Static Optimization

Stéphane Dupraz

An optimization problem in $\mathbb{R}^n$ consists in maximizing or minimizing an **objective function** $f : \mathbb{R}^n \to \mathbb{R}$ over a **constraint set** $\mathcal{D} \subseteq \mathbb{R}^n$:

$$\max_{x \in \mathbb{R}^n} / \min f(x) \text{ s.t. } x \in \mathcal{D}$$

Because minimizing f is equivalent to maximizing $-f$, all results on minimization are easy to derive from results on maximization, and from now on we restrict to maximization problems.

Our goal is to determine not only the maximum of f, but also the set of its maximizers $argmax(f(x), x \in \mathcal{D})$. A solution—a maximum and maximizer—does not necessarily exist, and when it does a maximizer is not necessarily unique. We have seen however that Weierstrass theorem guarantees the existence of solutions if $\mathcal{D}$ is compact—closed and bounded in $\mathbb{R}^n$—and f is continuous.

Weierstrass theorem (bis repetita) Let $f : \mathcal{D} \to \mathbb{R}^n$, $\mathcal{D} \subseteq \mathbb{R}$.

If f is continuous and $\mathcal{D}$ is compact (closed and bounded), then f has a maximum.

Although there is no reason to abandon looking for maxima if we cannot first establish their existence, it is a good rule to always start a maximization problem by checking whether Weierstrass theorem guarantees the existence of solutions.

In this chapter, we go beyond existence theorems and cover methods to explicitly find the solutions of an optimization problem. All the methods we cover rely on differentiation: in all that follows, we will assume that f is at least once-differentiable. We consider three types of maximization problems, depending on the constraint set $\mathcal{D}$:

1. Optimization on an open set: if $\mathcal{D}$ is an open set.
2. Equality constraints: if $\mathcal{D}$ can be written as $\mathcal{D} = \{x / g_1(x) = 0, ..., g_k(x) = 0\}$.
3. Inequality constraints: if $\mathcal{D}$ can be written as $\mathcal{D} = \{x / g_1(x) \geq 0, ..., g_k(x) \geq 0\}$.

1 Optimization on an open set

We first consider the case where the constraint set $\mathcal{D}$ is an open set. This is sometimes referred to as **unconstrained optimization**, which seems to restrict the case to the absence of constraints $\mathcal{D} = \mathbb{R}^n$. Instead, the terminology reflects that the constraints *"do not matter"* when $\mathcal{D}$ is an open set: when maxima are interior points of $\mathcal{D}$, the methods to find possible maxima are the same as if the problem had no constraints.

Note that if $\mathcal{D}$ is an open set of $\mathbb{R}^n$, it is not compact[1], so that we cannot use Weierstrass theorem to guarantee the existence of a maximum.

1.1 Necessary First-order Conditions

Assessing whether a function f reaches a maximum at a point x is not easy. Comparing $f(x)$ to the values that f takes at all other points is impossible. The method consists in restricting to differentiable functions and considering the variations of f around x as summed up by its derivative. As the derivative at a point x only informs about local variations, the method actually selects not only global maxima, but also local ones.

Definition 1.1. *Let (X, d) be a metric space and $f : X \to \mathbb{R}$ a function.*

- *f has a **local maximum at** $x_0 \in X$ iff there exists an open ball $B(x_0, \delta)$ around x_0 on which f has a maximum at x_0:*

$$\exists \delta / d(x, x_0) < \delta \Rightarrow f(x) \leq f(x_0).$$

- *f has a **strict local maximum at** $x_0 \in X$ iff there exists an open ball $B(x_0, \delta)$ around x_0 on which f has a maximum at x_0 and only at x_0:*

$$\exists \delta / 0 < d(x, x_0) < \delta \Rightarrow f(x) < f(x_0).$$

Obviously, a maximum—called a **global maximum** when there is ambiguity—is a local maximum. You can think of a local maximum as the top of a hill (and of Mount Everest as Earth's global maximum). Just before arriving to the top of the hill, you're going up, and just after, you're going down. So the top of the hill is a point where you're not going neither up nor down. This is the idea of the first-order condition:

[1] This is not obvious but it turns out that $\emptyset$ and $\mathbb{R}^n$ are the only two sets of $\mathbb{R}^n$ that are both open and closed. Since $\mathbb{R}^n$ is unbounded, no open set in $\mathbb{R}^n$ is compact.

Proposition 1.1. ***Necessary first-order conditions.***

Let $\mathcal{D}$ be an open set of $\mathbb{R}^n$, and $f : \mathcal{D} \to \mathbb{R}$ a differentiable function.

If f has a local maximum at x, then $f'(x) = 0$, i.e. all the partial derivatives are nil, $f'_i(x) = 0$ for all i.

We say that x is a **critical point** or **stationary point** of f when $f'(x) = 0$. So an alternative way of stating the result is : if x is a local maximum, then it is a critical point.

Proof. We show that $f'(x) = 0$ by proving $f'(x).h = 0$ for all h, that is that all directional derivatives are nil (showing it for the n partial derivatives would be enough). Fix a direction h. Because f has a local maximum at x, $f(x + th) - f(x) \leq 0$ for $|t|$ smaller than a δ. Consider the function $t \mapsto \frac{f(x+th)-f(x)}{t}$ for $t \in [0, \delta]$. It is negative, hence so is its limit at $t = 0$, $f'_h(x) \leq 0$. Considering the function $t \mapsto \frac{f(x+th)-f(x)}{t}$ for $t \in [-\delta, 0]$, we get the reverse equality $f'_h(x) \geq 0$. Hence $f'_h(x) = 0$. QED. □

Be careful that in general the first-order conditions select more than just the local maxima: being a critical point is a necessary but not sufficient condition for being a local maximum—let alone a global maximum. First, the set of critical points also contains local (including global) minima. But there are also critical points that are neither local maximum nor local minimum: they are called **saddle points**. As a counter-example, consider the function $x \mapsto x^3$; it has a saddle point at $x = 0$. The figure below sums up the different types of critical points. (Note that a point x can in principle be a local maximum and local minimum if f plateaus at x, and a global maximum and global minimum if f is constant).

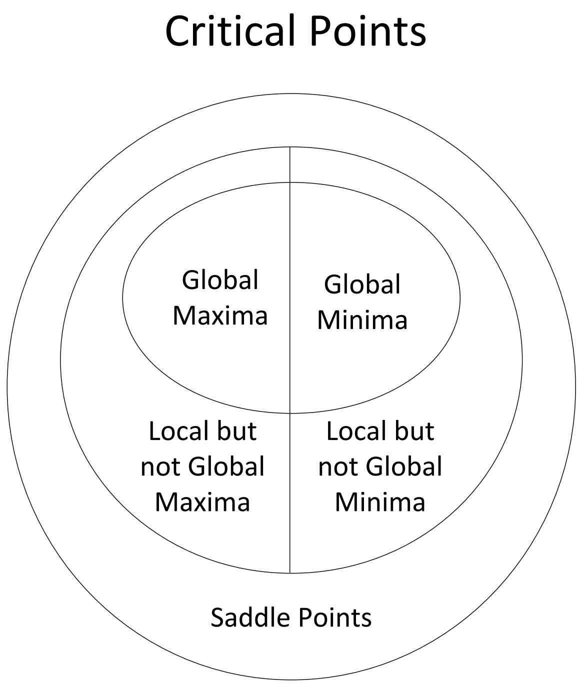

1.2 Necessary Second-order Conditions

For $\mathcal{C}^2$ functions, second-order conditions gives additional necessary conditions that differ between local maxima and local minima.

Proposition 1.2. ***Necessary second-order conditions.***
Let $\mathcal{D}$ be an open set of $\mathbb{R}^n$, and $f : \mathcal{D} \to \mathbb{R}$ a $\mathcal{C}^2$ function.

- *If f has a local maximum at x, then the hessian of f at x $H(x)$ is negative semi-definite.*
- *If f has a local minimum at x, then the hessian of f at x $H(x)$ is positive semi-definite.*

Proof. We *admit* the result. See e.g. Sundaram theorem 4.3 for a proof. □

The necessary second-order conditions are not extremely useful in practice though: the only thing they do is to restrict a bit more the set of possible candidates for a global maximum, excluding the points at which the hessian is not negative semi-definite. Typically however, the first-order conditions have already reduced the candidates to only a few, so that the second-order necessary conditions are not very useful.

Two final remarks. First, note that even when $\mathcal{D}$ is not an open set, we can apply the method to select the possible maxima on the interior of $\mathcal{D}$. Second, be careful that the necessary first and second order conditions do not allow us to conclude that a point is a maximum: they only identify possible candidates. We gather all sufficient conditions for local and global maxima at the end of the chapter.

2 Equality-constrained problems: the Theorem of Lagrange

The method of the previous section requires the maxima to be interior points of $\mathcal{D}$. This is obviously guaranteed if $\mathcal{D}$ is open, but more generally the methods can be used to find the maxima in the interior of $\mathcal{D}$ for any constraint set $\mathcal{D}$. However, this leaves out the maxima that lie on the frontier. We now extend the method of the previous section into methods that can find maxima in constraint sets that have a frontier. The first case we deal with is the case of equality constrains $\mathcal{D} = \{x, g_1(x) = 0, ..., g_k(x) = 0\}$.

Let us start with the issue of existence of maxima. As long as the function $g = (g_1, ..., g_k)'$ is continuous, $\mathcal{D}$ is closed since $\mathcal{D}$ is the inverse image of the closed set $\{0\}$ by a continuous function. If $\mathcal{D}$ is also bounded, then $\mathcal{D}$ is compact. If in addition f is continuous, the problem has a solution by Weierstrass theorem.

2.1 Can't we just "replace the constraints into the objective"?

Consider a canonical maximization problem, the utility maximization problem of a consumer. We assume that a consumer has preferences over n goods that can be represented by a utility function $U : \mathbb{R}^n \to \mathbb{R}$. The consumer decides how much to consume of each good subject to his budget constraint $\sum_{i=1}^{n} p_i x_i = w$, where p_i is the price of the good i and w is the consumer's income. Assuming the consumer behaves to get most satisfaction out of his consumption, his consumption is solution to the program:

$$\max_{x\in\mathbb{R}^n} U(x_1, ..., x_n) \text{ s.t. } \sum_{i=1}^{n} p_i x_i = w$$

The novelty relative to the previous section is in the constraint $\sum_i p_i x_i = w$. But do we really need a new method? Can't we just use the constraint to replace one variable in the objective and end up with an unconstrained program? For instance, we could replace x_n through $x_n = \frac{1}{p_n}\left(w - \sum_{i=1}^{n-1} p_i x_i\right)$ and use the method of the previous section to find the maxima of the unconstrained problem in $n-1$ variables $x_1, ..., x_{n-1}$:

$$\max_{x_1,...,x_{n-1}\in\mathbb{R}^{n-1}} U\left(x_1, ..., x_{n-1}, \frac{1}{p_n}\left(w - \sum_{i=1}^{n-1} p_i x_i\right)\right)$$

The first-order conditions give:

$$\forall i = 1, ..., n-1, U_i'(x) - \frac{p_i}{p_n} U_n'(x) = 0$$

The answer is: we absolutely can—in this example. But this is because the constraint lets us express x_n as an explicit function of $x_1, ..., x_{n-1}$. Remember from the discussion of the implicit function theorem that this is not always the case: sometimes we cannot turn the equation $g(x) = 0$ into an explicit function of k variables as a function of the $n-k$ others. The theorem of Lagrange lets us deal with constraints that only give implicit relationships between the optimization variables. As the proof of the theorem of Lagrange will make clear however, the idea of the theorem is precisely to replace k variables into the objective using the constraints. Its only sophistication is to use the implicit function theorem to do the replacement.

2.2 The theorem of Lagrange

Theorem 2.1. ***The Theorem of Lagrange (necessary first-order conditions)***
Let $f : \mathbb{R}^n \to \mathbb{R}$ be $\mathcal{C}^1$, and $g = (g_1, ..., g_k)' : \mathbb{R}^n \to \mathbb{R}^k$ be $\mathcal{C}^1$ functions, and consider the program:

$$\max_{x \in \mathbb{R}^n} f(x) \text{ s.t. } g_i(x) = 0 \text{ for all } i = 1, ..., k$$

*If x is a local maximum of f on $\mathcal{D}$ and x satisfies the **constraint qualification** $rank(g'(x)) = k$, then there exist k **Lagrange multipliers** $\lambda = (\lambda_1, ..., \lambda_k)' \in \mathbb{R}^k$ such that the first-order condition holds:*

$$f'(x) + \lambda' g'(x) = f'(x) + \sum_{i=1}^{k} \lambda_i g_i'(x) = 0.$$

*A point and associated Lagrange multiplier (x, λ) that satisfy the first order condition can be seen as a critical point of the **Lagrangian** defined as:*

$$L(x, \lambda) = f(x) + \lambda' g(x) = f(x) + \sum_{i=1}^{k} \lambda_i g_i(x)$$

(The derivative of L with respect to λ is simply the constraint $g(x) = 0$).

Proof. The idea is to apply the implicit function theorem to g in order to replace k variables in the objective using the constraints, and then take the first order condition of the resulting unconstrained program. To this end, let us divide the vector x between its $n-k$ first variables and its k last variables which we are going to replace: $x = (x_1, x_2)'$. Let $x^* = (x_1^*, x_2^*)$ be a local maximum of f.

Let us check the assumptions of the implicit function theorem. First, g is $\mathcal{C}^1$. Second, $g(x_1^*, x_2^*) = 0$, since $x^* \in \mathcal{D}$. Third, we need $g_2'(x_1^*, x_2^*)$ to be invertible. But this is precisely what the constraint qualification guarantees, provided a reordering of the variables in x^*. So by the implicit function theorem, we know there exists a $\mathcal{C}^1$ function $h : B(x_1^*, r) \to \mathbb{R}^k$ such that $x_2^* = h(x_1^*)$, $g(x_1, h(x_1)) = 0$ for all $x_1 \in B(x_1^*, r)$, and $h'(x_1^*) = -g_2'(x^*)^{-1} g_1'(x^*)$.

The assumption that x^* is a local maximum of f can now be phrased as saying that x_1^* is a local maximum of the $\mathcal{C}^1$ function $F(x_1) = f(x_1, h(x_1))$ on the open ball $B(x_1^*, r)$ without constraints. From the previous section, we know that x_1^* must satisfy the first-order condition:

$$\begin{aligned} f_1'(x^*) + f_2'(x^*)h'(x_1^*) &= 0 \\ f_1'(x^*) + f_2'(x^*)\Big(- g_2'(x^*)^{-1} g_1'(x^*)\Big) &= 0 \\ f_1'(x^*) + \Big(- f_2'(x^*) g_2'(x^*)^{-1}\Big) g_1'(x^*) &= 0 \end{aligned}$$

Defining $\lambda^* \equiv -f_2'(x^*)g_2'(x^*)^{-1}$, we get:

$$\begin{aligned} f_1'(x^*) + \lambda^* g_1'(x^*) &= 0 \\ f_2'(x^*) + \lambda^* g_2'(x^*) &= 0 \end{aligned}$$

Which is precisely the result. □

The constraint qualification guarantees that there exists a way to replace k variables by $n - k$ in the objective. If it is not satisfied at a local maximum, then a local maximum may well fail to satisfy the first-order condition stated in the Lagrange theorem. To select all possible maxima, we must therefore also look for the points where the constraint qualification fails, as they are possible maxima even if they do not satisfy the first-order conditions.

As an example, let us solve the consumer problem using the theorem of Lagrange. The constraint is $g(x_1, ..., x_n) = w - \sum_i p_i x_i$. Since $g'(x) = p'$, which has rank 1 assuming $p \neq 0$, the constraint qualification is satisfied everywhere. We form the Lagrangian:

$$L(x, \lambda) = U(x) + \lambda\Big(w - \sum_{i=1}^{n} p_i x_i\Big)$$

And look for solutions (x, λ) to the system:

$$\begin{aligned} \forall i = 1, ..., n, U_i'(x) &= \lambda p_i \\ \sum_{i=1}^{n} p_i x_i &= w \end{aligned}$$

We can eliminate λ using one first-order condition, for instance for x_n: $\lambda = \frac{U_i'(x)}{p_i}$. We end up with:

$$\forall i = 1, ..., n-1, U_i'(x) = \frac{p_i}{p_n} U_n'(x)$$
$$\sum_{i=1}^{n} p_i x_i = w$$

which is the same system as the one we obtain replacing the constraint in the objective.

Even in cases such as the consumer problem where we can do without the method of Lagrange, the method of Lagrange is to be preferred for interpretation purposes. First, it treats all variables symmetrically, which is usually easier to handle. Second and most importantly, we will see below that the Lagrange multipliers have an economic interpretation that makes the first-order conditions of the method of Lagrange economically insightful.

3 Inequality-constrained problems: the Theorem of Kuhn-Tucker

We now consider inequality-constrained programs: $\mathcal{D} = \{x, g_1(x) \geq 0, ..., g_k(x) \geq 0\}$. A word about existence again. As long as the function $g = (g_1, ..., g_k)'$ is continuous, $\mathcal{D}$ is closed since $\mathcal{D}$ is the inverse image of the closed set $(\mathbb{R}_+)^k$ by a continuous function. If $\mathcal{D}$ is also bounded, then $\mathcal{D}$ is compact. If in addition f is continuous, the problem has a solution by Weierstrass theorem.

With inequality constraint, the question appears of whether a constraint binds or not at a point $x \in \mathcal{D}$. We say that the constraint $g_i(x) \geq 0$ is **effective** or **active** or **binding** at x^* if it holds with equality at x^*, $g_i(x^*) = 0$. Note that an equality constraint $g(x) = 0$ can always be seen as two inequality constraints $g(x) \geq 0$ and $-g(x) \geq 0$, so that this section generalizes the previous one: the theorem of Kuhn and Tucker (and Karush) is going to generalize the theorem of Lagrange. The result is very similar, but there are three things to be careful about:

1. The signs of the Lagrange multipliers now matter.
2. Complementary slackness conditions are added to the necessary conditions.
3. The constraint qualification applies to the effective constraints only.

Theorem 3.1. ***The Theorem of Kuhn and Tucker (necessary first-order conditions)***
Let $f : \mathbb{R}^n \to \mathbb{R}$, and $g = (g_1, ..., g_k)' : \mathbb{R}^n \to \mathbb{R}^k$ be $\mathcal{C}^1$ functions, and consider the program:

$$\max_{x \in \mathbb{R}^n} f(x) \ \text{s.t.} \ g_i(x) \geq 0 \text{ for all } i = 1, ..., k$$

Let x be a local maximum of f on $\mathcal{D}$.
Let $g_{i_1}, ..., g_{i_e}$ be the e effective constraints at x; assume the constraint qualification $rank(g'_{i_1}(x), ..., g'_{i_e}) = e$ holds.
*Then there exist k Lagrange multipliers $\lambda \in \mathbb{R}^k$ such that the **Kuhn-Tucker conditions** hold:*

1. *The first-order conditions:*

$$L'_x(x, \lambda) = f'(x) + \lambda' g'(x) = f'(x) + \sum_{i=1}^{k} \lambda_i g'_i(x) = 0.$$

2. *The constraints are satisfied, all Lagrange multipliers are weakly positive, and the **complementary slackness conditions** are satisfied, meaning that for each i, either the constraint is effective at x,*

or $\lambda_i = 0$ *(or both).*

$$\forall i, g_i(x) \geq 0, \lambda_i \geq 0 \text{ and } g_i(x)\lambda_i = 0.$$

Proof. Part of the theorem of Kuhn and Tucker can be derived from the theorem of Lagrange by distinguishing cases depending on whether constraints bind or not. But showing that the Lagrange multipliers are positive requires more work. We *admit* the result. □

Solving the system defined by the Kuhn-Tucker conditions is usually laborious as it requires distinguishing cases based on which set of constraints is effective or not. Besides, just as with equality constraints, we must be careful to also look for points where the constraint qualification may fail. Finally, be careful how to write the constraints in the Lagrangian now that the sign of the Lagrange multipliers matter: if the constraint is $g_i(x) \geq 0$, then add $+\lambda_i g_i(x)$ to have $\lambda_i \geq 0$.

A final word about the case of **mixed constraints**, meaning problems with both equality and inequality constraints. As mentioned above, an equality constraint $g_i(x) = 0$ can always be seen as two inequality constraints $g_i(x) \geq 0$ and $g_i(x) \leq 0$. This would suggest to add two terms $\lambda_i g_i(x) - \mu_i g_i(x) = (\lambda_i - \mu_i) g_i(x)$, $\lambda_i, \mu_i \geq 0$ to the Lagrangian. Obviously, this is just akin to adding a single term $\lambda_i g_i(x)$ and dropping the requirement that λ_i be positive. Besides, we know that an equality constraint is effective, so we can drop the complementary slackness. To sum up, we treat equality constraints in mixed constraints programs just as in programs with equality constraints only.

4 Comparative Statics: the Envelope Theorems

4.1 The Envelope Theorems

In many instances, the objective function and/or constraint set of an optimization problem depends on a number of variables which are not maximized upon—for instances prices and income in the consumer problem. We are often interested in how the maximum and maximizers of the problem varies with these exogenous parameters. Let us note θ the vector of exogenous parameters that affects the objective function and the constraints, so that the objective function is now $f(x, \theta)$ and the constraint set $\mathcal{D}(\theta)$. We define the **value function** as the function that to each value of θ associates the maximum of the program for θ (assuming it exists).

$$V(\theta) \equiv \max_{x \in \mathcal{D}(\theta)} f(x, \theta)$$

In the same way, we can define a function that to each θ associates a maximizer $x^*(\theta)$ (if the solution if not unique, we need to pick one maximizer). The value function then appears as:

$$V(\theta) = f(x^*(\theta), \theta)$$

To know how V and x^* vary with θ, we wish to differentiate them. But are V and x^* differentiable? Simply assuming that f is differentiable is not enough to guarantee the differentiability of V and x^*, but there exist conditions on the primitives of the problem f and $\mathcal{D}$ that guarantee differentiability (we do not cover them here). Assuming f is differentiable, and provided $x^*(\theta)$ is differentiable, $V(\theta)$ is differentiable, and using the chain rule:

$$V'(\theta) = f_1'(x^*(\theta), \theta) \times x^{*'}(\theta) + f_2'(x^*(\theta), \theta)$$

This equation expresses in differential form the simple idea that V varies with θ through two effects. First, θ directly affects the function $x \mapsto f(x, \theta)$: for a fixed x, $f(x)$ changes. But there is a second, indirect effect: as θ varies, the point $x^*(\theta)$ at which f reaches its maximum varies, which affects V. The envelope theorems are the simple observations that in the cases we have seen, the second indirect effect is nil at first-order, because of the first-order conditions in the maximization program. This is most obvious in the case of optimization on an open set where the FOC is simply $f_1'(x^*(\theta), \theta) = 0$.

Theorem 4.1. ***Envelope theorem for optimization on an open set***

Consider the parameterized program for all $\theta \in \Theta$*:*

$$V(\theta) \equiv \max_{x \in \mathcal{D}(\theta)} f(x, \theta)$$

Let $x^*(\theta)$ *associates to each* θ *a maximizer of the* θ*-program.*

If f *and* $x^*(\theta)$ *are differentiable,* V *is differentiable and:*

$$V'(\theta) = f'_2(x^*(\theta), \theta)$$

The terminology "envelope theorem" comes from the graphical illustration of the theorem.

The envelope theorem can be generalized to inequality (including equality) constrained optimization programs. Although the value function $V(\theta) \equiv \max_{x/g(x,\theta) \geq 0} f(x, \theta)$ can still be written as:

$$V(\theta) = f(x^*(\theta), \theta),$$

this is no longer very helpful as the first-order condition of the maximization program is no longer $f'_1(x^*(\theta), \theta) = 0$ but $L'_1(x^*(\theta), \theta) = f'_1(x^*(\theta), \theta) + \lambda' g'_1(x^*(\theta), \theta) = 0$. However, since $\lambda' g(x^*(\theta), \theta) = 0$ because of the complementary slackness conditions, we can alternatively write:

$$V(\theta) = L(x^*(\theta), \theta)$$

The generalization of the envelope theorem follows:

Theorem 4.2. ***Envelope theorem for inequality-constrained programs***

Consider the parameterized program for all $\theta \in \Theta$*:*

$$V(\theta) \equiv \max_{x/g(x,\theta) \geq 0} f(x, \theta)$$

Let $L(x, \lambda, \theta) = f(x, \theta) + \lambda' g(x, \theta)$ *be the Lagrangian of the problem.*

Let $x^*(\theta)$ *associates to each* θ *a maximizer of the* θ*-program.*

If f, g and $x^*(\theta)$ are differentiable, V is differentiable and $x^*(\theta)$ is a solution to the KT-system:

$$V'(\theta) = f_2'(x^*(\theta), \theta) + \lambda' g_2'(x^*(\theta), \theta)$$

4.2 Marginal Interpretation of the Lagrange multipliers

An application of the envelope theorem allows to give an interpretation to the Lagrange multipliers. Consider maximizing the function f under the constraint $g(x) \leq c$, $c \in \mathbb{R}$. The real number c parameterizes the program. The Lagrangian of the problem is:

$$L(x, \lambda, c) = f(x) - \lambda(g(x) - c)$$

Under differentiability assumptions, the envelope theorem gives:

$$V'(c) = \lambda$$

In words, the Lagrange multiplier λ appears as the marginal value—in terms of the objective function—of increasing c—or the **marginal value of relaxing the constraint** $g(x) \leq c$.

For instance, consider the consumer program. The single constraint is the budget constraint and c is the consumer's income w. The Lagrange multiplier of the program appears as the marginal value of income. The first-order condition:

$$U_i'(x) = \lambda p_i$$

can now be interpreted as stating that the marginal utility of the good x_i must at the optimum be equal to the marginal value of one dollar of wealth, times the number of dollar that acquiring one additional good i costs. To decide to much to spend on good i, the consumer needs only know the price p_i and λ, which indicates the opportunity cost of spending money of good i instead of anything else.

5 Some Sufficient Conditions

Up to now, we have only covered *necessary* conditions for a points x to be a maximum. But even if a maximum necessarily satisfies the condition, a point that satisfies the condition needs not be a global, not even local maximum. In other words, the method selects possible maxima, but does not allow us to conclude that any of these candidates is a maximum. So far, we cannot claim that we have a method to determine the maxima of a function.

One situation in which we can conclude is when we have proven the existence of maxima independently, for instance through Weierstrass theorem. Then, if there is a single point that satisfies the necessary conditions, we can conclude it is the maximum. If the existence of a maximum has been proven and we have several points that satisfy the necessary conditions, then we can just compare the value of f at these points to determine which one is the global maximum. Be careful that if we are doing constrained optimization, we also need to include as candidates for a maximum the points at which the constraint qualification does not hold.

Otherwise, we need sufficient conditions to conclude. We first see some sufficient conditions to be a local extremum, then sufficient conditions to be a global extremum.

5.1 Sufficient Conditions for Local Extrema

Consider first optimization on an open set. First and second order conditions are almost sufficient for a local extremum. However, whereas the second-order conditions required the hessian of f to be negative *semi-definite*, the sufficient condition below requires the hessian to be negative *definite*. This stronger assumption guarantees x to be a *strict* local maximum.

Proposition 5.1. *Let $\mathcal{D}$ be an open set of $\mathbb{R}^n$, and $f : \mathcal{D} \to \mathbb{R}$ a $\mathcal{C}^2$ function.*

- *If $f'(x) = 0$ and $H(x)$ is negative definite, then x is a strict local maximum.*
- *If $f'(x) = 0$ and $H(x)$ is positive definite, then x is a strict local minimum.*

Proof. We *admit* the result. See e.g. Sundaram theorem 4.3 for a proof. □

The result generalizes to equality-constrained optimization.

Proposition 5.2. *Let $f : \mathbb{R}^n \to \mathbb{R}$, and $g = (g_1, ..., g_k)' : \mathbb{R}^n \to \mathbb{R}^k$ be $\mathcal{C}^2$ functions, and consider the* *program:*

$$\max_{x\in\mathbb{R}^n} f(x) \text{ s.t. } g_i(x) = 0 \text{ for all } i = 1, ..., k$$

- *If* $rank(g'(x)) = k$, $f'(x) + \lambda' g'(x) = 0$ *for some* $\lambda \in \mathbb{R}^k$ *and for all* z *such that* $g'(x)z = 0$, $z' H_L(x) z < 0$, *then* x *is a strict local maximum.*

- *If* $rank(g'(x)) = k$, $f'(x) + \lambda' g'(x) = 0$ *for some* $\lambda \in \mathbb{R}^k$ *and for all* z *such that* $g'(x)z = 0$, $z' H_L(x) z > 0$, *then* x *is a strict local minimum.*

Proof. We *admit* the result. See e.g. Sundaram theorem 5.4 for a proof. □

Remember that we saw in linear algebra how to check such a condition using bordered matrices. Here we would form the bordered matrix:

$$B(x) = \begin{pmatrix} 0 & g'(x) \\ (g'(x))' & H_L(x, \lambda) \end{pmatrix}$$

5.2 Sufficient Conditions for Global Extrema

Concavity and quasi-concavity (convexity and quasi-convexity for minimization problems) play an essential role in sufficient conditions for global maxima. Let us start with some results that do not rely on differentiability assumptions. The first observation is that quasi-concavity puts a lot of restrictions on the set of maximizers. Indeed, the set of maximizers of a function f is its upper-level set of level M, where M is the maximum of f, $\{x/f(x) = M\}$. Since all the upper-level sets of a quasi-concave functions are convex:

Proposition 5.3.

- *If* $f : \mathcal{D} \to \mathbb{R}$ *is quasi-concave (in particular concave), then the set of its maximizers is convex.*

- *If* $f : \mathcal{D} \to \mathbb{R}$ *is strictly quasi-concave (in particular strictly concave), then it has at most one maximizer.*

Proof. The proof of the second part is by contradiction. If x and y, $x \neq y$ are two maximizers of f strictly quasi-concave, then $f(\frac{1}{2}x + \frac{1}{2}y) > \min(f(x), f(y)) = f(x) = f(y)$, which contradict that x and y are maximizers of f. □

Besides, for concave functions and strictly quasi-concave functions (quasi-concavity is not enough), local maxima are necessarily global maxima:

Proposition 5.4.

- *If $f : \mathcal{D} \to \mathbb{R}$ is concave, then any local maximum of f is a global maximum.*
- *If $f : \mathcal{D} \to \mathbb{R}$ is strictly quasi-concave, then any local maximum of f is a global maximum.*

Proof. We prove the contrapositives. Assume x is not a global maximum of f: there exists $y \in \mathcal{D}$ such that $f(y) > f(x)$. But then if f is concave, for any $\lambda \in [0,1)$, $f(\lambda x + (1-\lambda)y) \geq \lambda f(x) + (1-\lambda)f(y) > f(x)$. And if f is strictly quasi-concave, for any $\lambda \in [0,1)$, $f(\lambda x + (1-\lambda)y) > f(x)$. So in both cases, however close to x (however close λ is to 1), there is a point that gives a strictly higher value of f: x is not a local maximum. □

If f is in addition differentiable, concavity means that for all x, y, $f(y) - f(x) \leq f'(x)(y-x)$. It follows that if x is a critical point, $f(x) \geq f(y)$ for all y. In other words, the first-order condition for an interior maximum become also sufficient.

Proposition 5.5. *Let $f : \mathcal{D} \to \mathbb{R}$, $\mathcal{D}$ an open set of $\mathbb{R}^n$.*

If f is differentiable and concave, then x is a maximum of f iff $f'(x) = 0$.

The result actually generalizes to unconstrained optimization. In this case, it is actually possible to require only the quasi-concavity of f, provided $f'(x) \neq 0$.

Proposition 5.6. *Let $f : \mathbb{R}^n \to \mathbb{R}$, and $g = (g_1, ..., g_k)' : \mathbb{R}^n \to \mathbb{R}^k$ be $\mathcal{C}^1$ functions, and consider the program:*

$$\max_{x \in \mathbb{R}^n} f(x) \text{ s.t. } g_i(x) \geq 0 \text{ for all } i = 1, ..., k$$

Let x be a point that satisfies the Kuhn-Tucker conditions. If:

- *the g_i's are quasi-concave.*
- *f is concave OR (f is quasi-concave AND $f'(x) \neq 0$)*

then x is a maximizer.

Proof. We *admit* the result. See e.g. Sundaram theorem 8.13 for a proof. □

6 An Introduction to Dynamic Programming

6.1 Static vs. Dynamic Optimization

In the consumer problem we have used as an illustration, the consumption of all goods is taking place simultaneously: the consumer decides what goods to spend his income on at a given date. Another central decision a consumer has to face is when to spend his income: the **consumption/saving decision**. Assume the consumer lives between 0 and T and that its utility depends on its consumption at all periods $U(c_0, c_1, ..., c_T)$ (this time we abstract from the decision of what goods to consume at each period and do as if there existed a single good at each period). At each period, the consumer receives an exogenous income y_t. Thanks to the existence of a financial market however, the consumer needs not consume his income at each period: he can save and borrow at an interest rate r to disentangle its consumption at t from its income at t. Noting s_t the household's savings at t, he starts each period with a financial wealth $(1+r)s_{t-1}$ inherited from his past saving behavior, plus his income y_t. He divides it between his consumption c_t and savings s_t. In other words, its **flow budget constraints** at every period are:

$$c_t + s_t = y_t + (1+r)s_{t-1}, \text{ for all } t = 0, ..., T.$$

A last constraint is that the consumer must have repaid its debts at period T:

$$s_T \geq 0$$

The consumer chooses c_t and s_t at each period $t = 0, ..., T$ in order to maximize his utility $U(c_0, c_1, ..., c_T)$. To complete the description of the problem, we need to specify the exogenous financial wealth s_{-1} the consumer starts with in period 0. Based on the story-telling behind the two consumer problems, it seems appealing to call the first problem **static**, and the second one **dynamic**. And this is indeed what we do. From a formal perspective however, the two problems are not different. We can easily turn the consumption/savings decision into a problem of the same form as the first one. Multiply each flow budget constraint at t by $(\frac{1}{1+r})^t$ and sum them all. This way, the savings variables s_t cancel out and we obtain the **intertemporal budget constraint**:

$$\left(\frac{1}{1+r}\right)^T s_T + \sum_{\tau=0}^{T} \left(\frac{1}{1+r}\right)^\tau c_\tau = \sum_{\tau=0}^{T} \left(\frac{1}{1+r}\right)^\tau y_\tau + s_{-1}(1+r)$$

The intertemporal budget constraint constraint has exactly the form $p'x = w$ of the first consumer problem, if we define the income as the **discounted sum of life-time income**, plus initial financial wealth:

$$w \equiv \sum_{\tau=0}^{T} \left(\frac{1}{1+r} \right)^{\tau} y_{\tau} + s_{-1}(1+r)$$

and $p_{\tau} = \left(\frac{1}{1+r} \right)^{\tau}$. (We do have savings at T left, but a consumer without bequest motive will choose not to leave any unspent wealth on his last day: $s_T = 0$). Conversely, we could write the first "static" consumer problem in a dynamic fashion, creating a wealth variable that keeps track of how much the consumer has left in his wallet when going from one shop to the other.

The bottom line is that **static** or **dynamic** optimization refers less to an intrinsic nature of a problem, than to a way to look at it. With alternative ways to look at the problem come alternative ways to solve it. The methods we have covered so far belong to **static optimization**. **Dynamic optimization** gathers methods that see a problem as a **sequence** of simpler problems, just as in the initial version of the consumption/savings problem. As the two examples make clear, some problems lend themselves easily to a dynamic interpretation. For such problem, a dynamic method is more intuitive and insightful. This is especially true for problems that involve time, such as the consumption/savings problems. To be sure, this is where the term "dynamic" comes from.

6.2 Finite v. infinite of variables (finite v. infinite horizon)

Quite apart from whether we look at a problem as a static of dynamic one, optimization problems can have a **finite or infinite number of variables** over which to maximize. All the results we have covered so far apply to $\mathbb{R}^n$, that is a finite number of variables; new results need to be derived for infinite number of variables. When we look at a problem as a dynamic one, it usually has a temporal interpretation, so that the distinction translates into whether the problem has a **finite or infinite horizon**. Although dynamic programming can also be applied to finite-horizon problems, it is most popular for infinite horizon problems. The next subsection considers an example of infinite-horizon dynamic programming.

6.3 Stationary Discounted Dynamic Programming: an Example

6.3.1 Bellman Principle of Optimality

Instead of the consumer/savings problem, we consider the following infinite-horizon problem as an illustration. The treatment is a little informal; see e.g. Stokey and Lucas for rigorous derivations of the results and a precise statement of the assumptions that are required. An agent has access to a production function f that takes capital k_{t-1} has a single input to produce an output $y_t = f(k_{t-1})$. The agent lives in autarky, not able to save nor borrow on financial markets. However, he can transfer resources across time by taking part of today's output y_t to serve as capital k_t in tomorrow's production: the agent divides the output y_t between his consumption c_t, and k_{t+1}. For a given k_{-1}, the agent solves:

$$\max_{(c_t)_{t=0}^{\infty},(k_t)_{t=0}^{\infty}} U(c_0, ..., c_t, ...)$$
$$\text{s.t. } c_t + k_t = f(k_{t-1}) \text{ for all } t \geq 0$$

(We will not discuss assumptions that guarantee the existence of a maximum). First, we make restricting assumptions of the form of the utility function. We assume U is additively separable and more precisely takes the form:

$$U(c_0, ..., c_t, ...) = \sum_{t=0}^{\infty} \beta^t u(c_t)$$

where $\beta \in (0, 1)$ is the **discount factor** that accounts for the impatience of consumers.

The constraints have the same sequential form as the flow budget constraint in the previous example: it presents the decision of allocating consumption between all periods as a series of simple decisions between consumption and investment. But the objective function is not in such a sequential form: it depends on the consumption at all periods, not on c_t and k_t only. The idea of dynamic programming is to rewrite the problem so that the objective depends only on these. In order to do so, consider the value function $V_t(k_{t-1})$ that gives the maximum of the program starting at t for an initial stock of capital k_{t-1} inherited from period $t-1$.

$$V_t(k_{t-1}) = \max_{(c_\tau)_{\tau=t}^{\infty},(k_\tau)_{\tau=t}^{\infty}} \sum_{\tau=t}^{\infty} \beta^{\tau-t} u(c_\tau)$$
$$\text{s.t. } c_\tau + k_\tau = f(k_{\tau-1}) \text{ for all } \tau \geq t$$

Now, a bit informally, note that:

$$\begin{aligned}
V_t(k_{t-1}) &= \max_{\substack{(k_\tau)_{\tau=t}^\infty \\ (c_\tau)_{\tau=t}^\infty \\ c_\tau + k_\tau = f(k_{\tau-1}) \text{ for all } \tau \geq t}} \left\{ u(c_t) + \beta \left(\sum_{\tau=t+1}^{\infty} \beta^{\tau-(t+1)} u(c_\tau) \right) \right\} \\
&= \max_{\substack{c_t, k_t \\ c_t + k_t = f(k_{t-1})}} \left\{ u(c_t) + \beta \max_{\substack{(k_\tau)_{\tau=t+1}^\infty \\ (c_\tau)_{\tau=t+1}^\infty \\ c_\tau + k_\tau = f(k_{\tau-1}) \text{ for all } \tau \geq t+1}} \left\{ \left(\sum_{\tau=t+1}^{\infty} \beta^{\tau-(t+1)} u(c_\tau) \right) \right\} \right\} \\
V_t(k_{t-1}) &= \max_{\substack{c_t, k_t \\ c_t + k_t = f(k_{t-1})}} \{ u(c_t) + \beta V_{t+1}(k_t) \}
\end{aligned}$$

The final line is called the **Bellman equation**. It expresses the simple insight at the core of dynamic programming: the **Bellman principle of optimality**. In words, it says that to pick the optimal whole path of consumption and capital from today (period t) on, it is enough to optimally allocate between c_t and k_t today, provided one picks the optimal whole path of consumption and capital from tomorrow (period $(t+1)$) on. To do so however, one needs to know the value function $V_{t+1}(k_t)$, hence to have solved future programs.

6.3.2 Functional equation

Because of the infinite-horizon and the discounted form of the utility in this example, at each t the consumer faces exactly the same program. It follows that we can drop the time index in the Bellman equation. Noting $c = c_t$, $k = k_{t-1}$ and $k' = k_t$:

$$V(k) = \max_{\substack{c, k' \\ c + k' = f(k)}} \{ u(c) + \beta V(k') \}$$

Replacing the constraint in the objective (but we could just as well use a Lagrangian):

$$V(k) = \max_{k'} \{ u(f(k) - k') + \beta V(k') \}$$

Both the left-hand and right-hand sides are functions of k: written this way, the value function appears as the solution to a **functional equation**. More precisely, it is the fixed point to a function of functions T:

$$T : V \mapsto T(V) : \left[k \mapsto \max_{k'} \{ u(f(k) - k') + \beta V(k') \} \right]$$

To prove the existence and uniqueness of V, as well as to solve for V numerically, we rely on the contraction mapping theorem. To do so, we need to specify the metric space on which we define T. In principle, we would like to consider the vector space of functions from $\mathbb{R}$ to $\mathbb{R}$, or continuous functions from $\mathbb{R}$ to $\mathbb{R}$, endowed with some norm.

- However, to define the norm $||V|| = \sup_{k\in\mathbb{R}} |V(k)|$, we need to restrict to the vector subspace of bounded functions $\mathcal{B}(\mathbb{R})$, or bounded and continuous functions $\mathcal{BC}(\mathbb{R})$.
- It turns out that for any set X, the vector space of bounded functions from X to $\mathbb{R}$ is complete, as well as the vector space of bounded and continuous functions from X to $\mathbb{R}$.
- However, we need to make sure that T maps $\mathcal{B}(\mathbb{R})$ into itself, or $\mathcal{BC}(\mathbb{R})$ into itself, which requires some assumptions. We do not explicit them here.

Now, to prove that T is a contraction, we rely on Blackwell's sufficient conditions.

Proposition 6.1. ***Blackwell's Sufficient Conditions for a Contraction***
Let $(X, ||.||)$ *be a normed vector space and* $(B(X), ||.||_\infty)$ *be the metric space of bounded functions from* X *to* $\mathbb{R}$ *endowed with the sup norm* $||f||_\infty = \sup_{x\in X} ||f(x)||$. *Let* $T : B(X) \to B(X)$ *be a mapping of functions. If* T *satisfies:*

1. ***Monotonicity****: for all* $f, g \in B(X)$, $f \leq g \Rightarrow T(f) \leq T(g)$.
2. ***Discounting****: there exists* $\beta \in (0, 1)$ *such that for all* $f \in B(X)$ *and all* $a \geq 0$*:*

$$T(f + a) \leq T(f) + \beta a$$

then T *is a contraction of modulus* β.

In the proposition, we note $f \leq g$ when $f(x) \leq g(x)$ for all $x \in X$, and note a the function constant to a.

Proof. The proof is straightforward once we write $f - g \leq ||f - g||_\infty$ as $f \leq g + ||f, g||_\infty$.

$$\begin{aligned} f \leq g + ||f, g||_\infty &\Rightarrow T(f) \leq T(g + ||f - g||_\infty) \text{ by monotonicity} \\ &\Rightarrow T(f) \leq T(g) + \beta||f - g||_\infty \text{ by discounting} \end{aligned}$$

Similarly, $g - f \leq ||f - g||_\infty$ implies $T(f) \leq T(g) + \beta||f - g||_\infty$, so that $||T(f) - T(g)||_\infty \leq \beta||f - g||_\infty$. QED. □

Let us check our example satisfies Blackwell's sufficient conditions. First, monotonicity: assume $V(k) \leq W(k)$ for all k. Then:

$$\begin{aligned}\forall k, k', u(f(k) - k') + \beta V(k') &\leq u(f(k) - k') + \beta W(k') \\ &\leq \max_{k'}\{u(f(k) - k') + \beta W(k')\} = T(W(k))\end{aligned}$$

Hence:

$$\forall k, T(V(k)) = \max_{k'}\{u(f(k) - k') + \beta V(k')\} \leq T(W(k))$$

Second, discounting:

$$\begin{aligned}\forall k, k', u(f(k) - k') + \beta(V(k') + a) &= u(f(k) - k') + \beta V(k') + \beta a \\ &\leq \max_{k'}\{u(f(k) - k') + \beta V(k')\} + \beta a = T(V(k)) + \beta a\end{aligned}$$

Hence:

$$\forall k, T((V + a)(k)) = \max_{k'}\{u(f(k) - k') + \beta(V(k') + a)\} \leq T(V(k)) + \beta a$$