the *gradient*, such that

$$\lim_{\mathbf{d} \rightarrow \mathbf{0}} \frac{f(\mathbf{x} + \mathbf{d}) - f(\mathbf{x}) - \nabla f(\mathbf{x}) \cdot \mathbf{d}}{\|\mathbf{d}\|} = 0$$

And

$$\nabla f(\mathbf{x}) = \left[\frac{\partial f(\mathbf{x})}{\partial x_1}, \dots, \frac{\partial f(\mathbf{x})}{\partial x_n} \right]^T \in \mathbb{R}^n \quad \frac{\partial f(\mathbf{x})}{\partial x_i} = \lim_{h \rightarrow 0} \frac{f(\mathbf{x} + h\mathbf{e}_i) - f(\mathbf{x})}{h}$$

f is *differentiable* over an open set $\mathcal{U} \in \mathcal{X}$ if it is differentiable at every point in the set. If, in addition, the components of the gradient are continuous over \mathcal{U} , then f is *continuously differentiable* over \mathcal{U} .

- If, for a point $\mathbf{x} \in \text{int } \mathcal{X}$, each component of the gradient is differentiable, we say f is *twice differentiable* at \mathbf{x} , and we define the *Hessian Matrix* $\nabla^2 f(\mathbf{x}) \in \mathbb{R}^{n \times n}$ by

$$\nabla^2 f(\mathbf{x}) = \left[\frac{\partial^2 f(\mathbf{x})}{\partial x_i \partial x_j} \right]_{ij}$$

If f is twice continuously differentiable in a neighborhood of \mathbf{x} , then the Hessian is symmetric.

- Suppose f is twice continuously differentiable over a neighborhood $N_r(\mathbf{x})$, then for all $\mathbf{d} \in N_r(\mathbf{0})$

$$f(\mathbf{x} + \mathbf{d}) = f(\mathbf{x}) + \nabla f(\mathbf{x})^T \mathbf{d} + \frac{1}{2} \mathbf{d}^T \nabla^2 f(\mathbf{x}) \mathbf{d} + o(\|\mathbf{d}\|^2)$$

(Formally, this means that for every $C > 0$, there exists a neighborhood around $\mathbf{0}$ such that the estimate of $f(\mathbf{x} + \mathbf{d})$ differs from the real value by no more than $C \|\mathbf{d}\|^2$.)

- Consider a vector-valued function $\mathbf{F} : \mathcal{X} \rightarrow \mathbb{R}^m, \mathcal{X} \subset \mathbb{R}^n$ and a point $\mathbf{x} \in \text{int } \mathcal{X}$. We define the *gradient* to be the matrix $\nabla \mathbf{F}(\mathbf{x}) \in \mathbb{R}^{n \times m}$ with

$$\nabla \mathbf{F}(\mathbf{x}) = \left[\nabla F_1(\mathbf{x}), \dots, \nabla F_m(\mathbf{x}) \right] \quad \nabla F(\mathbf{x})_{ij} = \frac{\partial F_j(\mathbf{x})}{\partial x_i}$$

- The chain rule states that for interior points, if $\mathbf{h}(\mathbf{x}) = \mathbf{g}(f(\mathbf{x}))$, then

$$\nabla \mathbf{h}(\mathbf{x}) = \nabla f(\mathbf{x}) \nabla \mathbf{g}(f(\mathbf{x}))$$

• *Linear algebra – Kernels and Images*

- Consider a matrix $A \in \mathbb{R}^{m \times n}$. Then
 - $\ker A = \{ \mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{0} \}$
 - $\text{im } A = \{ \mathbf{y} \in \mathbb{R}^m : \mathbf{y} = A\mathbf{x}, \mathbf{x} \in \mathbb{R}^n \}$
- Given a set $\mathcal{S} \in \mathbb{R}^n$, $\mathcal{S}^\perp = \{ \mathbf{x} \in \mathbb{R}^n : \mathbf{x} \cdot \mathbf{y} = 0 \ \forall \mathbf{y} \in \mathcal{S} \}$

- **Lemma:** $\text{im } A = [\ker(A^\top)]^\perp$. In other words, given $\mathbf{z} \in \mathbb{R}^m$,

$$\mathbf{z} = A\mathbf{x} \text{ for some } \mathbf{x} \in \mathbb{R}^n \Leftrightarrow \mathbf{z} \cdot \mathbf{y} = 0 \ \forall \mathbf{y} \text{ with } A^\top \mathbf{y} = 0$$

- **Sets, etc...**

- **Affine sets**

- **Definition:** A set $\mathcal{C} \subset \mathbb{R}^n$ is *affine* if, for all points $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{C}$ and a scalar $\lambda \in \mathbb{R}$, $\lambda \mathbf{x}_1 + (1 - \lambda)\mathbf{x}_2 \in \mathcal{C}$.

- **Example:** The empty space, a line and any subspace are affine. Similarly, $\mathcal{C} = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} = \mathbf{b}\}$ is affine.

- **Definition:** Given a set of points $\mathcal{X} \subset \mathbb{R}^n$, the *affine hull* $\text{aff } \mathcal{X}$ is the set of points $\lambda_1 \mathbf{x}_1 + \dots + \lambda_k \mathbf{x}_k$, where $k \geq 1$, $\{\mathbf{x}_i\} \subset \mathcal{X}$ and $\lambda_i = 1$. The affine hull is affine and is the smallest affine set containing \mathcal{X} .

- **Convex sets**

- **Definition:** The set \mathcal{C} is *convex* if, for all points $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{C}$ and a scalar $\lambda \in (0,1)$, $\lambda \mathbf{x}_1 + (1 - \lambda)\mathbf{x}_2 \in \mathcal{C}$. Clearly, affine sets are also convex.

- **Definition:** Given a set of points $\mathcal{X} \subset \mathbb{R}^n$, the *convex hull* $\text{conv } \mathcal{X}$ is the set of points $\lambda_1 \mathbf{x}_1 + \dots + \lambda_k \mathbf{x}_k$, where $k \geq 1$, $\{\mathbf{x}_i\} \subset \mathcal{X}$, $\lambda_i \geq 0$ and $\lambda_i = 1$.

- **Theorem (scalar multiplication):** if $\mathcal{C} \subset \mathbb{R}^n$ is convex and $\alpha \in \mathbb{R}$, then $\alpha\mathcal{C} = \{\alpha\mathbf{x} : \mathbf{x} \in \mathcal{C}\}$ is convex.

- **Theorem (vector sum):** If $\mathcal{C}, \mathcal{D} \subset \mathbb{R}^n$ are convex sets, then the set $\mathcal{C} + \mathcal{D} = \{\mathbf{x} + \mathbf{y} : \mathbf{x} \in \mathcal{C}, \mathbf{y} \in \mathcal{D}\}$ is also convex.

- **Theorem (affine transformations):** If $\mathcal{C} \subset \mathbb{R}^n$ is a convex set, $A \in \mathbb{R}^{m \times n}$ is a matrix and $\mathbf{b} \in \mathbb{R}^m$ is a vector, then the set $\{A\mathbf{x} + \mathbf{b} : \mathbf{x} \in \mathcal{C}\}$ is a convex subset of \mathbb{R}^m .

- **Theorem:** If \mathcal{K} is an arbitrary collection of convex sets, then the intersection $\bigcap_{\mathcal{C} \in \mathcal{K}} \mathcal{C}$ is also convex.

- **Miscellaneous definitions**

- **Definition:** A set \mathcal{P} is a *polyhedron* if it is of the form $\mathcal{P} = \{\mathbf{x} \in \mathbb{R}^n : \mathbf{A}\mathbf{x} \leq \mathbf{b}\}$, for some $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a vector $\mathbf{b} \in \mathbb{R}^m$. Polyhedra are convex.
- **Definition:** A set $\mathcal{C} \subset \mathbb{R}^n$ is a *cone* if for all $\mathbf{x} \in \mathcal{C}, \lambda \geq 0, \lambda\mathbf{x} \in \mathcal{C}$. If the cone is convex (ie, if $\lambda_1\mathbf{x}_1 + \lambda_2\mathbf{x}_2 \in \mathcal{C}$ for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{C}, \lambda_1, \lambda_2 \geq 0$), it is a *convex cone*.
 - **Example:** The *conic hull* of the points $\mathcal{X} \subset \mathbb{R}^n$ consists of the points $\lambda_1\mathbf{x}_1 + \dots + \lambda_k\mathbf{x}_k$, where $k \geq 1, \{\mathbf{x}_i\} \subset \mathcal{X}$ and $\boldsymbol{\lambda} \geq \mathbf{0}$. It is a convex cone.
 - **Example:** Given a norm on \mathbb{R}^n , the *norm cone*

$$\{(\mathbf{x}, t) \in \mathbb{R}^{n+1} : \|\mathbf{x}\| \leq t\}$$
 is a convex cone in \mathbb{R}^{n+1} . When the norm in question is the 2-norm, the cone is called the *second-order cone*.
 - **Example:** The set of positive semidefinite matrices

$$S_+^n = \{X \in \mathbb{R}^{n \times n} : X^\top = X, X \succeq 0\}$$
 is a convex cone in $\mathbb{R}^{n \times n}$ called the *positive semidefinite cone*.
- **Hyperplanes and halfspaces**
 - **Definition:** A hyperplane is a set of the form $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{a} \cdot \mathbf{x} = b\}$, where $\mathbf{a} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ is called the *normal vector*. Hyperplanes are affine and therefore convex.
 - **Definition:** A halfspace is a set of the form $\{\mathbf{x} \in \mathbb{R}^n : \mathbf{a} \cdot \mathbf{x} \leq b\}$. Halfspaces are convex but not affine.
- **Definition:** A *norm* is a real valued function $\|\cdot\|$ on \mathbb{R}^n such that
 - $\|\mathbf{x}\| = 0 \Leftrightarrow \mathbf{x} = \mathbf{0}$
 - For all $\mathbf{x} \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}, \|\lambda\mathbf{x}\| = |\lambda|\|\mathbf{x}\|$
 - For all $\mathbf{x}_1, \mathbf{x}_2 \in \mathbb{R}^n, \|\mathbf{x}_1 + \mathbf{x}_2\| \leq \|\mathbf{x}_1\| + \|\mathbf{x}_2\|$

Examples of norms:

- The L2-norm (Euclidean norm): $\|\mathbf{x}\|_2 = \sqrt{\sum_{i=1}^n x_i^2} = \sqrt{\mathbf{x} \cdot \mathbf{x}}$
- $\|\mathbf{x}\|_{\Gamma} = \sqrt{\mathbf{x}^T \Gamma \mathbf{x}}$ (when $\Gamma \succ 0$ and symmetric)
- The p -norm: $\|\mathbf{x}\| = \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}$ for $p \geq 1$
- $\|\mathbf{x}\|_{\infty} = \max\{|x_1|, \dots, |x_n|\}$

Given a norm, the (closed) ball with centre \mathbf{x}_0 and radius $r \geq 0$ is $\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x} - \mathbf{x}_0\| \leq r\}$, and it is convex.


- **Theorem (projection):** Let $\mathcal{C} \subset \mathbb{R}^n$ be a *closed* and non-empty convex set, and consider the Euclidean norm. Fix the vector $\mathbf{x} \in \mathbb{R}^n$. Consider the problem

$$\begin{aligned} \min \quad & \|\mathbf{z} - \mathbf{x}\| \\ \text{s.t.} \quad & \mathbf{z} \in \mathcal{C} \subset \mathbb{R}^n \end{aligned}$$

For every $\mathbf{x} \in \mathbb{R}^n$, the problem has a unique global minimum \mathbf{x}^* called the *projection* of \mathbf{x} onto \mathcal{C} . A vector $\mathbf{x}' \in \mathcal{C}$ is equal to \mathbf{x}^* if and only if

$$(\mathbf{x} - \mathbf{x}') \cdot (\mathbf{z} - \mathbf{x}') \leq 0 \quad \forall \mathbf{z} \in \mathcal{C}$$

Geometrically, the angle between $\mathbf{x}' \rightarrow \mathbf{x}$ and $\mathbf{x}' \rightarrow \mathbf{z}$ must be larger than 90° for all points in the set:



Proof: Existence follows from the fact $\|\mathbf{z} - \mathbf{x}\|$ is coercive and \mathcal{C} is closed. Uniqueness follows because minimizing $\|\mathbf{z} - \mathbf{x}\|$ is equivalent to minimizing $\|\mathbf{z} - \mathbf{x}\|^2 = \mathbf{z} \cdot \mathbf{z} - 2\mathbf{z} \cdot \mathbf{x} + \mathbf{x} \cdot \mathbf{x}$, which is strictly convex.

Now, consider that $\nabla f(\mathbf{x}^*) = 2(\mathbf{x}^* - \mathbf{x})$. By necessary and sufficient conditions for convex optimization problems (derived later), the condition in the theorem must hold.

Application: Suppose we want to approximate $f(\mathbf{x})$ over a set of points $\{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ using $g(\mathbf{x}) = \sum_{\ell=1}^k r_\ell \phi_\ell(\mathbf{x})$, where the ϕ_ℓ are basis functions and \mathbf{r} is a vector of weights. One way to do this is to solve the problem

$$\begin{aligned} \min \quad & \sum_{i=1}^m [f(\mathbf{x}_i) - g(\mathbf{x}_i)]^2 \\ \text{s.t.} \quad & g(\cdot) \text{ is a linear combination of } \{\phi_\ell(\cdot)\} \end{aligned}$$

Consider the matrix $\Phi_{i,\ell} = \phi_\ell(\mathbf{x}_i)$ and the vector \mathbf{y} , $y_i = f(\mathbf{x}_i)$. This problem is equivalent to

$$\begin{aligned} \min \quad & \|\mathbf{y} - \mathbf{z}\| \\ \text{s.t.} \quad & \mathbf{z} \in \{\Phi \mathbf{r} : \mathbf{r} \in \mathbb{R}^k\} \end{aligned}$$

This is a projection problem, and so a unique optimizer exists.

Existence of solutions

- **Theorem – Sufficient Conditions** (Weierstrass): Consider the problem $\min f(\mathbf{x})$ s.t. $x \in \mathcal{C} \subset \mathbb{R}^n$. Then if
 - \mathcal{C} is non-empty
 - f is lower semicontinuous over \mathcal{C}

and one of the following conditions hold:

1. \mathcal{C} is compact
2. \mathcal{C} is closed, and f is coercive
3. There exists a scalar γ such that the level set $\mathcal{C}(\gamma) = \{\mathbf{x} \in \mathcal{C} : f(\mathbf{x}) \leq \gamma\}$ is nonempty and compact.

then the set of optimal minimizing solutions of f is non-empty and compact.

Proof:

- 1 → 3: define $f^* = \inf_{x \in \mathcal{C}} f(x) \in \mathbb{R} \cup \{-\infty\}$ (this always exists). Then, given $\gamma > f^*$, the level set $\{\mathbf{x} \in \mathcal{C} : f(\mathbf{x}) \leq \gamma\}$ must be non-empty. By the continuity of f , it is also closed. Thus, since \mathcal{C} is compact, so is this set.
- 2 → 3: Define $\mathcal{C}(\gamma) = \{\mathbf{x} \in \mathcal{C} : f(\mathbf{x}) \leq \gamma\}$. Since f is coercive, $\mathcal{C}(\gamma)$ is non-empty and bounded for any $\gamma > f^*$. Furthermore, since the domain of f (ie: \mathcal{C}) is

closed, $(-\infty, \gamma]$ is closed and $\mathcal{C}(\gamma) = f^{-1}((-\infty, \gamma])$, $\mathcal{C}(\gamma)$ is also closed. Thus, $\mathcal{C}(\gamma)$ is compact.

- 3: Given a sequence of real numbers $\{\gamma_k\}$ with $\gamma_k \downarrow f^*$, the set of optimal solutions is

$$\mathcal{X}^* = \bigcap_{k=1}^{\infty} \mathcal{C}(\gamma_k)$$

By a theorem stated above, the intersection of these nested, non-empty compact sets is also non-empty, and compact.

Note that the lower semicontinuity of f was only used in proving that $\mathcal{C}(\gamma)$ is *closed*.

- **Example:** consider $\min \frac{1}{2} \mathbf{x}^\top \Gamma \mathbf{x} - \mathbf{b}^\top \mathbf{x}, \mathbf{x} \in \mathbb{R}^n$. If λ is the smallest eigenvalue of Γ , we have $\frac{1}{2} \mathbf{x}^\top \Gamma \mathbf{x} - \mathbf{b}^\top \mathbf{x} \geq \frac{\lambda}{2} \|\mathbf{x}\|^2 - \|\mathbf{b}\| \|\mathbf{x}\|$, which is coercive if $\lambda > 0$. Thus, a solution exists if $\Gamma \succ 0$ (ie: it is positive definite).

Unconstrained local optimality

- In this section, we consider the problem $\min f(\mathbf{x})$ s.t. $\mathbf{x} \in \mathcal{C} \subset \mathbb{R}^n$, but we focus on minima that lie in the interior of \mathcal{C} . In other words, unconstrained minima.
- **Local optimality**
 - **Theorem – necessary conditions:** let $\mathbf{x}^* \in \text{int} \mathcal{C}$ be an unconstrained local minimum. Then

- If f is continuously differentiable in a neighborhood of \mathbf{x}^* , then

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

- If f is twice continuously differentiable in a neighborhood of \mathbf{x}^* , then

$$\nabla^2 f(\mathbf{x}^*) \succeq 0 \quad [\text{Positive semidefinite}]$$

Geometrically, we simply require that the tangent at the said point be horizontal, and also that the tangent *underestimate* the curve in the neighborhood of the stationary point (in other words, the curve should lie above the tangent).

Proof:

- First order: fix $\mathbf{d} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$. By the definition of the gradient,

$$\lim_{\alpha \rightarrow 0} \frac{f(\mathbf{x}^* + \alpha \mathbf{d}) - f(\mathbf{x}^*) - \alpha \nabla f(\mathbf{x}^*) \cdot \mathbf{d}}{\alpha \|\mathbf{d}\|} = 0$$

$$\lim_{\alpha \rightarrow 0} \frac{f(\mathbf{x}^* + \alpha \mathbf{d}) - f(\mathbf{x}^*)}{\alpha} = \nabla f(\mathbf{x}^*) \cdot \mathbf{d}$$

If \mathbf{x}^* is a global optimum, the LHS must be positive for small enough α . Thus, $\nabla f(\mathbf{x}^*) \cdot \mathbf{d} \geq 0$. Since \mathbf{d} is arbitrary, we must have $\nabla f(\mathbf{x}^*) = 0$.

- Second order: fix $\mathbf{d} \in \mathbb{R}^n$. For sufficiently small α :

$$\begin{aligned} f(\mathbf{x}^* + \alpha \mathbf{d}) - f(\mathbf{x}^*) &= \alpha \nabla f(\mathbf{x}^*) \cdot \mathbf{d} + \frac{1}{2} \alpha^2 \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} + o(\alpha^2) \\ &= \frac{1}{2} \alpha^2 \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} + o(\alpha^2) \end{aligned}$$

If \mathbf{x}^* is a global optimum, the LHS must be positive for small enough α , and so

$$\begin{aligned} \frac{1}{2} \alpha^2 \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} + o(\alpha^2) &\geq 0 \\ \frac{1}{2} \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} + \frac{o(\alpha^2)}{\alpha^2} &\geq 0 \end{aligned}$$

Taking limits as $\alpha \rightarrow 0$:

$$\mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} \geq 0$$

Since \mathbf{d} is arbitrary, this leads to our result.

- **Theorem – sufficient conditions**: Consider a point $\mathbf{x}^* \in \text{int } \mathcal{C}$. If f is twice continuously differentiable in a neighborhood of \mathbf{x}^* , and

$$\nabla f(\mathbf{x}^*) = 0 \quad \nabla^2 f(\mathbf{x}^*) \succ 0$$

Then \mathbf{x}^* is a *strict unconstrained local minimum*. The geometric interpretation is as above – the only difference is that we now require a positive definite instead of a positive *semidefinite* matrix.

Proof: Let $\lambda > 0$ be the smallest eigenvalue of $\nabla^2 f(\mathbf{x}^*)$, and let $\mathbf{d} \in N_r(\mathbf{0}) \setminus \{\mathbf{0}\}$

$$\begin{aligned} f(\mathbf{x}^* + \mathbf{d}) - f(\mathbf{x}^*) &= \nabla f(\mathbf{x}^*) \cdot \mathbf{d} + \frac{1}{2} \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} + o(\|\mathbf{d}\|^2) \\ &= \frac{1}{2} \mathbf{d}^T \nabla^2 f(\mathbf{x}^*) \mathbf{d} + o(\|\mathbf{d}\|^2) \\ &\geq \frac{1}{2} \lambda \|\mathbf{d}\|^2 + o(\|\mathbf{d}\|^2) \\ &= \left(\frac{\lambda}{2} + \frac{o(\|\mathbf{d}\|^2)}{\|\mathbf{d}\|^2} \right) \|\mathbf{d}\|^2 \end{aligned}$$

Now, for any $\gamma \in (0, \lambda)$, there exists $\varepsilon \in (0, r]$ such that

$$\frac{\lambda}{2} + \frac{o(\|\mathbf{d}^2\|)}{\|\mathbf{d}\|^2} \geq \frac{\gamma}{2} \quad \forall \mathbf{d} \text{ with } \|\mathbf{d}\| < \varepsilon$$

And this means that

$$f(\mathbf{x}^* + \mathbf{d}) \geq f(\mathbf{x}^*) + \frac{\gamma}{2} \|\mathbf{d}\|^2 > f(\mathbf{x}^*)$$

- *Using the necessary conditions*

- Verify there is a global minimum (using the existence theorem).
- Find the set of possible unconstrained local minima using $\nabla f(\mathbf{x}) = \mathbf{0}$.
- Compare these points with all points on the boundary $\mathcal{C} \setminus \text{int} \mathcal{C}$.
- **Example:** Consider $\min_{\mathbf{x} \in \mathbb{R}^n} \frac{1}{2} \mathbf{x}^\top \Gamma \mathbf{x} - \mathbf{b}^\top \mathbf{x}$ and $\Gamma \succ 0$. By an earlier theorem, global minima must exist. Furthermore, $\mathcal{C} \setminus \text{int} \mathcal{C}$ is empty, and so the global minimum must be an unconstrained local minimum. The first order necessary conditions immediately allow us to characterize that point as $\Gamma \mathbf{x}^* - \mathbf{b} = \mathbf{0}$.

- *Sensitivity analysis*

- Consider the problem $\min f(\mathbf{x}, \mathbf{a})$ s.t. $\mathbf{x} \in \mathbb{R}^n$. We let \mathbf{x}^* be a local optimum, and $f^*(\mathbf{a}) = f(\mathbf{x}^*(\mathbf{a}), \mathbf{a})$. The first-order conditions are

$$\nabla_{\mathbf{x}} f(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) = \mathbf{0}$$

Taking the derivative with respect to \mathbf{a} , we obtain

$$\nabla_{\mathbf{x}^*}(\mathbf{a}) \nabla_{\mathbf{x}\mathbf{x}}^2 f(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) + \nabla_{\mathbf{x}\mathbf{a}}^2 f(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) = \mathbf{0}$$

From this expression, we can obtain expressions for the sensitivity of the optimum, and of the optimal value:

$$\nabla_{\mathbf{x}^*}(\mathbf{a}) = -\nabla_{\mathbf{x}\mathbf{a}}^2 f(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) \left\{ \nabla_{\mathbf{x}\mathbf{x}}^2 f(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) \right\}^{-1}$$

$$\nabla f^*(\mathbf{a}) = \nabla_{\mathbf{a}} f(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) = \nabla_{\mathbf{x}^*}(\mathbf{a}) \nabla_{\mathbf{x}} f(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) + \nabla_{\mathbf{a}} f(\mathbf{x}^*(\mathbf{a}), \mathbf{a}) = \nabla_{\mathbf{a}} f(\mathbf{x}^*(\mathbf{a}), \mathbf{a})$$

- The implicit function theorem tells us when this exists.


Constrained local optimality

- Consider the problem $\min f(\mathbf{x})$ s.t. $\mathbf{x} \in \mathcal{C} \subset \mathbb{R}^n$. We are interested in characterizing local minima that are *not* in $\text{int} \mathcal{C}$. We will assume, though, that f is continuously differentiable in a neighborhood of the point considered.

- **Definition:** The set of *descent directions* of f at \mathbf{x}^* is the set $\mathcal{D}(\mathbf{x}^*) = \{d \in \mathbb{R}^n : \nabla f(\mathbf{x}^*) \cdot d < 0\}$.
- **Definition:** The *tangent cone* $\mathcal{T}(\mathbf{x}^*)$ of the constraint set \mathcal{C} at \mathbf{x}^* is the set of directions $d \in \mathbb{R}^n$ such that either
 - $d = \mathbf{0}$
 - There exists a sequence $\{\mathbf{x}_k\} \subset \mathcal{C}, \mathbf{x}_k \rightarrow \mathbf{x}^*$ such that

$$\frac{\mathbf{x}_k - \mathbf{x}^*}{\|\mathbf{x}_k - \mathbf{x}^*\|} \rightarrow \frac{d}{\|d\|}$$

Geometrically, this is simply the statement that d is tangent to \mathcal{C} if and only if there is some walk we can take in \mathcal{C} that leads us to \mathbf{x}^* and that ends up being in direction d . For example:



- **Theorem – necessary condition:** If \mathbf{x}^* is a local minimum, there is no descent direction in the tangent cone:

$$\mathcal{D}(\mathbf{x}^*) \cap \mathcal{T}(\mathbf{x}^*) = \emptyset$$

Geometrically, the tangent cone contains the “directions in which we can move”. The set of descent directions contains the “directions in which we can improve our objective”. If any direction fulfils both these conditions, then we can clearly improve on our current point.

Proof: Consider $d \in \mathcal{T}(\mathbf{x}^*) \setminus \{\mathbf{0}\}$ and an appropriate sequence $\{\mathbf{x}_k\} \subset \mathcal{C}, \mathbf{x}_k \rightarrow \mathbf{x}^*$. Define

$$\zeta_k = \frac{\mathbf{x}_k - \mathbf{x}^*}{\|\mathbf{x}_k - \mathbf{x}^*\|} - \frac{d}{\|d\|} \rightarrow 0 \qquad \mathbf{d}_k = d + \|d\| \zeta_k \rightarrow d$$

Now, if $\tilde{\mathbf{x}}_k$ is a point on the line segment between \mathbf{x}_k and \mathbf{x}^* , the mean value theorem tells us that

$$f(\mathbf{x}_k) = f(\mathbf{x}^*) + \nabla f(\tilde{\mathbf{x}}_k) \cdot (\mathbf{x}_k - \mathbf{x}^*)$$

Note, however, that we can write

$$(\mathbf{x}_k - \mathbf{x}^*) = \frac{\|\mathbf{x}_k - \mathbf{x}^*\|}{\|\mathbf{d}\|} (\|\mathbf{d}\| \zeta_k + \mathbf{d}) = \frac{\|\mathbf{x}_k - \mathbf{x}^*\|}{\|\mathbf{d}\|} \mathbf{d}_k$$

And so we can re-write the above as

$$f(\mathbf{x}_k) = f(\mathbf{x}^*) + \frac{\|\mathbf{x}_k - \mathbf{x}^*\|}{\|\mathbf{d}\|} \nabla f(\tilde{\mathbf{x}}_k) \cdot \mathbf{d}_k$$

Now, if $\mathbf{d} \in \mathcal{D}(\mathbf{x}^*)$ as well, then $\nabla f(\mathbf{x}^*) \cdot \mathbf{d} < 0$. The strict inequality implies that this is also true in a neighborhood of \mathbf{x}^* , and so for k large enough, we get $f(\mathbf{x}_k) < f(\mathbf{x}^*)$. This contradicts the local minimality of \mathbf{x}^* .

- Unfortunately, \mathcal{T} is hard to characterize algebraically, unless we focus on the particular example where \mathcal{C} is the intersection of equality constraints.

Equality constrained optimization

- Consider the problem $\min f(\mathbf{x})$ s.t. $\mathbf{h}(\mathbf{x}) = \mathbf{0}, \mathbf{x} \in \mathbb{R}^n$ where $h : \mathbb{R}^n \rightarrow \mathbb{R}^m$. We assume the f and h_i are continuously differentiable in a neighborhood of the local minimum.
- In this particular case, we will show we can characterize \mathcal{T} in a simple way. The intuition behind our result is that for any feasible $\mathbf{x}, \mathbf{d} \in \mathbb{R}^n$ and $\alpha > 0$

$$\mathbf{h}(\mathbf{x} + \alpha \mathbf{d}) \approx \mathbf{h}(\mathbf{x}) + \alpha \nabla \mathbf{h}(\mathbf{x})^\top \mathbf{d} = \alpha \nabla \mathbf{h}(\mathbf{x})^\top \mathbf{d}$$

So intuitively, one might expect that any direction for which $\nabla \mathbf{h}(\mathbf{x})^\top \mathbf{d} = \mathbf{0}$ to maintain feasibility. We now formalize this statement...

- **Definition:** the cone of first-order feasible variations at $\mathbf{x}^* \in \mathbb{R}^n$ is the set

$$\mathcal{V}(\mathbf{x}^*) = \left\{ \mathbf{d} \in \mathbb{R}^n : \nabla \mathbf{h}(\mathbf{x}^*)^\top \mathbf{d} = \mathbf{0} \right\} = \left[\ker \nabla \mathbf{h}(\mathbf{x}^*)^\top \right]$$

Note that $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*) \Rightarrow -\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$. As such, $\mathcal{V}(\mathbf{x}^*)$ is actually a *subspace* of \mathbb{R}^n .

- **Definition:** A point $\mathbf{x}^* \in \mathbb{R}^n$ is a *regular point* if it is feasible and the constraint gradients $\nabla h_i(\mathbf{x}^*)$ are linearly independent. In other words, $|\nabla \mathbf{h}(\mathbf{x}^*)| \neq \mathbf{0}$. If $m > n$, no regular points exist, and if $m = 1$, this reduces to $\nabla h_1(\mathbf{x}^*) \neq \mathbf{0}$.
- **Lemma (regularity):** Let \mathbf{x}^* be a regular point. Then $\mathcal{T}(\mathbf{x}^*) = \mathcal{V}(\mathbf{x}^*)$

Proof: This theorem is hard. The intuition behind the proof is

- Consider the curve we would trace if we were sitting at a point \mathbf{x}^* and we started walking forward or backwards *while staying on the constraint* (ie: while keeping the constraint satisfied). We'll start by showing that for any direction

$\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$, there is such a path that starts by walking forward or backward along the direction \mathbf{d} .

- Once we've established this, the result is relatively easy, because the path constitutes a “walk” fully contained in our set \mathcal{C} which eventually ends up being in the direction \mathbf{d} . It's therefore in \mathcal{T} .

And now the painful details! First, let's find the curve in question:

- Begin by choosing $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$. Given a scalar t , consider the curve $\mathbf{x}(t) = \mathbf{x}^* + t\mathbf{d}$. This satisfies our requirement that we be moving either side of \mathbf{x}^* , and that we start by going in direction \mathbf{d} . However, there's no guarantee we stay on the constraints.
- Instead, consider the path $\mathbf{x}(t) = \mathbf{x}^* + t\mathbf{d} + \nabla\mathbf{h}(\mathbf{x}^*)\mathbf{u}(t)$ for some unknown vector $\mathbf{u}(t) \in \mathbb{R}^m$. This seems sensible – we are correcting our path to reflect how \mathbf{h} might change. For $\mathbf{x}(t)$ to be “valid”, we require it to satisfy the m equations

$$\boxed{\mathbf{h}(\mathbf{x}^* + t\mathbf{d} + \nabla\mathbf{h}(\mathbf{x}^*)\mathbf{u}(t)) = \mathbf{0}}$$

For $t = 0$, $\mathbf{u}(0) = \mathbf{0}$ is clearly a solution.

Now, take the gradient of the boxed equation with respect to \mathbf{u} and evaluate it at $(t, \mathbf{u}) = \mathbf{0}$. We get

$$\nabla\mathbf{h}(\mathbf{x}^*)^\top \nabla\mathbf{h}(\mathbf{x}^*)$$

Since the columns of $\nabla\mathbf{h}(\mathbf{x}^*)$ are linearly independent, this matrix is invertible.

The two results above allow us to use the implicit function theorem to deduce that a solution $\mathbf{u}(t)$ to the boxed equation exists for all $t \in (-\tau, \tau)$, for some τ .

Thus, we have managed to find a curve $\mathbf{x}(t)$ that keeps us on the constraints and that is defined over $t \in (-\tau, \tau)$ with $\mathbf{x}(0) = \mathbf{x}^*$ (this implies that the curve represents moving forward and backward from \mathbf{x}^*).

- All we now need to prove is that the initial direction in which we move is \mathbf{d} . To do that, differentiate the boxed equation above with respect to t and evaluate at $t = 0$. We get

$$\left(\mathbf{d}^\top + \dot{\mathbf{u}}(0)^\top \nabla \mathbf{h}(\mathbf{x}^*)^\top\right) \nabla \mathbf{h}(\mathbf{x}^*) = \mathbf{0}$$

But since $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$, we also have that $\mathbf{d}^\top \nabla \mathbf{h}(\mathbf{x}^*) = \mathbf{0}$, and so $\dot{\mathbf{u}}(0) = \mathbf{0}$, and $\dot{\mathbf{x}}(0) = \mathbf{d}$, as required.

- [It will be useful for later to note that if $\mathbf{h}(\cdot)$ is twice continuously differentiable, then so is $\mathbf{x}(\cdot)$. Though I'm not quite sure how to prove that result].

We now have our elusive curve! Let's now prove the theorem.

- $\boxed{\mathcal{V}(\mathbf{x}^*) \subset \mathcal{T}(\mathbf{x}^*)}$: choose $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*) \setminus \{\mathbf{0}\}$ and let $\mathbf{x}(t)$ be the curve discussed above. Take a sequence $t_k \subset (0, \tau), t_k \rightarrow 0$, so that $\mathbf{x}(t_k) \neq \mathbf{x}^*$. Then, by the mean value theorem, there is some $\tilde{t} \in [0, t_k]$ such that

$$\begin{aligned} \mathbf{x}(t_k) - \mathbf{x}(0) &= \dot{\mathbf{x}}(\tilde{t})(t_k - 0) \\ \frac{\mathbf{x}(t_k) - \mathbf{x}^*}{\|\mathbf{x}(t_k) - \mathbf{x}^*\|} &= \frac{\dot{\mathbf{x}}(\tilde{t})}{\|\dot{\mathbf{x}}(\tilde{t})\|} \end{aligned}$$

As $t_k \rightarrow 0$ and therefore $\tilde{t} \rightarrow 0$, this tends to

$$\rightarrow \frac{\dot{\mathbf{x}}(0)}{\|\dot{\mathbf{x}}(0)\|} = \frac{\mathbf{d}}{\|\mathbf{d}\|}$$

So $\mathbf{d} \in \mathcal{T}(\mathbf{x}^*)$.

- $\boxed{\mathcal{T}(\mathbf{x}^*) \subset \mathcal{V}(\mathbf{x}^*)}$: consider $\mathbf{d} \in \mathcal{T}(\mathbf{x}^*) \setminus \{\mathbf{0}\}$ and an associated sequence $\{\mathbf{x}_k\}$ in the feasible set, as defined in the definition of $\mathcal{T}(\mathbf{x}^*)$. By the mean value theorem, there is some $\tilde{\mathbf{x}} \in [\mathbf{x}_k, \mathbf{x}^*]$ such that

$$\mathbf{h}(\mathbf{x}_k) - \mathbf{h}(\mathbf{x}^*) = \nabla \mathbf{h}(\tilde{\mathbf{x}})^\top (\mathbf{x}_k - \mathbf{x}^*)$$

But since \mathbf{x}^* and every \mathbf{x}_k are in the feasible set, $\mathbf{h}(\mathbf{x}_k) = \mathbf{h}(\mathbf{x}^*) = \mathbf{0}$, and

$$0 = \mathbf{h}(\mathbf{x}_k) - \mathbf{h}(\mathbf{x}^*) = \nabla \mathbf{h}(\tilde{\mathbf{x}})^\top (\mathbf{x}_k - \mathbf{x}^*) = \nabla \mathbf{h}(\tilde{\mathbf{x}})^\top (\mathbf{x}_k - \mathbf{x}^*)$$

$$\begin{aligned} \nabla \mathbf{h}(\tilde{\mathbf{x}})^\top \frac{\mathbf{x}_k - \mathbf{x}^*}{\|\mathbf{x}_k - \mathbf{x}^*\|} &= 0 \\ \rightarrow \nabla \mathbf{h}(\mathbf{x}^*)^\top \mathbf{d} &= \mathbf{0} \end{aligned}$$

And so $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$.

{Done! Take a deep breath!}

- **Theorem – necessary condition:** If \mathbf{x}^* is a local minimum that is a regular point, then there is no descent direction that is also a first order feasible variation:

$$\nabla f(\mathbf{x}^*) \cdot \mathbf{d} = 0 \quad \forall \mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$$

Or in other words, we require $\nabla f(\mathbf{x}^*)$ to be in $\mathcal{V}(\mathbf{x}^*)^\perp$:

$$\nabla f(\mathbf{x}^*) \in \mathcal{V}(\mathbf{x}^*)^\perp = \left[\ker \nabla \mathbf{h}(\mathbf{x}^*)^\top \right]^\perp = \text{im } \nabla \mathbf{h}(\mathbf{x}^*)$$

Or in other words, there exists $\boldsymbol{\lambda} \in \mathbb{R}^m$ such that $\nabla f(\mathbf{x}^*) = \nabla \mathbf{h}(\mathbf{x}^*) \boldsymbol{\lambda}$.

Proof: Since \mathbf{x}^* is a local minimum, $\mathcal{D}(\mathbf{x}^*) \cap \mathcal{T}(\mathbf{x}^*) = \emptyset$, and since \mathbf{x}^* is regular, $\mathcal{D}(\mathbf{x}^*) \cap \mathcal{V}(\mathbf{x}^*) = \emptyset$.

Now, assume $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$ – by what we have such said, $\mathbf{d} \notin \mathcal{D}(\mathbf{x}^*)$, and so $\nabla f(\mathbf{x}^*) \cdot \mathbf{d} \geq 0$. However, since we also have $-\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$, we must have $\nabla f(\mathbf{x}^*) \cdot \mathbf{d} = 0$.

For the last part of the theorem, note that $\text{im } A = (\ker A^\top)^\perp$, as proved in the introductory section of these notes.


- The last part of the previous theorem is important, because it provides a “simple” way to characterize the tangent cone, and a “recipe” to find optimal points. This can be formalized further using...
- *...Lagrange Multipliers*
 - **Theorem – necessary conditions:** If \mathbf{x}^* is a local minimum that is a regular point, then there exists a unique vector $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ called a *Lagrange multiplier* such that

$$\nabla f(\mathbf{x}^*) + \boldsymbol{\lambda}^{*\top} \nabla \mathbf{h}(\mathbf{x}^*) = \nabla f(\mathbf{x}^*) + \sum_{i=1}^m \lambda_i^* \nabla h_i(\mathbf{x}^*) = \mathbf{0}$$

In addition, if f and \mathbf{h} are twice continuously differentiable

$$\mathbf{d}^\top \left(\nabla^2 f(\mathbf{x}^*) + \sum_{i=1}^m \lambda_i^* \nabla^2 h_i(\mathbf{x}^*) \right) \mathbf{d} \geq 0 \quad \forall \mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$$

There is an interesting geometrical interpretation of the first-order condition. It effectively states that $\nabla f(\mathbf{x}^*)$ [the direction in which we might increase our objective] must be a linear combination of the $\nabla h_i(\mathbf{x}^*)$ [the *perpendicular* to the constraints $h_i(\mathbf{x}^*) = 0$]. Since we cannot move along any of those perpendiculars without leaving the constraints, we clearly cannot move along $\nabla f(\mathbf{x}^*)$. Here is an example, in which $\nabla f(\mathbf{x})$ is constant:



Proof: The existence of λ^* is simply a restatement of the previous theorem. The uniqueness of λ^* follows from the fact that the columns of $\nabla h(\mathbf{x}^*)$ are linearly independent.

For the second-order condition, consider a $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$, and use the first part of the regularity lemma to define a path $\mathbf{x}(t)$ either side of \mathbf{x}^* , which always stays on the constraints and such that $\dot{\mathbf{x}}(0) = \mathbf{d}$. Now, define $g(t) = f(\mathbf{x}(t))$ and take a double derivative

$$\ddot{g}(t) = \dot{\mathbf{x}}(t)^\top \nabla^2 f(\mathbf{x}(t)) \dot{\mathbf{x}}(t) + \ddot{\mathbf{x}}(t)^\top \nabla f(\mathbf{x}(t))$$

Since all points $\mathbf{x}(t)$ satisfy the constraints of the problem, and \mathbf{x}^* is a local minimum, $t = 0$ must be an unconstrained *local* minimum of $g(t)$. Thus

$$\ddot{g}(0) = \mathbf{d}^\top \nabla^2 f(\mathbf{x}^*) \mathbf{d} + \ddot{\mathbf{x}}(0)^\top \nabla f(\mathbf{x}^*) \geq 0$$

Finally, consider $\ell(t) = \lambda^{*\top} h(\mathbf{x}(t)) = 0$ and differentiate it twice, to get

$$\ddot{\ell}(0) = \mathbf{d}^\top \left(\sum_{i=1}^m \lambda_i \nabla^2 h_i(\mathbf{x}^*) \right) \mathbf{d} + \ddot{\mathbf{x}}(0)^\top \nabla h(\mathbf{x}^*) \lambda^* = 0$$

Finally, add the last two equations, and apply the first order condition.

- o We define the *Lagrangian* as

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}) = f(\mathbf{x}) + \boldsymbol{\lambda} \cdot \mathbf{h}(\mathbf{x})$$

The first and second order conditions then reduce to

$$\begin{aligned} \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*) &= 0 \\ \mathbf{d}^\top \nabla_{\mathbf{xx}}^2 \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*) \mathbf{d} &\geq 0 \quad \forall \mathbf{d} \in \mathcal{V}(\mathbf{x}^*) \end{aligned}$$

And the feasibility condition is given by

$$\nabla_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*) = 0$$

This formulation of the first order conditions also has an interesting interpretation. We effectively allow the constraint $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ to be broken, *but* we associate a cost $\boldsymbol{\lambda}$ with breaking it. We then apply unconstrained first-order conditions to the resulting problem.

- As such, we have the following simple recipe to solve an equality-constrained problem (assuming f and h are continuously differentiable on \mathbb{R}^n):
 - Check that global minima exist
 - Find the set of $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ satisfying $\nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*) = 0$ and $\nabla_{\boldsymbol{\lambda}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*) = 0$.
 - Find the set of non-regular points
 - The global minima are amongst these points.

A few examples:

- $\min f(\mathbf{x}) = x_1 + x_2$ s.t. $h(\mathbf{x}) = x_1^2 + x_2^2 - 2 = 0, \mathbf{x} \in \mathbb{R}^2$. In this case, note that $\nabla h = (2x_1, 2x_2)$; thus, provided the constraints are met, all points will be regular. In this case, $\mathcal{L}(\mathbf{x}, \lambda) = x_1 + x_2 + \lambda(x_1^2 + x_2^2 - 2)$, and since all points are regular, we simply need to consider all points satisfying the first-order conditions and choose the smallest one.
- Consider the problem

$$\begin{aligned} \min \quad & f(\mathbf{x}) = x_1 + x_2 \\ \text{s.t.} \quad & h_1(\mathbf{x}) = (x_1 - 1)^2 + x_2^2 - 1 = 0 \\ & h_2(\mathbf{x}) = (x_1 - 2)^2 + x_2^2 - 4 = 0 \\ & \mathbf{x} \in \mathbb{R}^2 \end{aligned}$$

The only feasible point is $(0, 0)$ – and so it is the global minimum. However, it is not a regular point. And indeed, we find that there is no solution to the first-order necessary conditions.

- **Constraint qualification:** In summary, what we have proved so far is

$$\begin{aligned} & \mathbf{x}^* \text{ local minimum} \\ & \Rightarrow \mathcal{D}(\mathbf{x}^*) \cap \mathcal{T}(\mathbf{x}^*) = \emptyset \\ & \Rightarrow \mathcal{D}(\mathbf{x}^*) \cap \mathcal{V}(\mathbf{x}^*) = \emptyset \\ & \Rightarrow \text{Lagrange multipliers exist} \end{aligned}$$

In proving this sequence of statements, we made use of the fact that \mathbf{x}^* was a regular point. However, it is possible to prove the existence of Lagrange

multipliers under weaker assumptions called *constraint qualifications*. If the constraints are linear, for example, Lagrange multipliers are guaranteed to exist. The weakest form of constraint qualification is *quasiregularity*, which requires that $\mathcal{V}(\mathbf{x}^*) = \mathcal{T}(\mathbf{x}^*)$.

- **Theorem – Sufficient Conditions:** Assume that f and \mathbf{h} are both twice continuously differentiable, and that $\mathbf{x}^* \in \mathbb{R}^n$ and $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ satisfy

$$\begin{aligned}\nabla_{\mathbf{x}} L(\mathbf{x}^*, \boldsymbol{\lambda}^*) &= \mathbf{0} \\ \nabla_{\boldsymbol{\lambda}} L(\mathbf{x}^*, \boldsymbol{\lambda}^*) &= \mathbf{0} \\ \mathbf{d}^\top \nabla_{\mathbf{xx}}^2 L(\mathbf{x}^*, \boldsymbol{\lambda}^*) \mathbf{d} &> 0 \quad \forall \mathbf{d} \in \mathcal{V}(\mathbf{x}^*) \setminus \{\mathbf{0}\}\end{aligned}$$

Then \mathbf{x}^* is a strict local minimum.

Proof: The second condition above implies that \mathbf{x}^* is clearly feasible. Suppose it is not a strict local minimum; then there exists a sequence $\{\mathbf{x}_k\} \subset \mathbb{R}^n$ such that $\mathbf{x}_k \neq \mathbf{x}^*$ and $\mathbf{x}_k \rightarrow \mathbf{x}^*$ which lies entirely in the feasible region of the problem [ie: $\mathbf{h}(\mathbf{x}_k) = \mathbf{0}$] and $f(\mathbf{x}_k) \leq f(\mathbf{x}^*)$. We define, for some \mathbf{d}

$$\mathbf{d}_k = \frac{\mathbf{x}_k - \mathbf{x}^*}{\|\mathbf{x}_k - \mathbf{x}^*\|} \rightarrow \mathbf{d} \quad \delta_k = \|\mathbf{x}_k - \mathbf{x}^*\| \rightarrow 0$$

Now, by the mean value theorem, there exists $\tilde{\mathbf{x}} \in [\mathbf{x}^*, \mathbf{x}_k]$ with

$$\mathbf{h}(\mathbf{x}_k) - \mathbf{h}(\mathbf{x}^*) = \nabla \mathbf{h}(\tilde{\mathbf{x}})^\top (\mathbf{x}_k - \mathbf{x}^*) = \nabla \mathbf{h}(\tilde{\mathbf{x}})^\top (\delta_k \mathbf{d}_k)$$

But since \mathbf{x}^* and \mathbf{x}_k are feasible, $\mathbf{h}(\mathbf{x}_k) = \mathbf{h}(\mathbf{x}^*) = \mathbf{0}$, so.

$$\nabla \mathbf{h}(\tilde{\mathbf{x}})^\top \mathbf{d}_k = 0$$

Taking the limit as $k \rightarrow \infty$, we get $\nabla \mathbf{h}(\mathbf{x}^*)^\top \mathbf{d} = 0$, and so $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*)$.

Now, we know that

$$\mathbf{h}(\mathbf{x}_k) = \mathbf{0} \quad \mathbf{f}(\mathbf{x}_k) - f(\mathbf{x}^*) \leq 0$$

Using a second order Taylor expansion (with remainder) with some set of $\hat{\mathbf{x}}^i \in [\mathbf{x}_k, \mathbf{x}^*]$, we can re-write these as

$$\begin{aligned}h_i(\mathbf{x}_k) &= h_i(\mathbf{x}^*) + \delta_k \nabla h_i(\mathbf{x}^*) \cdot \mathbf{d}_k + \frac{1}{2} \delta_k^2 \mathbf{d}_k^\top \nabla^2 h_i(\hat{\mathbf{x}}^i) \mathbf{d}_k = 0 \\ f(\mathbf{x}_k) - f(\mathbf{x}^*) &= \boxed{\delta_k \nabla f(\mathbf{x}^*) \cdot \mathbf{d}_k + \frac{1}{2} \delta_k^2 \mathbf{d}_k^\top \nabla^2 f(\hat{\mathbf{x}}^0) \mathbf{d}_k} \leq 0\end{aligned}$$

We can modify the first set of equations slightly by remembering that $\mathbf{h}(\mathbf{x}^*) = \mathbf{0}$, and multiplying both sides of the equation by λ_i^* . This gives

$$h_i(\mathbf{x}_k) = \boxed{\delta_k \lambda_i^* \nabla h_i(\mathbf{x}^*) \cdot \mathbf{d}_k + \frac{1}{2} \delta_k^2 \mathbf{d}_k^\top \lambda_i^* \nabla^2 h_i(\hat{\mathbf{x}}^i) \mathbf{d}_k = 0}$$

Adding these $m + 1$ equations, we get

$$\begin{aligned} \delta_k \left(\nabla f(\mathbf{x}^*) + \sum_{i=1}^m \lambda_i^* \nabla h_i(\mathbf{x}^*) \right) \cdot \mathbf{d}_k + \frac{1}{2} \delta_k^2 \mathbf{d}_k^\top \left(\nabla^2 f(\hat{\mathbf{x}}^0) + \sum_{i=1}^m \lambda_i^* \nabla^2 h_i(\hat{\mathbf{x}}^i) \right) \mathbf{d}_k &\leq 0 \\ \delta_k \nabla_{\mathbf{x}} L(\mathbf{x}^*, \boldsymbol{\lambda}^*) \cdot \mathbf{d}_k + \frac{1}{2} \delta_k^2 \mathbf{d}_k^\top \left(\nabla^2 f(\hat{\mathbf{x}}^0) + \sum_{i=1}^m \lambda_i^* \nabla^2 h_i(\hat{\mathbf{x}}^i) \right) \mathbf{d}_k &\leq 0 \end{aligned}$$

Noting that, by the first order conditions, $\nabla_{\mathbf{x}} L(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ and then dividing by $\frac{1}{2} \delta_k^2$ and taking the limit as $k \rightarrow \infty$, this becomes

$$\mathbf{d}^\top \left(\nabla^2 f(\mathbf{x}^*) + \sum_{i=1}^m \lambda_i^* \nabla^2 h_i(\mathbf{x}^*) \right) \mathbf{d} \leq 0$$

But since $\mathbf{d} \in \mathcal{V}(\mathbf{x}^*) \setminus \{\mathbf{0}\}$, this violates our assumed second order condition.

- We now consider an application of these conditions. Consider the program

$$\min_{\mathbf{x} \in \mathbb{R}^n} \sigma^2 = \mathbf{x}^\top \Gamma \mathbf{x} \text{ s.t. } \mathbf{1}^\top \mathbf{x} = 1, \boldsymbol{\mu}^\top \mathbf{x} = \bar{\mu}$$

which might represent minimizing the variance in a portfolio while keeping total sales equal to 1 unit, and keeping the expected return equal to a certain value $\bar{\mu}$.

The first-order conditions give

$$2\Gamma \mathbf{x}^* + \lambda_1^* \mathbf{1} + \lambda_2^* \boldsymbol{\mu} = \mathbf{0} \quad \mathbf{1}^\top \mathbf{x}^* = 1, \boldsymbol{\mu}^\top \mathbf{x}^* = \bar{\mu}$$

From the first equation, we obtain

$$\begin{aligned} \mathbf{x}^* &= -\frac{1}{2} \Gamma^{-1} \mathbf{1} \lambda_1^* - \frac{1}{2} \Gamma^{-1} \boldsymbol{\mu} \lambda_2^* \\ \mathbf{1}^\top \mathbf{x}^* &= -\frac{1}{2} \mathbf{1}^\top \Gamma^{-1} \mathbf{1} \lambda_1^* - \frac{1}{2} \mathbf{1}^\top \Gamma^{-1} \boldsymbol{\mu} \lambda_2^* = 1 \\ \boldsymbol{\mu}^\top \mathbf{x}^* &= -\frac{1}{2} \boldsymbol{\mu}^\top \Gamma^{-1} \mathbf{1} \lambda_1^* - \frac{1}{2} \boldsymbol{\mu}^\top \Gamma^{-1} \boldsymbol{\mu} \lambda_2^* = \bar{\mu} \end{aligned}$$

The last two equations are a system of equations for $(\lambda_1^*, \lambda_2^*)$:

$$-\frac{1}{2} \begin{pmatrix} \mathbf{1}^\top \Gamma^{-1} \mathbf{1} & \mathbf{1}^\top \Gamma^{-1} \boldsymbol{\mu} \\ \boldsymbol{\mu}^\top \Gamma^{-1} \mathbf{1} & \boldsymbol{\mu}^\top \Gamma^{-1} \boldsymbol{\mu} \end{pmatrix} \begin{pmatrix} \lambda_1^* \\ \lambda_2^* \end{pmatrix} = \begin{pmatrix} 1 \\ \bar{\mu} \end{pmatrix}$$

this system is nonsingular provided that $\Gamma \succ 0$ and $\mathbf{1}$ and $\boldsymbol{\mu}$ are linearly independent. We then get

$$\begin{pmatrix} \lambda_1^* \\ \lambda_2^* \end{pmatrix} = \begin{pmatrix} \eta_1 + \zeta_1 \bar{\mu} \\ \eta_2 + \zeta_2 \bar{\mu} \end{pmatrix}$$


Where the constants depend on Γ and $\boldsymbol{\mu}$. Now, using the first equation in the FOCs, we obtain, for some vectors \mathbf{v} and \mathbf{w}

$$\mathbf{x}^* = \bar{\mu} \mathbf{v} + \mathbf{w}$$

which finally gives, for some constants α, β, γ depending on Γ and $\boldsymbol{\mu}$

$$\sigma^2 = (\bar{\boldsymbol{\mu}}\mathbf{v} + \mathbf{w})^\top \Gamma (\bar{\boldsymbol{\mu}}\mathbf{v} + \mathbf{w}) = (\alpha\bar{\mu} + \beta)^2 + \gamma$$

This implies that of all the possible asset combinations, those that provide the maximum expected returns for a particular variance lie on a parabola that looks like this:



The upper part of this parabola is called the *efficient frontier*.

- **Sensitivity analysis**

- We now look at what happens when we slightly vary our constraint.
- **Theorem:** Consider a family of equality-constrained problems $\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x})$ s.t. $\mathbf{h}(\mathbf{x}) = \mathbf{u}$. Suppose there exists a local pair $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ satisfying the second-order sufficient conditions when $\mathbf{u} = \mathbf{0}$, and \mathbf{x}^* is regular. Then there exists a neighborhood $N_\epsilon(0) \subset \mathbb{R}^m$ of $\mathbf{u} = \mathbf{0}$ and a function $\mathbf{x}^*(\cdot)$ defined on that neighborhood such that

- $\mathbf{x}^*(\mathbf{0}) = \mathbf{x}^*$, and for each $\mathbf{u} \in N_\epsilon(0)$, $\mathbf{x}^*(\mathbf{u})$ is a strict local minimum.
- $\mathbf{x}^*(\mathbf{u})$ is continuously differentiable.
- If $p(\mathbf{u}) = f(\mathbf{x}^*(\mathbf{u}))$, then $\nabla p(\mathbf{0}) = -\boldsymbol{\lambda}^*$.

The last part of the theorem is the interesting one – it tells us that changing \mathbf{u} by an infinitesimal amount around our solution point will lead in a change of $-\boldsymbol{\lambda}^*$ in the objective.

Proof: Consider, for $\mathbf{u} \in \mathbb{R}^m$, the following system of equations in $(\mathbf{x}, \boldsymbol{\lambda})$, representing the first-order sufficient conditions:

$$\begin{aligned} \nabla f(\mathbf{x}) + \nabla \mathbf{h}(\mathbf{x})\boldsymbol{\lambda} &= \mathbf{0} \\ \mathbf{h}(\mathbf{x}) - \mathbf{u} &= \mathbf{0} \end{aligned}$$

Let's carefully calculate the gradient of this system

$$\begin{bmatrix} \nabla_x \{ \nabla f(\mathbf{x}) + \nabla \mathbf{h}(\mathbf{x}) \boldsymbol{\lambda} \} & \nabla_x \{ \mathbf{h}(\mathbf{x}) - \mathbf{u} \} \\ \nabla_\lambda \{ \nabla f(\mathbf{x}) + \nabla \mathbf{h}(\mathbf{x}) \boldsymbol{\lambda} \} & \nabla_\lambda \{ \mathbf{h}(\mathbf{x}) - \mathbf{u} \} \end{bmatrix} = \begin{bmatrix} \nabla_{xx} L(\mathbf{x}, \boldsymbol{\lambda}) & \nabla_x \mathbf{h}(\mathbf{x}) \\ \nabla \mathbf{h}(\mathbf{x})^\top & 0 \end{bmatrix}$$

Clearly, the system above has a solution at $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$, and at that point, the gradient is given by

$$J = \begin{bmatrix} \nabla_{xx} L(\mathbf{x}^*, \boldsymbol{\lambda}^*) & \nabla_x \mathbf{h}(\mathbf{x}^*) \\ \nabla \mathbf{h}(\mathbf{x}^*)^\top & 0 \end{bmatrix}$$

Now, if this matrix were singular, there would exist some nonzero vector $\mathbf{x} = (\mathbf{y}^\top, \mathbf{z}^\top)^\top$ such that $J\mathbf{x} = \mathbf{0}$. This implies that

$$\begin{aligned} \nabla_{xx}^2 L(\mathbf{x}^*, \boldsymbol{\lambda}^*) \mathbf{y} + \nabla_x \mathbf{h}(\mathbf{x}^*) \mathbf{z} &= \mathbf{0} \\ \nabla \mathbf{h}(\mathbf{x}^*)^\top \mathbf{y} &= \mathbf{0} \end{aligned}$$

The second equation implies that $\mathbf{y} \in \mathcal{V}(\mathbf{x}^*)$. Multiplying the first equation by \mathbf{y}^\top and using the second equation implies that we would then have to have $\mathbf{y}^\top \nabla_{xx}^2 L(\mathbf{x}^*, \boldsymbol{\lambda}^*) \mathbf{y} = \mathbf{0}$ for some $\mathbf{y} \in \mathcal{V}(\mathbf{x}^*)$; this violates the second-order sufficient conditions. Thus, our matrix is nonsingular.

Using the implicit function theorem, this implies that we can define $(\mathbf{x}^*(\mathbf{u}), \boldsymbol{\lambda}^*(\mathbf{u}))$ satisfying first-order conditions for all \mathbf{u} in some $N_\varepsilon(\mathbf{0})$. Second-order conditions follow, for \mathbf{u} sufficiently close to $\mathbf{0}$, by continuity assumptions.

For the last part of our theorem, consider the FOCs in terms of this $(\mathbf{x}^*(\mathbf{u}), \boldsymbol{\lambda}^*(\mathbf{u}))$, and multiply them by $\nabla \mathbf{x}^*(\mathbf{u})$:

$$\begin{aligned} \nabla \mathbf{x}^*(\mathbf{u}) \left[\nabla f(\mathbf{x}^*(\mathbf{u})) \right] + \nabla \mathbf{x}^*(\mathbf{u}) \left[\nabla \mathbf{h}(\mathbf{x}^*(\mathbf{u})) \right] \boldsymbol{\lambda}^*(\mathbf{u}) &= \mathbf{0} \\ \nabla_u f(\mathbf{x}^*(\mathbf{u})) + \nabla \mathbf{x}^*(\mathbf{u}) \left[\nabla \mathbf{h}(\mathbf{x}^*(\mathbf{u})) \right] \boldsymbol{\lambda}^*(\mathbf{u}) &= \mathbf{0} \end{aligned}$$

Note also that differentiating $\mathbf{h}(\mathbf{x}^*(\mathbf{u})) = \mathbf{u}$ with respect to \mathbf{u} yields $\nabla_u \mathbf{x}^*(\mathbf{u}) \nabla_x \mathbf{h}(\mathbf{x}^*(\mathbf{u})) = I$. Using that result, we can re-write the previous equation as

$$\begin{aligned} \nabla_u f(\mathbf{x}^*(\mathbf{u})) + I \boldsymbol{\lambda}^*(\mathbf{u}) &= \mathbf{0} \\ \nabla_u f(\mathbf{x}^*(\mathbf{u})) &= -\boldsymbol{\lambda}^*(\mathbf{u}) \end{aligned}$$

As required.

Inequality constrained optimization

- Consider the program

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \text{ s.t. } \mathbf{h}(\mathbf{x}) = \mathbf{0}, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}$$

Where $f: \mathbb{R}^n \rightarrow \mathbb{R}$, $\mathbf{h}: \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $\mathbf{g}: \mathbb{R}^n \rightarrow \mathbb{R}^r$. Note also that when we say $\mathbf{x} \leq \mathbf{0}$, we mean that every component of \mathbf{x} is less than or equal to 0.

We also assume that f , \mathbf{h} and \mathbf{g} are continuously differentiable on \mathbb{R}^n , though the necessary and sufficient conditions also hold if these are defined and continuously differentiable only in a neighborhood of the local minimum.

- Definition:** Given a feasible point $\mathbf{x}^* \in \mathbb{R}^n$, the set of active inequality constraints $\mathcal{A}(\mathbf{x}^*)$ is defined as

$$\mathcal{A}(\mathbf{x}^*) = \{j : g_j(\mathbf{x}^*) = 0\} \subseteq \{1, \dots, r\}$$

- Lemma:** Let \mathbf{x}^* be a local minimum for the inequality constrained program above (IneqCP). Then it is also a local minimum for the following equality constrained program (EqCP):

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \text{ s.t. } \mathbf{h}(\mathbf{x}) = \mathbf{0}, g_j(\mathbf{x}) = 0 \quad \forall j \in \mathcal{A}(\mathbf{x}^*)$$

Proof: Suppose \mathbf{x}^* is not a local optimum for EqCP. Then there exists a sequence of points $\{\mathbf{x}_k\}$ feasible for EqCP such that $\mathbf{x}_k \rightarrow \mathbf{x}^*$ and $f(\mathbf{x}_k) < f(\mathbf{x}^*)$.

Since \mathbf{g} is continuous, we have $\mathbf{g}(\mathbf{x}_k) \rightarrow \mathbf{g}(\mathbf{x}^*)$. In particular, if $j \notin \mathcal{A}(\mathbf{x}^*)$, $g_j(\mathbf{x}_k) \rightarrow g_j(\mathbf{x}^*) < 0$. Thus, for sufficiently large k , \mathbf{x}_k is feasible for IneqCP. This contradicts the local optimality of \mathbf{x}^* for IneqCP.

- Definition:** Consider the inequality constrained program above. A point $\mathbf{x}^* \in \mathbb{R}^n$ is a *regular point* if it is feasible and if the set of constraint gradients

$$\{\nabla h_i(\mathbf{x}^*) : 1 \leq i \leq m\} \cup \{\nabla g_j(\mathbf{x}^*) : j \in \mathcal{A}(\mathbf{x}^*)\}$$

are linearly independent.

- Definition:** The cone $\mathcal{V}^{\text{Eq}}(\mathbf{x}^*)$ at the point $\mathbf{x}^* \in \mathbb{R}^n$ is the set of vectors $\mathbf{d} \in \mathbb{R}^n$ such that

$$\begin{aligned} \mathbf{d} \cdot \nabla h_i(\mathbf{x}^*) &= 0 & \forall 1 \leq i \leq m \\ \mathbf{d} \cdot \nabla g_j(\mathbf{x}^*) &= 0 & \forall j \in \mathcal{A}(\mathbf{x}^*) \end{aligned}$$

Intuitively, it is the set of points in which we can move while keeping the equality constraints and the active inequality constraints satisfied.

- **Theorem (Karush-Kuhn-Tucker Necessary Conditions):** If \mathbf{x}^* is a local minimum that is a regular point, then there exists unique Lagrange multiplier vectors $\boldsymbol{\lambda}^* \in \mathbb{R}^m$ and $\boldsymbol{\mu}^* \in \mathbb{R}^r$ such that

$$\begin{aligned} \nabla f(\mathbf{x}^*) + \sum_{i=1}^m \lambda_i^* \nabla h_i(\mathbf{x}^*) + \sum_{j=1}^r \mu_j^* \nabla g_j(\mathbf{x}^*) &= \mathbf{0} \\ \boldsymbol{\mu}^* &\geq \mathbf{0} \\ \mu_j^* &= 0 \quad \forall j \notin \mathcal{A}(\mathbf{x}^*) \end{aligned}$$

In addition, if f , \mathbf{h} and \mathbf{g} are twice continuously differentiable, then

$$\mathbf{d}^\top \left(\nabla^2 f(\mathbf{x}^*) + \sum_{i=1}^m \lambda_i^* \nabla^2 h_i(\mathbf{x}^*) + \sum_{j=1}^r \mu_j^* \nabla^2 g_j(\mathbf{x}^*) \right) \mathbf{d} \geq 0 \quad \forall \mathbf{d} \in \mathcal{V}^{\text{EQ}}(\mathbf{x}^*)$$

There is an interesting geometrical explanation of the first-order conditions. First, consider that

- Because the h are equalities, we cannot move along $\nabla h_i(\mathbf{x})$ [the perpendiculars to the equality constraints] in *any* direction (either forwards or backwards).
- Because the active constraints are \leq inequalities, we can move along $\nabla g_i(\mathbf{x}), i \in \mathcal{A}(\mathbf{x})$ [the perpendiculars to the active inequality constraints] only in the *negative* direction. Moving along $+\nabla g_i(\mathbf{x})$ would leave the feasible region.
- Inactive constraints are not tight, and so we can move both forwards and backwards along $\nabla g_i(\mathbf{x}), i \notin \mathcal{A}(\mathbf{x})$.

All the first-order condition states is that there is no way to move in such a way that we remain in the feasible region and decrease $f(\mathbf{x}^*)$. Mathematically, the conditions states that $-\nabla f(\mathbf{x}^*)$ [the direction which would decrease $f(\mathbf{x}^*)$] must be composed entirely of

- A linear combination of the perpendiculars to the equality constraints (along which we cannot move).
- A *positive* linear combination of the perpendiculars to the active inequality constraints (along which we cannot move).
- [No perpendiculars to the inactive inequality constraints, because we *could* move along those].

Proof: Everything follows from applying Lagrange multipliers to the equality constrained program, except for the fact that $\mu_j^* \geq 0$ for all $j \in \mathcal{A}(\mathbf{x}^*)$.

To see why this is the case, assume that this does *not* hold for some $j \in \mathcal{A}(\mathbf{x}^*)$. Then let $\mathcal{C}_j \subset \mathbb{R}^n$ be the set of points feasible for all *other* active constraints, and let $\mathcal{V}_j^{\text{EQ}}(\mathbf{x}^*)$ be the corresponding cone of first-order feasible directions:

$$\begin{aligned} \mathcal{C}_j &= \left\{ \mathbf{x} : \mathbf{h}(\mathbf{x}) = \mathbf{0}, g_\ell(\mathbf{x}) = 0 \quad \forall \ell \in \mathcal{A}(\mathbf{x}^*) \setminus \{j\} \right\} \\ \mathcal{V}_j^{\text{EQ}}(\mathbf{x}^*) &= \left\{ \mathbf{d} : \mathbf{d} \cdot \nabla \mathbf{h}(\mathbf{x}^*) = \mathbf{0}, \mathbf{d} \cdot \nabla g_\ell(\mathbf{x}^*) = 0 \quad \forall \ell \in \mathcal{A}(\mathbf{x}^*) \setminus \{j\} \right\} \end{aligned}$$

By regularity, there must exist a $\mathbf{d} \in \mathcal{V}_j^{\text{EQ}}(\mathbf{x}^*)$ with $\mathbf{d} \cdot \nabla g_j(\mathbf{x}^*) < 0$ (if it were the case that $\mathbf{d} \cdot \nabla g_j(\mathbf{x}^*) = 0 \quad \forall \mathbf{d} \in \mathcal{V}_j^{\text{EQ}}(\mathbf{x}^*)$, then the point would not be regular, and if $\mathbf{d}' \cdot \nabla g_j(\mathbf{x}^*) > 0$ for some $\mathbf{d}' \in \mathcal{V}_j^{\text{EQ}}(\mathbf{x}^*)$, then just take $\mathbf{d} = -\mathbf{d}' \in \mathcal{V}_j^{\text{EQ}}(\mathbf{x}^*)$).

Then, by the regularity lemma, there exists a curve $\mathbf{x}(t) \in \mathcal{C}_j$ with $\mathbf{x}(0) = \mathbf{x}^*$, $\dot{\mathbf{x}}(0) = \mathbf{d}$ and such that $\mathbf{x}(t)$ is feasible for small t . Now, if we let $\ell(t) = f[\mathbf{x}(t)]$, we find that

$$\dot{\ell}(t) = \dot{\mathbf{x}}(t)^\top \nabla f[\mathbf{x}(t)]$$

Evaluating at $t = 0$

$$\begin{aligned} \dot{\ell}(0) &= \mathbf{d}^\top \nabla f(\mathbf{x}^*) \\ &= -\mathbf{d}^\top \left\{ \sum_{i=1}^m \lambda_i^* h_i(\mathbf{x}^*) + \sum_{j=1}^m \mu_j^* g_j(\mathbf{x}^*) \right\} \\ &= -\mathbf{d}^\top \nabla g_j(\mathbf{x}^*) \mu_j^* \\ &< 0 \end{aligned}$$

This contradicts the local optimality of \mathbf{x}^* .

- Practically, we typically define a Lagrangian

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = f(\mathbf{x}) + \boldsymbol{\lambda} \cdot \mathbf{h}(\mathbf{x}) + \boldsymbol{\mu} \cdot \mathbf{g}(\mathbf{x})$$

The first-order conditions then reduce to

$$\begin{aligned} \nabla_x \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) &= \mathbf{0} & \mathbf{h}(\mathbf{x}^*) &= \mathbf{0} & \mathbf{g}(\mathbf{x}^*) &\leq \mathbf{0} \\ & & \boldsymbol{\mu}^* &\geq \mathbf{0} \\ \mu_j^* g_j(\mathbf{x}^*) &= 0 & \forall 1 \leq j \leq r \end{aligned}$$

For any given problem, then, we

- Prove global minima exist
- Find the set of $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ satisfying these FOC

- Find the set of non-regular points
- Choose the point with the lowest objective.

For example, consider the problem

$$\min_{\mathbf{x} \in \mathbb{R}^3} \frac{1}{2}(x_1^2 + x_2^2 + x_3^2) \text{ s.t. } x_1 + x_2 + x_3 \leq -3$$

The objective and constraints are continuously differentiable, and minima exist (by coerciveness). The Lagrangian is

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \frac{1}{2}(x_1^2 + x_2^2 + x_3^2) + \mu(x_1 + x_2 + x_3 + 3)$$

The first-order conditions are

$$\begin{aligned} \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = \mathbf{0} &\Rightarrow x_1^* + \mu^* = x_2^* + \mu^* = x_3^* + \mu^* = 0 \\ \mathbf{g}(\mathbf{x}^*) \leq \mathbf{0} &\Rightarrow x_1^* + x_2^* + x_3^* \leq -3 \\ \mu_j^* g_j(\mathbf{x}^*) = 0 &\Rightarrow \mu^*(x_1^* + x_2^* + x_3^* + 3) = 0 \end{aligned}$$

The solution is $\mathbf{x}^* = (-1, -1, -1)$ and $\mu^* = 1$ which satisfies $\mu \geq 0$. Furthermore, all points are regular, so this is the global minimum.

- **Theorem (KKT Sufficient Conditions):** Assume that f , \mathbf{h} and \mathbf{g} are twice continuously differentiable, and that $\mathbf{x}^* \in \mathbb{R}^n, \boldsymbol{\lambda}^* \in \mathbb{R}^m, \boldsymbol{\mu}^* \in \mathbb{R}^r$ satisfy

$$\begin{aligned} \nabla_{\mathbf{x}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) = \mathbf{0} \quad \mathbf{h}(\mathbf{x}^*) = \mathbf{0} \quad \mathbf{g}(\mathbf{x}^*) \leq \mathbf{0} \\ \boldsymbol{\mu}^* \geq \mathbf{0} \quad \mu_j^* = 0 \quad \forall j \notin \mathcal{A}(\mathbf{x}^*) \\ \mathbf{d}^\top \nabla_{xx}^2 \mathcal{L}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*) \mathbf{d} > 0 \quad \forall \mathbf{d} \in \mathcal{V}^{\text{EQ}}(\mathbf{x}^*) \setminus \{\mathbf{0}\} \end{aligned}$$

Assume also that

$$\mu_j^* > 0 \quad \forall j \in \mathcal{A}(\mathbf{x}^*)$$

Then \mathbf{x}^* is a strict local minimum.

Proof: We follow the equality case. Suppose \mathbf{x}^* is not a strict local minimum. Then there exists $\{\mathbf{x}_k\} \subset \mathbb{R}^n, \mathbf{h}(\mathbf{x}_k) = \mathbf{0}, \mathbf{g}(\mathbf{x}_k) \leq \mathbf{0}, \mathbf{x}_k \neq \mathbf{x}^*, \mathbf{x}_k \rightarrow \mathbf{x}^*$ with $f(\mathbf{x}_k) \leq f(\mathbf{x}^*)$. We define

$$\mathbf{d}_k = \frac{\mathbf{x}_k - \mathbf{x}^*}{\|\mathbf{x}_k - \mathbf{x}^*\|} \quad \delta_k = \|\mathbf{x}_k - \mathbf{x}^*\|$$

Without loss of generality, assume $\mathbf{d}_k \rightarrow \mathbf{d}$. Using the same mean-value-theorem argument as in the sufficient conditions proof for Lagrange multipliers, we find that $\nabla \mathbf{h}(\mathbf{x}^*)^\top \mathbf{d} = \mathbf{0}$.

Now, consider the \mathbf{g} . If $j \in \mathcal{A}(\mathbf{x}^*)$, then $g_j(\mathbf{x}^*) = 0$, and so since \mathbf{x}_k is feasible, $g_j(\mathbf{x}_k) - g_j(\mathbf{x}^*) \leq 0$. Furthermore, by the mean-value-theorem,

$$g_j(\mathbf{x}_k) - g_j(\mathbf{x}^*) = \nabla g_j(\tilde{\mathbf{x}}_k)^\top (\mathbf{x}_k - \mathbf{x}^*) \rightarrow \nabla g_j(\mathbf{x}^*)^\top \mathbf{d}$$

for some $\tilde{\mathbf{x}}_k \in [\mathbf{x}_k, \mathbf{x}^*]$. Together, these imply that

$$\nabla g_j(\mathbf{x}^*)^\top \mathbf{d} \leq 0$$

If the inequality is tight for all $j \in \mathcal{A}(\mathbf{x}^*)$, then $\mathbf{d} \in \mathcal{V}^{\text{EQ}}(\mathbf{x}^*)$, and we proceed as in the equality case.

If, on the other hand, $\nabla g_j(\mathbf{x}^*)^\top \mathbf{d} < 0$ for some $j \in \mathcal{A}(\mathbf{x}^*)$, then

$$\mathbf{d}^\top \nabla f(\mathbf{x}^*) = -\mathbf{d}^\top [\nabla \mathbf{h}(\mathbf{x}^*) \cdot \boldsymbol{\lambda}^* + \nabla \mathbf{g}(\mathbf{x}^*) \cdot \boldsymbol{\mu}^*] > 0$$

[The inequality holds because $\nabla g_j(\mathbf{x}^*)^\top \mathbf{d} < 0$ whereas \mathbf{d} dotted with every other gradient of \mathbf{h} and \mathbf{g} is 0, as we found above]. However, by the definition of our sequence, $f(\mathbf{x}_k) - f(\mathbf{x}^*) \leq 0$, and by the mean value theorem,

$$f(\mathbf{x}_k) - f(\mathbf{x}^*) = \nabla f(\tilde{\mathbf{x}}_k)^\top (\mathbf{x}_k - \mathbf{x}^*) \rightarrow \nabla f(\mathbf{x}^*)^\top \mathbf{d}$$

For some $\tilde{\mathbf{x}}_k \in [\mathbf{x}_k, \mathbf{x}^*]$. This implies that $\nabla f(\mathbf{x}^*)^\top \mathbf{d} \leq 0$. This, together with the statement $\mathbf{d}^\top \nabla f(\mathbf{x}^*) > 0$, is a contradiction.

- Practical point of note: the problem above was a **min** problem with \leq inequalities, and it resulted in a *positive* Lagrange multiplier. If either of these facts changed (eg: max and/or \geq), the sign of the Lagrange multiplier required would be flipped. For consistency, it is then customary to change the way we define the Lagrangian to ensure the Lagrange multiplier remains positive; specifically, the term $+\mu g_i(\mathbf{x})$ changes to $-\mu g_i(\mathbf{x})$.

- **Sensitivity analysis**

- Consider the optimization problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \text{ s.t. } \mathbf{h}(\mathbf{x}) = \mathbf{u}, \mathbf{g}(\mathbf{x}) \leq \mathbf{v}$$

- **Theorem:** Suppose there exists a triple $(\mathbf{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ satisfying the second order sufficient conditions when $(\mathbf{u}, \mathbf{v}) = (\mathbf{0}, \mathbf{0})$, with \mathbf{x}^* regular. Then there exists a neighborhood N of $(\mathbf{u}, \mathbf{v}) = (\mathbf{0}, \mathbf{0})$ and a function $\mathbf{x}^*(\cdot, \cdot)$ defined on N such that

- $\mathbf{x}^*(\mathbf{0}, \mathbf{0}) = \mathbf{x}^*$, and for each $(\mathbf{u}, \mathbf{v}) \in N$, $\mathbf{x}^*(\mathbf{u}, \mathbf{v})$ is a strict local minimum.

- $\mathbf{x}^*(\cdot, \cdot)$ is continuously differentiable.
- If $p(\mathbf{u}, \mathbf{v}) = f[\mathbf{x}^*(\mathbf{u}, \mathbf{v})]$

$$\nabla_{\mathbf{u}} p(\mathbf{0}, \mathbf{0}) = -\boldsymbol{\lambda}^* \quad \nabla_{\mathbf{v}} p(\mathbf{0}, \mathbf{0}) = -\boldsymbol{\mu}^*$$

• **Existence of Lagrange Multipliers**

- **Theorem (Farkas' Lemma):** Consider a matrix $A \in \mathbb{R}^{m \times n}$. Then a vector $\mathbf{z} \in \mathbb{R}^m$ satisfies

$$\mathbf{z} \cdot \mathbf{y} \leq 0 \text{ for all } \mathbf{y} \in \mathbb{R}^m \text{ such that } A^\top \mathbf{y} = \mathbf{0}$$

if and only if

$$\mathbf{z} = A\mathbf{x} \text{ for some } \mathbf{x} \in \mathbb{R}^n \text{ with } \mathbf{x} \geq \mathbf{0}$$

- **Theorem:** Consider the standard inequality optimization program above. Let \mathbf{x}^* be a local minimum. Then there exists Lagrange multipliers $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ satisfying the FOCs

$$\begin{aligned} \nabla f(\mathbf{x}^*) + \nabla \mathbf{h}(\mathbf{x}^*)\boldsymbol{\lambda}^* + \nabla \mathbf{g}(\mathbf{x}^*)\boldsymbol{\mu}^* &= \mathbf{0} \\ \boldsymbol{\mu}^* \geq \mathbf{0} \quad \mu_j^* g_j(\mathbf{x}^*) &= 0 \quad \forall 1 \leq j \leq r \end{aligned}$$

If and only if

$$\mathcal{D}(\mathbf{x}^*) \cap \mathcal{V}^{\text{EQ}}(\mathbf{x}^*) = \emptyset$$

Proof: Suppose there are no equality constraints. Then $\mathcal{D}(\mathbf{x}^*) \cap \mathcal{V}^{\text{EQ}}(\mathbf{x}^*) = \emptyset$ is equivalent to

$$\mathbf{d} \notin \mathcal{D}(\mathbf{x}^*) \text{ for all } \mathbf{d} \text{ such that } \mathbf{d} \in \mathcal{V}^{\text{EQ}}(\mathbf{x}^*)$$

{ \mathbf{d} must be “uphill” for every \mathbf{d} {that keeps us in the feasible set}}

$$\nabla f(\mathbf{x}^*) \cdot \mathbf{d} \geq 0 \text{ for all } \mathbf{d} \text{ such that } \nabla g_j(\mathbf{x}^*) \cdot \mathbf{d} \leq 0 \quad \forall j \in \mathcal{A}(\mathbf{x}^*)$$

By Farkas' Lemma, this is equivalent to

$$\nabla f(\mathbf{x}^*) + \nabla \mathbf{g}(\mathbf{x}^*)\boldsymbol{\mu} = \mathbf{0}$$

For some $\boldsymbol{\mu}$ with $\boldsymbol{\mu} \geq \mathbf{0}$ and $\mu_j = 0$ if $j \notin \mathcal{A}(\mathbf{x}^*)$.

If there are equality constraints $\mathbf{h}(\mathbf{x}) = \mathbf{0}$, we can replace them with inequality constraints $\mathbf{h}(\mathbf{x}) \leq \mathbf{0}$ and $-\mathbf{h}(\mathbf{x}) \leq \mathbf{0}$.

- **Constraint qualification:** In summary, what we have proved so far is

$$\begin{aligned}
 & \mathbf{x}^* \text{ local minimum} \\
 & \Rightarrow \mathcal{D}(\mathbf{x}^*) \cap \mathcal{T}(\mathbf{x}^*) = \emptyset \\
 & \Rightarrow \mathcal{D}(\mathbf{x}^*) \cap \mathcal{V}(\mathbf{x}^*) = \emptyset \\
 & \Rightarrow \text{Lagrange multipliers exist}
 \end{aligned}$$

In proving this sequence of statements, we made use of the fact that \mathbf{x}^* was a regular point. However, it is possible to prove the existence of Lagrange multipliers under weaker assumptions called *constraint qualifications*. If the constraints are linear, for example, Lagrange multipliers are guaranteed to exist. The weakest form of constraint qualification is *quasiregularity*, which requires that $\mathcal{V}(\mathbf{x}^*) = \mathcal{T}(\mathbf{x}^*)$.

- **Theorem (Linear Constraint Qualification):** Consider the optimization program

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \text{ s.t. } A\mathbf{x} \leq \mathbf{b}$$

Suppose that \mathbf{x}^* is a local minimum. Then \mathbf{x}^* is quasi-regular and Lagrange multipliers exist. This trivially also applies to linear equality constraints, and can be extended to linear equality constraints and concave inequality constraints.

- **Theorem (General Sufficiency Condition):** Consider the optimization program

$$\min_{\mathbf{x} \in \Omega} f(\mathbf{x}) \text{ s.t. } \mathbf{g}(\mathbf{x}) \leq \mathbf{0}$$

Let $\mathbf{x}^* \in \mathbb{R}^n$ be a feasible point, and $\boldsymbol{\mu}^* \in \mathbb{R}^r$ be a vector such that

$$\begin{aligned}
 & \boldsymbol{\mu}^* \geq \mathbf{0} \\
 & \mu_j^* = 0 \quad \forall j \notin \mathcal{A}(\mathbf{x}^*) \\
 & \mathbf{x}^* \in \operatorname{argmin}_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu}^*)
 \end{aligned}$$

Then \mathbf{x}^* is a global minimum. Note that no differentiability or continuity assumptions were made.

Proof:


$$\begin{aligned}
 f(\mathbf{x}^*) &= f(\mathbf{x}^*) + \boldsymbol{\mu}^* \cdot \mathbf{g}(\mathbf{x}^*) \\
 &= \min_{\mathbf{x} \in \Omega} \{f(\mathbf{x}) + \boldsymbol{\mu}^* \mathbf{g}(\mathbf{x})\} \\
 &\leq \min_{\mathbf{x} \in \Omega, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}} \{f(\mathbf{x}) + \boldsymbol{\mu}^* \mathbf{g}(\mathbf{x})\} \\
 &\leq \min_{\mathbf{x} \in \Omega, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}} f(\mathbf{x})
 \end{aligned}$$

Optimization over Convex Sets

- Consider the problem $\min f(\mathbf{x})$ subject to $\mathbf{x} \in \mathcal{C} \subset \mathbb{R}^n$ where \mathcal{C} is a convex set.
- **Theorem (necessary conditions):** If \mathbf{x}^* is a local minimum and f is continuously differentiable in a neighborhood of \mathbf{x}^* , then

$$\nabla f(\mathbf{x}^*) \cdot (\mathbf{x} - \mathbf{x}^*) \geq 0 \quad \forall \mathbf{x} \in \mathcal{C}$$

Geometrically, at a local minimum, the gradient of the objective function must make an acute angle with any improving direction (or else, we could improve along that direction), and since the set is convex, $\mathbf{x} - \mathbf{x}^*$ must be an improving direction for any $\mathbf{x} \in \mathcal{C}$, since the line between \mathbf{x} and \mathbf{x}^* is in \mathcal{C} :



Proof: Suppose that there exists an $\mathbf{x} \in \mathcal{C}$ with $\nabla f(\mathbf{x}^*) \cdot (\mathbf{x} - \mathbf{x}^*) < 0$. Consider applying the mean value theorem to the function $g(\varepsilon) = f[\mathbf{x}^* + \varepsilon(\mathbf{x} - \mathbf{x}^*)]$:

$$g(\varepsilon) - g(0) = g'(c)[\varepsilon - 0]$$

$$f[\mathbf{x}^* + \varepsilon(\mathbf{x} - \mathbf{x}^*)] = f(\mathbf{x}^*) + \varepsilon \nabla f(\mathbf{x}^* + \varepsilon s(\mathbf{x} - \mathbf{x}^*)) \cdot (\mathbf{x} - \mathbf{x}^*)$$

Where $s \in (0,1)$. Since ∇f is continuous, and $\nabla f(\mathbf{x}^*) \cdot (\mathbf{x} - \mathbf{x}^*) < 0$, then fall all sufficiently small ε , this will also be true of $\nabla f(\mathbf{x}^* + \varepsilon s(\mathbf{x} - \mathbf{x}^*)) \cdot (\mathbf{x} - \mathbf{x}^*) < 0$. This, however, implies that $f(\mathbf{x}^*) > f[\mathbf{x}^* + \varepsilon(\mathbf{x} - \mathbf{x}^*)]$. Since \mathcal{C} is convex, this point is in \mathcal{C} , and so this contradicts the optimality of \mathbf{x}^* .

- **Theorem (optimality for convex optimization):** Suppose $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex over \mathcal{C} . Then
 1. Any local minimum of f is also a global minimum.
 2. If f is strictly convex, then there exists at most one global minimum.

Proof: Consider:

1. Suppose \mathbf{x}^* is a local minimum, and there exists some $\mathbf{x} \neq \mathbf{x}^*$ with $f(\mathbf{x}) < f(\mathbf{x}^*)$. Then for any $\lambda \in [0,1)$, then

$$f[\lambda \mathbf{x}^* + (1 - \lambda)\mathbf{x}] \leq \lambda f(\mathbf{x}^*) + (1 - \lambda)f(\mathbf{x}) < f(\mathbf{x}^*)$$

Now:

- Since \mathcal{C} is convex, $\lambda \mathbf{x}^* + (1 - \lambda)\mathbf{x} \in \mathcal{C}$
- For every small r , there will be a λ such that $\lambda \mathbf{x}^* + (1 - \lambda)\mathbf{x} \in B_r(\mathbf{x}^*)$. The fact the value of f at that point is more than at \mathbf{x}^* contradicts our local optimality assumption.

2. Suppose $\mathbf{x}_0 \neq \mathbf{x}_1$ are two global minima. Then

$$f\left(\frac{1}{2}(\mathbf{x}_0 + \mathbf{x}_1)\right) < \frac{1}{2}(f(\mathbf{x}_0) + f(\mathbf{x}_1)) = f(\mathbf{x}_0) = f(\mathbf{x}_1)$$

This means there is a point at which f is strictly less than at \mathbf{x}_0 and \mathbf{x}_1 , which contradicts our local optimality assumption.

- **Theorem (necessary & sufficient conditions):** If $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is convex over \mathcal{C} and differentiable and $\mathbf{x}^* \in \mathcal{C}$ is a feasible point, then \mathbf{x}^* is globally optimal if and only if

$$\nabla f(\mathbf{x}^*) \cdot (\mathbf{x} - \mathbf{x}^*) \geq 0 \quad \forall \mathbf{x} \in \mathcal{C}$$

Proof: Necessity follows from the previous theorem. For sufficiency, note that

$$f(\mathbf{x}) \geq f(\mathbf{x}^*) + \nabla f(\mathbf{x}^*) \cdot (\mathbf{x} - \mathbf{x}^*) \geq f(\mathbf{x}^*) \quad \forall \mathbf{x} \in \mathcal{C}$$

Duality

- **Supporting & Separating hyperplanes**


- **Definition:** The hyperplane $\{\mathbf{x} \in \mathbb{R}^n : \boldsymbol{\mu} \cdot \mathbf{x} = b\}$ with $\boldsymbol{\mu} \in \mathbb{R}^n \setminus \{0\}$ supports the convex set $\mathcal{C} \subset \mathbb{R}^n$ at the point $\bar{\mathbf{x}}$ if

$$\boldsymbol{\mu} \cdot \mathbf{x} \geq \boldsymbol{\mu} \cdot \bar{\mathbf{x}} = b \quad \forall \mathbf{x} \in \mathcal{C}$$

or equivalently

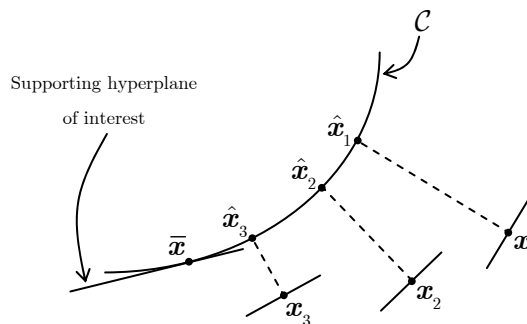
$$\inf_{\mathbf{x} \in \mathcal{C}} \boldsymbol{\mu} \cdot \mathbf{x} \geq \boldsymbol{\mu} \cdot \bar{\mathbf{x}} = b$$

This statement requires that (a) $\bar{\mathbf{x}}$ lies on the hyperplane, and (b) \mathcal{C} lies entirely on one side of the hyperplane.



- **Theorem (supporting hyperplane):** Let $C \subset \mathbb{R}^n$ be a convex set and $\bar{x} \in \mathbb{R}^n$ be a point that is not in the interior of C . Then there exists a supporting hyperplane at \bar{x} .

Proof: The intuition behind this proof will be to consider a sequence $\{\mathbf{x}_k\}$ of points *outside* the closure of C converging to \bar{x} . This is possible, since $\bar{x} \notin \text{int} C$. We consider the projection of each of these \mathbf{x}_k onto \bar{C} (denoted $\hat{\mathbf{x}}_k$), and we consider the hyperplane *perpendicular* to the line from \mathbf{x}_k to $\hat{\mathbf{x}}_k$:



We then show that

- The set C lies on one side of each of these hyperplanes.
- These hyperplanes tend to the supporting hyperplane of interest. In other words, \bar{x} lies on the limiting supporting hyperplane.

Now, to the formal stuff. Define $\bar{C} = \text{cl } C$, and note that \bar{C} is convex. Let $\{\mathbf{x}_k\}$ be a sequence of points such that $\mathbf{x}_k \notin \bar{C}, \mathbf{x}_k \rightarrow \bar{x}$, and for each \mathbf{x}_k , let $\hat{\mathbf{x}}_k$ be its projection onto \bar{C} . Then, by the projection theorem,

$$(\hat{\mathbf{x}}_k - \mathbf{x}_k) \cdot (\mathbf{x} - \hat{\mathbf{x}}_k) \geq 0 \quad \forall \mathbf{x} \in \bar{C}$$

And this means that for all k and $\mathbf{x} \in \bar{C}$,

$$\begin{aligned} (\hat{\mathbf{x}}_k - \mathbf{x}_k) \cdot \mathbf{x} &\geq (\hat{\mathbf{x}}_k - \mathbf{x}_k) \cdot \hat{\mathbf{x}}_k \\ &= (\hat{\mathbf{x}}_k - \mathbf{x}_k) \cdot (\hat{\mathbf{x}}_k - \mathbf{x}_k) + (\hat{\mathbf{x}}_k - \mathbf{x}_k) \cdot \mathbf{x}_k \\ &\geq (\hat{\mathbf{x}}_k - \mathbf{x}_k) \cdot \mathbf{x}_k \end{aligned}$$

More succinctly

$$(\hat{\mathbf{x}}_k - \mathbf{x}_k) \cdot \mathbf{x} \geq (\hat{\mathbf{x}}_k - \mathbf{x}_k) \cdot \mathbf{x}_k = \text{constant} \quad \forall \mathbf{x} \in \bar{\mathcal{C}}$$

In other words, $\bar{\mathcal{C}}$ lies on one side of each of those hyperplanes.

But now, set

$$\boldsymbol{\mu}_k = \frac{\hat{\mathbf{x}}_k - \mathbf{x}_k}{\|\hat{\mathbf{x}}_k - \mathbf{x}_k\|}$$

Then the equation above can be written as

$$\boldsymbol{\mu}_k \cdot \mathbf{x} \geq \boldsymbol{\mu}_k \cdot \mathbf{x}_k \quad \forall k, \mathbf{x} \in \mathcal{C}$$

Since $\|\boldsymbol{\mu}_k\| = 1$, the sequence $\{\boldsymbol{\mu}_k\}$ is bounded has a non-zero subsequential limit


$\boldsymbol{\mu}$. Letting $k \rightarrow \infty$, we get $\boldsymbol{\mu}_k \rightarrow \boldsymbol{\mu}$ and $\mathbf{x}_k \rightarrow \bar{\mathbf{x}}$, so

$$\boldsymbol{\mu} \cdot \mathbf{x} \geq \boldsymbol{\mu} \cdot \bar{\mathbf{x}} \quad \forall k, \mathbf{x} \in \bar{\mathcal{C}}$$

As required.

- **Theorem (separating hyperplane):** Let $\mathcal{C}_1, \mathcal{C}_2 \in \mathbb{R}^n$ be two disjoint non-empty convex sets. Then there exists a hyperplane that separates them; ie: a vector $\boldsymbol{\mu} \in \mathbb{R}^n, \boldsymbol{\mu} \neq 0$ and a scalar $b \in \mathbb{R}$ with

$$\boldsymbol{\mu} \cdot \mathbf{x}_1 \leq b \leq \boldsymbol{\mu} \cdot \mathbf{x}_2 \quad \forall \mathbf{x}_1 \in \mathcal{C}_1, \mathbf{x}_2 \in \mathcal{C}_2$$



Proof: Consider the convex set

$$\mathcal{D} = \mathcal{C}_1 - \mathcal{C}_2 = \{x_1 - x_2 : x_1 \in \mathcal{C}_1, x_2 \in \mathcal{C}_2\}$$


Since the two sets are disjoint, $\mathbf{0} \notin \mathcal{D}$. Thus, by the supporting hyperplane theorem, there exists a vector $\boldsymbol{\mu} \neq \mathbf{0}$ with

$$\mathbf{0} \leq \boldsymbol{\mu}^\top (\mathbf{x}_1 - \mathbf{x}_2) \quad \forall \mathbf{x}_1 \in \mathcal{C}_1, \mathbf{x}_2 \in \mathcal{C}_2$$


Setting $b = \sup_{\mathbf{x}_2 \in \mathcal{C}_2} \boldsymbol{\mu}^\top \mathbf{x}_2$, we obtain the desired result.

- **Theorem (strictly separating hyperplane):** Let $\mathcal{C} \subset \mathbb{R}^n$ be a closed convex set and $\bar{\mathbf{x}} \notin \mathcal{C}$ a point. Then there exists a hyperplane that *strictly* separates the point and the set. In other words, $\exists \boldsymbol{\mu} \in \mathbb{R}^n \setminus \{0\}, b \in \mathbb{R}$ such that

$$\boldsymbol{\mu} \cdot \bar{\mathbf{x}} < b < \inf_{\mathbf{x} \in \mathcal{C}} \boldsymbol{\mu} \cdot \mathbf{x}$$



Proof: Define $r = \min_{\mathbf{x} \in \mathcal{C}} \|\mathbf{x} - \bar{\mathbf{x}}\|$. This will be > 0 because \mathcal{C} is closed, and $\bar{\mathbf{x}} \notin \mathcal{C}$. Now, let $\bar{\mathcal{C}} = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x} - \bar{\mathbf{x}}\| \leq r/2\}$. Clearly, \mathcal{C} and $\bar{\mathcal{C}}$ are disjoint, so we can apply the separating hyperplane theorem. Diagrammatically:



- **Corollary:** If $\mathcal{C} \subsetneq \mathbb{R}^n$ is a closed convex set, then it is the intersection of all closed halfspaces that contain it.

Proof: Let \mathcal{H} be the collection of all closed halfspaces containing \mathcal{C} , and let $\bar{H} = \bigcap_{H \in \mathcal{H}} H$.

Since $\mathcal{C} \neq \mathbb{R}^n$, the strictly separating hyperplane theorem implies that \mathcal{H} is non-empty (since there is a point $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{x} \notin \mathcal{C}$) and clearly, $\mathcal{C} \subset \bar{H}$.

Now, suppose there exists an $\mathbf{x} \in \bar{H}$ with $\mathbf{x} \notin \mathcal{C}$, then

- By the strictly separating hyperplane theorem, there exists a halfplane that strictly separates \mathbf{x} and \mathcal{C} , with \mathbf{x} on the opposite side of the hyperplane as \mathcal{C} .
- Thus, $\mathbf{x} \notin \bar{H}$. This is a contradiction.


• **Farkas' Lemma**

○ **Lemma (Farkas):** Consider a matrix $A \in \mathbb{R}^{m \times n}$. Given a vector $\mathbf{z} \in \mathbb{R}^m$, the following statements are equivalent


1. $\mathbf{z} \cdot \mathbf{y} \leq 0$ for all $\mathbf{y} \in \mathbb{R}^m$ with $A^\top \mathbf{y} \leq \mathbf{0}$.
2. $\mathbf{z} = A\mathbf{x}$ for some $\mathbf{x} \in \mathbb{R}^n$ with $\mathbf{x} \geq \mathbf{0}$.

Geometrically, the two parts of this theorem are as follows:

1. \mathbf{z} makes an obtuse angle with all vectors \mathbf{y} that make an obtuse angle with every column of A .



2. \mathbf{z} lies in the cone formed by the columns of A .



Proof. Consider:

- 2 \rightarrow 1 Dot both sides of the equation with \mathbf{y} to get


$$\mathbf{z} \cdot \mathbf{y} = (A\mathbf{x}) \cdot \mathbf{y} = \mathbf{x}A^\top \mathbf{y} \leq 0$$

Where the last equality follows because \mathbf{x} is positive and $A^\top \mathbf{y} \leq \mathbf{0}$, by assumption.

- 1 \rightarrow 2 Suppose \mathbf{z} satisfies (1) [it makes an obtuse angle to \mathbf{y} , which makes an obtuse angle to all the columns of A], but that there is no $\mathbf{x} \geq \mathbf{0}$ with $\mathbf{z} = A\mathbf{x}$ [it lies outside the cone defined by the columns of A].

Then define $\mathcal{C} = \{Ax : x \geq 0\}$ to be the closed, convex cone formed by the columns of A . Clearly, $z \notin \mathcal{C}$.

By the strictly separating hyperplane theorem, there exists $y \in \mathbb{R}^m \setminus \{0\}$ such that $y \cdot z > y \cdot c \forall c \in \mathcal{C}$:



We will show that this y forms an obtuse angle with all the columns of A , but an acute angle with z (thus contradicting [1]).

- **Obtuse angle with the columns of A :** Note that $\lambda A^{\text{col } i} \in \mathcal{C}$, and so $y \cdot z > \lambda y \cdot A^{\text{col } i}$. This means that regardless how far we go along any of the edges of the cone, the point we reach must still be on a different side of the hyperplane to z . This implies that the hyperplane must be “titled away” from every edge of the cone (the columns of A), and that the normal *must* make an obtuse angle to the said column. More formally,

$$y \cdot A^{\text{col } i} < \frac{y \cdot z}{\lambda} \xrightarrow{\lambda \rightarrow \infty} 0$$

- **Acute angle with z :** Note that since $0 \in \mathcal{C}$, $y \cdot z > 0$ and so y must form an acute angle with z .

We have therefore reached a contradiction.

- We quickly digress to cover an interesting application of Farkas’ Lemma: arbitrage-free pricing.
 - Consider a market with n assets. A portfolio is a vector $x \in \mathbb{R}^n$ in which x_i describes the amount of asset i purchased.

- The current price of asset i is v_i , and so the current price of portfolio \mathbf{x} is $\mathbf{v} \cdot \mathbf{x}$.
- The future is modeled by m different scenarios, described by a matrix $R \in \mathbb{R}^{m \times n}$. In scenario m , asset n yields R_{mn} . The future payoffs are therefore $R\mathbf{x}$.
- **Definition:** An *arbitrage opportunity* is a portfolio \mathbf{x} such that

$$\mathbf{v} \cdot \mathbf{x} < 0 \quad R\mathbf{x} \geq \mathbf{0}$$

In other words, the portfolio costs a negative amount to buy right now and will always yield nonnegative returns. We say a market is *consistent* if and only if it has no arbitrage opportunity.

- By Farkas' Lemma, a market is consistent if and only if there exists a vector $\mathbf{q} \geq \mathbf{0}$ with $\mathbf{v} = R^\top \mathbf{q}$. Suppose there exists such a vector, and define

$$r = \frac{1}{\mathbf{1} \cdot \mathbf{q}} - 1 \quad \boldsymbol{\pi} = (1 + r)\mathbf{q}$$

$\boldsymbol{\pi}$ is then a probability distribution, because $\boldsymbol{\pi} \geq \mathbf{0}$ and $\mathbf{1} \cdot \boldsymbol{\pi} = 1$. We also have that

$$v_i = \frac{1}{1 + r} (\boldsymbol{\pi} \cdot \mathbf{R}^{\text{col } i})$$

In other words, the prices today are the discounted value of the future prices under the distribution $\boldsymbol{\pi}$. We call $\boldsymbol{\pi}$ the *risk-neutral distribution*.

• *The primal and the dual*

- **Definition:** Consider the *primal problem* $\min f(\mathbf{x})$ s.t. $\mathbf{g}(\mathbf{x}) \leq \mathbf{0}, \mathbf{x} \in \Omega \subset \mathbb{R}^n$, and define

$$f^* = \inf_{\mathbf{x} \in \Omega, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}} f(\mathbf{x})$$

Assumption: The feasible set is non-empty and the optimal cost is bounded below. In other words, $-\infty < f^* < \infty$.

Definition (Lagrangian function): For $\mathbf{x} \in \Omega, \boldsymbol{\mu} \in \mathbb{R}^r$

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\mu}) = f(\mathbf{x}) + \boldsymbol{\mu} \cdot \mathbf{g}(\mathbf{x})$$

We say $\boldsymbol{\mu}^*$ is a *geometric multiplier*

- $\boldsymbol{\mu}^* \geq \mathbf{0}$
- $f^* = \inf_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu}^*)$
- **Definition:** The set $\mathcal{S} \subset \mathbb{R}^{r+1}$ of *constraint-cost pairs* is defined by


$$\mathcal{S} = \left\{ (\mathbf{g}(\mathbf{x}), f(\mathbf{x})) : \mathbf{x} \in \Omega \right\}$$

Given $(\boldsymbol{\mu}, \mu_0) \in \mathbb{R}^{r+1} \setminus \{0\}$, the *hyperplane passing through* $(\bar{\mathbf{z}}, \bar{w})$ is

$$\left\{ (\mathbf{z}, w) : \boldsymbol{\mu} \cdot \mathbf{z} + \mu_0 w = \boldsymbol{\mu} \cdot \bar{\mathbf{z}} + \mu_0 \bar{w} \right\}$$


The positive/negative halfspace has the $=$ replaced by \geq/\leq . The hyperplane is *non-vertical* if $\mu_0 \neq 0$. In such a case, it can be normalized such that $\mu_0 = 1$.

- **Lemma (visualization):**
 - The hyperplane with normal $(\boldsymbol{\mu}, 1)$ that passes through $(\mathbf{g}(\bar{\mathbf{x}}), f(\bar{\mathbf{x}}))$ intercepts the vertical axis $\left\{ (0, w) \in \mathbb{R}^{r+1} : w \in \mathbb{R} \right\}$ at the level $L(\mathbf{x}, \boldsymbol{\mu})$.




In other words, the simple interpretation of $L(\mathbf{x}, \boldsymbol{\mu})$ is the point at which a given hyperplane crosses the vertical axis.


- Amongst all hyperplanes with normal $(\boldsymbol{\mu}, 1)$ that contain \mathcal{S} in the positive halfspace, the highest interception of the vertical axis is attained at $\inf_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu})$.



- μ^* is a geometric multiplier if and only if $\mu^* \geq 0$ and, among all hyperplanes with normal $(\mu^*, 1)$ that contain S in the positive halfspace, the highest interception of the vertical axis is attained at f^* (in other words $\inf_{x \in \Omega} L(x, \mu^*) = f^*$).



Here is an example of a problem in which a geometric multiplier does *not* exist:



In this case, there is no hyperplane that contains the whole of S in one of its halfspaces that also passes through the point $(0, f^*)$. Thus, there is no μ for which $\inf_{x \in \Omega} L(x, \mu) = f^*$ because, as we saw in part (2),

$\inf_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu})$ considers the highest intercept for planes that contain the whole of \mathcal{S} in one of their halfspaces.

Proof:

- The hyperplane is the set (\mathbf{z}, w) satisfying

$$\boldsymbol{\mu} \cdot \mathbf{z} + w = \boldsymbol{\mu} \cdot \mathbf{g}(\mathbf{x}) + f(\mathbf{x})$$

Clearly, if $\mathbf{z} = \mathbf{0}$, we must have $w = L(\mathbf{x}, \boldsymbol{\mu})$.

- The hyperplane with normal $(\boldsymbol{\mu}, 1)$ that intercepts the axis at level c is the set (\mathbf{z}, w) with

$$\boldsymbol{\mu} \cdot \mathbf{z} + w = c$$

If \mathcal{S} lies in the positive halfspace, then

$$L(\mathbf{x}, \boldsymbol{\mu}) = \boldsymbol{\mu} \cdot \mathbf{g}(\mathbf{x}) + f(\mathbf{x}) \geq c \quad \forall \mathbf{x} \in \Omega$$


Thus, the maximum intercept is $\inf_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu})$.

- We need to show that $f^* = \inf_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu}^*)$. It is obvious from part (2) that this is true if and only if among all hyperplanes with normal $(\boldsymbol{\mu}^*, 1)$, the highest interception with the vertical axis is at f^* .

- Theorem:** Let $\boldsymbol{\mu}^*$ be a geometric multiplier. Then \mathbf{x}^* is a global minimum if and only if \mathbf{x}^* is feasible and

$$\mathbf{x}^* \in \operatorname{argmin}_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu}^*) \quad \boldsymbol{\mu}^* \cdot \mathbf{g}_j(\mathbf{x}^*) = 0 \quad \forall 1 \leq j \leq r$$

Geometrically, the first statement is that \mathbf{x}^* is, indeed, the value of \mathbf{x} that minimizes L at that value of $\boldsymbol{\mu}$ (and recall that $f^* = \inf_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu}^*)$), and the second is that either \mathbf{x}^* is on the boundary of the feasible set (in which case we can “improve no further”) or that the geometric multiplier is horizontal (in which case the minimum is attained on the interior of the feasible set).



Proof: Assume \mathbf{x}^* is a global minimum. Then

$$f^* = f(\mathbf{x}^*) \geq f(\mathbf{x}^*) + \boldsymbol{\mu}^* \cdot \mathbf{g}(\mathbf{x}^*) = L(\mathbf{x}^*, \boldsymbol{\mu}^*) \geq \inf_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu}^*) = f^*$$

Since the two ends are equal, we have equality the whole way through, and

$$\mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}^*) = \inf_{\mathbf{x} \in \Omega} \mathcal{L}(\mathbf{x}, \boldsymbol{\mu}^*)$$

Similarly, we have $f(\mathbf{x}^*) \geq f(\mathbf{x}^*) + \boldsymbol{\mu}^* \cdot \mathbf{g}(\mathbf{x}^*)$. Thus, since $\boldsymbol{\mu}^* \geq \mathbf{0}$, we must have


$$\boldsymbol{\mu}^* \cdot \mathbf{g}(\mathbf{x}^*) = 0$$

Conversely, if $\mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}^*) = \inf_{\mathbf{x} \in \Omega} \mathcal{L}(\mathbf{x}, \boldsymbol{\mu}^*)$ and $\boldsymbol{\mu}^* \cdot \mathbf{g}(\mathbf{x}^*) = 0$:

$$f(\mathbf{x}^*) = f(\mathbf{x}^*) + \boldsymbol{\mu}^* \cdot \mathbf{g}(\mathbf{x}^*) = L(\mathbf{x}^*, \boldsymbol{\mu}^*) = \min_{\mathbf{x} \in \Omega} L(\mathbf{x}, \boldsymbol{\mu}^*) = f^*$$

- o We define the *dual function* $q(\boldsymbol{\mu}) = \inf_{\mathbf{x} \in \Omega} \mathcal{L}(\mathbf{x}, \boldsymbol{\mu})$; it is concave over its domain, which is convex [because q is a point-wise minimum of concave – actually linear – functions]. Geometrically, q simply does the following:
 - Take a hyperplane with normal $(\boldsymbol{\mu}, 1)$
 - Push it up as far as possible until it hits the set \mathcal{S} .
 - Find the value at which it intercepts the $f(\mathbf{x})$ axis. By the second part of the visualization lemma, this is $q(\boldsymbol{\mu})$.

Diagrammatically:




- o We then define the *dual problem* as

$$\max q(\boldsymbol{\mu}) \text{ s.t. } \boldsymbol{\mu} \geq \mathbf{0}$$

And the dual optimal value is given by $q^* = \max_{\boldsymbol{\mu} \geq \mathbf{0}} q(\boldsymbol{\mu})$. [It is possible, however, that $q(\boldsymbol{\mu}) = -\infty$ for all $\boldsymbol{\mu} \geq \mathbf{0}$, in which case the dual problem is infeasible and $q^* = -\infty$].

Geometrically, the dual problem takes all of the snugly fitting hyperplanes defined by different values of μ and finds the one with the highest intercept:



[Parenthetically, recall that we proved above that a convex set can be described as the intersection of half-spaces containing it. If this is true, then the dual problem can – in some sense – be seen as “exploring” this set via its halfspaces rather than directly].

- We now consider a number of examples to illustrate the concepts above
 - Consider the program

$$\min_{\mathbf{x} \in \mathbb{R}_+^2} f(\mathbf{x}) = x_1 - x_2 \text{ s.t. } g(\mathbf{x}) = x_1 + x_2 - 1 \leq 0$$

It turns out the solution to this problem is $f^* = -1, \mathbf{x}^* = (0, 1)$.

Now, the Lagrangian is given by

$$\mathcal{L}(\mathbf{x}, \mu) = x_1 - x_2 + \mu(x_1 + x_2 - 1)$$

with $\mu \geq 0$. The dual function is

$$\begin{aligned} q(\mu) &= \min_{\mathbf{x} \in \mathbb{R}_+^2} \mathcal{L}(\mathbf{x}, \mu) \\ &= \min_{\mathbf{x} \in \mathbb{R}_+^2} (\mu + 1)x_1 + (\mu - 1)x_2 - \mu \end{aligned}$$

Clearly, this is only bounded if both the coefficients of x_1 and x_2 are positive. This means we need $\mu \geq 1$. Thus

$$q(\mu) = \begin{cases} -\infty & \mu < 1 \\ -\mu & \mu \geq 1 \end{cases}$$

Maximising this function, we clearly find $\mu^* = 1 \Rightarrow q^* = -1 = f^*$. Thus, there is no duality gap.

- Consider the program

$$\min_{\mathbf{x} \in \{(x_1, x_2) \in \mathbb{R}^2 : x_2 \geq 0\}} f(\mathbf{x}) = |x_1| + x_2 \text{ s.t. } g(\mathbf{x}) = x_1 \leq 0$$

It turns out the solution to this problem is $f^* = 0, \mathbf{x}^* = (0, 0)$.

Now, the Lagrangian is given by $\mathcal{L}(\mathbf{x}^*, \mu) = |x_1| + x_2 + \mu x_1$, with $\mu \geq 0$. If $|\mu| \leq 1$, the first term will dominate over the last term, and $x_1 = 0$ will be the bounded solution. If, on the other hand, μ falls outside this range, the last term dominates, and the Lagrangian can be shrunk ad infinitum, Thus

$$q(\mu) = \begin{cases} -\infty & |\mu| > 1 \\ 0 & |\mu| \leq 1 \end{cases}$$

The optimal solution is clearly $\mu \in [0, 1]$ (since $\mu \geq 0$), with $q^* = 0 = f^*$. So once again, no duality gap. Notice that primal degeneracy leads to dual non-uniqueness; see linear programming notes for more on this.

- Consider the program

$$\min_{x \in \mathbb{R}} f(x) = x \text{ s.t. } g(x) = x^2 \leq 0$$

It turns out the solution to this problem is $f^* = 0, x = 0$.

The Lagrangian is $\mathcal{L}(x, \mu) = x + \mu x^2$ with $\mu \geq 0$ and

$$q(\mu) = \begin{cases} -\frac{1}{4\mu} & \mu > 0 \\ -\infty & \mu \leq 0 \end{cases}$$

In this case, $q^* = 0$, so there is no duality gap, but there is also no optimal dual solution; it is attained as $\mu \rightarrow \infty$.

- Consider the program

$$\min_{x \in \{0, 1\}} f(x) = -x \text{ s.t. } g(x) = x - \frac{1}{2} \leq 0$$

By inspection, the solution is $f^* = 0, x^* = 0$.

The Lagrangian is $\mathcal{L}(x, \mu) = (\mu - 1)x - \frac{1}{2}\mu$ with $\mu \geq 0$ and


$$q(\mu) = \begin{cases} -\frac{1}{2}\mu & \mu \geq 1 \\ \frac{1}{2}\mu - 1 & \mu < 1 \end{cases} = \min \left\{ -\frac{\mu}{2}, \frac{\mu}{2} - 1 \right\}$$

And so the solution of the dual problem is $q^* = -\frac{1}{2}$. There is a duality gap.

- *Weak duality*

- **Theorem (Weak Duality):** $q^* \leq f^*$

Geometrically, this is an obvious statement. Whatever the shape of the set \mathcal{S} , the “snug hyperplane” with the highest cutting point will be lower than, or at the same place as f^* :



Proof: Consider feasible \mathbf{x} and $\boldsymbol{\mu}$; in other words, $\boldsymbol{\mu} \geq \mathbf{0}$, $\mathbf{x} \in \Omega$ and $\mathbf{g}(\mathbf{x}) \leq \mathbf{0}$.

$$q(\boldsymbol{\mu}) = \inf_{\mathbf{z} \in \Omega} L(\mathbf{z}, \boldsymbol{\mu}) \leq f(\mathbf{x}) + \boldsymbol{\mu}^\top \mathbf{g}(\mathbf{x}) \leq f(\mathbf{x})$$

Thus

$$q^* = \sup_{\boldsymbol{\mu} \geq \mathbf{0}} q(\boldsymbol{\mu}) \leq \inf_{\mathbf{x} \in \Omega, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}} f(\mathbf{x}) = f^*$$

- **Definition (Duality gap):** If $q^* = f^*$, we say there is no *duality gap*. If $q^* < f^*$ then there is a duality gap.
- **Theorem:** If there is no duality gap, the set of geometric multipliers is equal to the set of dual optimal solutions. If there is a duality gap, the set of geometric multipliers is empty.

Proof: By definition, $\boldsymbol{\mu}^* \geq \mathbf{0}$ is a geometric multiplier if and only if $f^* = \inf_{\mathbf{x} \in \Omega} \mathcal{L}(\mathbf{x}, \boldsymbol{\mu}^*) = q(\boldsymbol{\mu}^*) \leq q^*$. By weak duality, this holds if and only if there is no duality gap.

- Note also that if the primal problem is unbounded, $f^* = -\infty$. By weak duality, this means that $q(\boldsymbol{\mu}) = -\infty \forall \boldsymbol{\mu} \geq \mathbf{0}$, and so the dual problem is infeasible. Similarly, if the dual problem is feasible, the primal problem is bounded. However, if the primal is infeasible, we can say nothing about the dual.

- *Primal & Dual Optimality Conditions*

- **Theorem (optimality conditions):** $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ is an optimal solution/geometric multiplier pair to the problem $\min_{\mathbf{x} \in \Omega} f(\mathbf{x})$ s.t $\mathbf{g}(\mathbf{x}) \leq \mathbf{0}$ if and only if

1. PRIMAL FEASIBILITY: $\mathbf{x} \in \Omega$ and $\mathbf{g}(\mathbf{x}) \leq \mathbf{0}$
2. DUAL FEASIBILITY: $\boldsymbol{\mu}^* \geq \mathbf{0}$
3. LAGRANGIAN OPTIMALITY: $\mathbf{x}^* \in \operatorname{argmin}_{\mathbf{x} \in \Omega} \mathcal{L}(\mathbf{x}, \boldsymbol{\mu}^*)$
4. COMPLEMENTARY SLACKNESS: $\mu_j^* g_j^*(\mathbf{x}^*) = 0 \quad \forall 1 \leq j \leq r$

Note that this theorem is only useful if there is no duality gap.

Proof: Clearly, if $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ is an optimal solution/geometric multiplier pair, 1 and 2 must hold. 3 and 4 hold by earlier theorems.

Conversely, if 1-4 hold, then

$$f^* \leq f(\mathbf{x}^*) \stackrel{\text{by (4)}}{=} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}^*) = \min_{\mathbf{x} \in \Omega} \mathcal{L}(\mathbf{x}, \boldsymbol{\mu}^*) \stackrel{\text{by (3)}}{=} q(\boldsymbol{\mu}^*) \leq q^*$$

By weak duality, the opposite inequality must hold. Therefore $f^* = q^*$, f^* is primal optimal and q^* is dual optimal.

- **Theorem (Saddle point):** $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ is an optimal solution/geometric multiplier pair of the problem above if and only if $\mathbf{x}^* \in \Omega, \boldsymbol{\mu}^* \geq \mathbf{0}$ and $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ is a *saddle point* of the Lagrangian in the sense that

$$\mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}) \leq \mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}^*) \leq \mathcal{L}(\mathbf{x}, \boldsymbol{\mu}^*) \quad \forall \mathbf{x} \in \Omega, \boldsymbol{\mu} \geq \mathbf{0}$$

Again, this theorem is only useful if there is no duality gap.

Proof: Two parts:

- Optimal \Rightarrow Saddle If $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ is an optimal solution/geometric multiplier pair, then from optimality condition (3), we have

$$\mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}^*) \leq \mathcal{L}(\mathbf{x}, \boldsymbol{\mu}^*) \quad \forall \mathbf{x} \in \Omega$$

Furthermore, since $\mathbf{g}(\mathbf{x}) \leq \mathbf{0}$, for any $\boldsymbol{\mu} \geq \mathbf{0}$ we must have $\mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}) \leq f(\mathbf{x}^*) = \mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}^*)$.

- Saddle \Rightarrow Optimal If $\mathbf{x}^* \in \Omega, \boldsymbol{\mu}^* \geq \mathbf{0}$ and $(\mathbf{x}^*, \boldsymbol{\mu}^*)$ is a saddle point, then

$$\sup_{\boldsymbol{\mu} \geq \mathbf{0}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}) = \sup_{\boldsymbol{\mu} \geq \mathbf{0}} \left\{ f(\mathbf{x}^* + \boldsymbol{\mu}^\top \mathbf{g}(\mathbf{x}^*)) \right\} = \begin{cases} f(\mathbf{x}^*) & \text{if } \mathbf{g}(\mathbf{x}^*) \leq \mathbf{0} \\ +\infty & \text{otherwise} \end{cases}$$

We show that each optimality condition holds:

- Condition 1: By the definition of a saddle point, $\mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}) \leq \mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}^*)$. Therefore, the first of the two options for $\sup_{\boldsymbol{\mu} \geq \mathbf{0}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu})$ options must be correct, and $\mathbf{g}(\mathbf{x}^*) \leq \mathbf{0}$.
- Condition 2: The statement of the theorem says that $\boldsymbol{\mu}^* \geq \mathbf{0}$.
- Condition 3: Obvious, by the definition of a saddle point.
- Condition 4: We have established that the first option for $\sup_{\boldsymbol{\mu} \geq \mathbf{0}} \mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu})$ above is correct. Thus, $\mathcal{L}(\mathbf{x}^*, \boldsymbol{\mu}^*) = f(\mathbf{x}^*)$ and this means that $\boldsymbol{\mu}^{*\top} \mathbf{g}(\mathbf{x}^*) = 0$. But since $\boldsymbol{\mu}^* \geq \mathbf{0}$ and $\mathbf{g}(\mathbf{x}^*) \leq \mathbf{0}$, we get $\mu_j^* g_j(\mathbf{x}^*) = 0$.

- **Extension to equality constraints**

- The theory above can be extended to equality constraints by simply introducing a pair of inequality constraints for each equality constraints. Thus, $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ becomes $\mathbf{h}(\mathbf{x}) \geq \mathbf{0}$ and $\mathbf{h}(\mathbf{x}) \leq \mathbf{0}$.
- It turns out this is equivalent to simply eliminating the non-negativity constraints on the Lagrange multipliers for equality constraints. See the discussion of the geometrical interpretation of the KKT conditions for some intuition behind this result.

- **Strong Duality**

- **Theorem (Linear Constraints)**: Consider the following problem:

$$(P) : \min_{\mathbf{x}} f(\mathbf{x}) \text{ s.t. } A\mathbf{x} \leq \mathbf{b}, \mathbf{x} \in \mathbb{R}^n$$

The Lagrangian is $\mathcal{L}(\mathbf{x}, \boldsymbol{\mu}) = f(\mathbf{x}) + \boldsymbol{\mu}^\top (A\mathbf{x} - \mathbf{b})$ and $q(\boldsymbol{\mu}) = \inf_{\mathbf{x} \in \mathbb{R}^n} \mathcal{L}(\mathbf{x}, \boldsymbol{\mu})$ and the dual is

$$\max_{\boldsymbol{\mu}} q(\boldsymbol{\mu}) \text{ s.t. } \boldsymbol{\mu} \geq \mathbf{0}, \boldsymbol{\mu} \in \mathbb{R}^m$$

If f is *convex* over \mathbb{R}^n and *continuously differentiable*, then if the primal has an optimal solution, then there is no duality gap, and at least one geometric multiplier exists.

Proof: Let \mathbf{x}^* be an optimal solution for the primal. Then, by the KKT conditions, there exists $\boldsymbol{\mu}^* \in \mathbb{R}^m$ such that

$$\boldsymbol{\mu}^* \geq \mathbf{0} \quad \boldsymbol{\mu}^{*\top} (A\mathbf{x}^* - \mathbf{b}) = 0 \quad \nabla f(\mathbf{x}^*) + A\boldsymbol{\mu}^* = \mathbf{0}$$

(We do not need to check for regularity, since the constraints are linear).

Now, since $\mathcal{L}(\mathbf{x}, \boldsymbol{\mu})$ is convex and $\nabla f(\mathbf{x}^*) + A\boldsymbol{\mu}^* = \mathbf{0}$, we have that

$$\mathbf{x}^* \in \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^n} \mathcal{L}(\mathbf{x}, \boldsymbol{\mu}^*)$$

As such

$$f(\mathbf{x}^*) = \min_{\mathbf{x} \in \mathbb{R}^n} \{f(\mathbf{x}) + \boldsymbol{\mu}^{*\top}(A\mathbf{x} - \mathbf{b})\} = q(\boldsymbol{\mu}^*)$$

By weak duality, however, we have that

$$q(\boldsymbol{\mu}^*) \leq q^* \leq f^* \leq f(\mathbf{x}^*)$$

However, by the previous statement, $f(\mathbf{x}^*) = q(\boldsymbol{\mu}^*)$, equality holds throughout, and so $f^* = q^*$.

- **Extension to equality constraints:** The above trivially extends to linear equality constraints. More generally, consider the problem

$$\min_{\mathbf{x} \in \mathbb{R}^n} f(\mathbf{x}) \text{ s.t. } A\mathbf{x} = \mathbf{b}, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}$$

With Lagrangian

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = f(\mathbf{x}) + \boldsymbol{\lambda}^\top(A\mathbf{x} - \mathbf{b}) + \boldsymbol{\mu}^\top \mathbf{g}(\mathbf{x})$$

(With $\boldsymbol{\mu} \geq \mathbf{0}$ and $\boldsymbol{\lambda}$ unrestricted).

Then, provided that

- There exists an optimal solution \mathbf{x}^*
- f and \mathbf{g} are continuously differentiable
- There exists multipliers $(\boldsymbol{\lambda}^*, \boldsymbol{\mu}^*)$ satisfying the KKT conditions (ie: some sort of regularity).
- f and \mathbf{g} are convex over \mathbb{R}^n

Then there is no duality gap and geometric multipliers exists.

(Note that we are effectively requiring inequality constraints to be convex and equality constraints to be linear. One way to look at this requirement is as a requirement that \mathcal{L} be convex. Indeed, since $\boldsymbol{\mu}$ is positive, $\boldsymbol{\mu}^\top \mathbf{g}(\mathbf{x})$ is convex for all \mathbf{g} . Since $\boldsymbol{\lambda}$ can take any value, however, it needs to multiply a linear function to retain convexity).

- **Theorem (Slater's Condition):** Consider the problem

$$\min_{\mathbf{x}} f(\mathbf{x}) \text{ s.t. } \mathbf{g}(\mathbf{x}) \leq \mathbf{0}, \mathbf{x} \in \Omega$$

Then, suppose that

- The problem is bounded, in other words

$$-\infty < f^* = \inf_{\mathbf{x} \in \Omega, \mathbf{g}(\mathbf{x}) \leq \mathbf{0}} f(\mathbf{x})$$

- The set Ω is convex, and f and \mathbf{g} are convex over Ω .
- There exists a vector $\bar{\mathbf{x}}$ with $\mathbf{g}(\bar{\mathbf{x}}) < \mathbf{0}$

Then there is no duality gap, and there exists at least one geometric multiplier.

Proof: Define

$$\mathcal{A} = \{(\mathbf{z}, w) \in \mathbb{R}^{r+1} : \exists \mathbf{x} \in \Omega \text{ with } \mathbf{g}(\mathbf{x}) \leq \mathbf{z}, f(\mathbf{x}) \leq w\}$$

By the convexity of \mathbf{g} , f and Ω , this set is convex.

Note also that $(\mathbf{0}, f^*)$ is not in the interior of \mathcal{A} . Otherwise, for some $\varepsilon > 0$, $(\mathbf{0}, f^* - \varepsilon) \in \mathcal{A}$, which contradicts the definition of f^* .

By the supporting hyperplane theorem, there exists a normal vector $(\boldsymbol{\mu}, \beta) \neq (\mathbf{0}, 0)$ such that

$$\beta f^* \leq \beta w + \boldsymbol{\mu} \cdot \mathbf{z} \quad \forall (\mathbf{z}, w) \in \mathcal{A}$$

Now, consider; if $(\mathbf{z}, w) \in \mathcal{A}$, then $(\mathbf{z}, w + \gamma) \in \mathcal{A}$, for all $\gamma \geq 0$. Thus, $\beta \geq 0$. Similarly, $\boldsymbol{\mu} \geq \mathbf{0}$.

Consider, however, that if $\beta = 0$, the equation above becomes

$$0 \leq \boldsymbol{\mu} \cdot \mathbf{z} \quad \forall (\mathbf{z}, w) \in \mathcal{A}$$

But since $(\mathbf{g}(\bar{\mathbf{x}}), f(\bar{\mathbf{x}})) \in \mathcal{A}$, this would mean that

$$0 \leq \boldsymbol{\mu} \cdot \mathbf{g}(\bar{\mathbf{x}}) \quad \forall (\mathbf{z}, w) \in \mathcal{A}$$

Since, by definition, $\mathbf{g}(\bar{\mathbf{x}}) < \mathbf{0}$, this means that $\boldsymbol{\mu} = \mathbf{0}$. This is a contradiction, which means that we must have $\beta > 0$.

We can therefore divide by β and normalize, so that $\beta = 1$; we then get


$$\begin{aligned} f^* &\leq w + \boldsymbol{\mu} \cdot \mathbf{z} && \forall (\mathbf{z}, w) \in \mathcal{A} \\ \Rightarrow f^* &\leq f(\mathbf{x}) + \boldsymbol{\mu} \cdot \mathbf{g}(\mathbf{x}) && \forall \mathbf{x} \in \Omega \end{aligned}$$

Minimizing over $\mathbf{x} \in \Omega$, we get


$$f^* \leq \inf_{x \in \Omega} [f(x) + \boldsymbol{\mu} \cdot \mathbf{g}(x)] = q(\boldsymbol{\mu}) \leq q^*$$

Thus, by weak duality, $f^* = q^*$ and $\boldsymbol{\mu}$ is a geometric multiplier.

Diagrammatically, if the set $\mathcal{S} = \{(\mathbf{g}(x), f(x)) : x \in \Omega\}$ and the function $f(x)$ look like this:



Then the set $\mathcal{A} = \{(z, w) \in \mathbb{R}^{r+1} : \exists x \in \Omega \text{ with } \mathbf{g}(x) \leq z, f(x) \leq w\}$ looks like




The line must be sloping as above (ie: $\boldsymbol{\mu}$ and β must be positive). Now, consider – if there is a point \bar{x} such that $g(\bar{x}) < 0$, it means $g_{\min} < 0$. Therefore, we can't have $\beta = 0$ (ie: the line can't be vertical). Now, the Lagrangian is basically the projection of any point in this set onto the vector $(\boldsymbol{\mu}, 1)$ – which is clearly minimized at f^* .

- o Note that the requirement that there be a $g(\bar{x}) < 0$ is crucial. In particular, this condition is *not* satisfied for problems involving equality constraints. These are harder to deal with and require a new definition

- **Definition (Relative interior):** Suppose $\mathcal{C} \subset \mathbb{R}^n$ is a convex set. The *relative interior* of \mathcal{C} is the set $\text{relint } \mathcal{C}$ of all $\mathbf{x} \in \mathbb{R}^n$ for which there exists an $\varepsilon > 0$ such that if $\mathbf{z} \in \text{aff } \mathcal{C}$ with $\|\mathbf{z} - \mathbf{x}\| < \varepsilon$, then $\mathbf{z} \in \mathcal{C}$.

For example, consider the set $\mathcal{C} = [0,1] \in \mathbb{R}^2$. This set has no interior, because there is no \mathbb{R}^2 ball lying totally in the set. However, it has a relative interior:



- **Theorem (Slater's Conditions with Mixed Constraints):** Consider the program (E is a matrix)

$$\min_{\mathbf{x} \in \Omega} f(\mathbf{x}) \text{ s.t. } \mathbf{g}(\mathbf{x}) \leq \mathbf{0}, A\mathbf{x} \leq \mathbf{b}, E\mathbf{x} = \mathbf{d}$$

Suppose that the optimal value f^* is finite and that

- Ω is the intersection of a convex set \mathcal{C} and a polyhedron
- The functions f and \mathbf{g} are convex over Ω
- There is a feasible vector $\bar{\mathbf{x}}$ with $\mathbf{g}(\bar{\mathbf{x}}) < \mathbf{0}$
- There is a vector \mathbf{x} with $A\mathbf{x} \leq \mathbf{b}$, $E\mathbf{x} = \mathbf{d}$, $\mathbf{x} \in \text{relint } \mathcal{C}$ and $\mathbf{x} \in \Omega$

Then there is no duality gap, and there exists at least one Lagrange multiplier.