

Lecture 6: General Optimization

Solving an optimization problem requires three separate steps. First, we show that a solution exists. Second, we find the points that satisfy the necessary conditions for a solution. Third, we determine whether any of those points is a global solution. Setting a derivative equal to zero addresses only the second step.

1 The problem and existence

1.1 The problem and the candidate list

Many optimization problems belong to a family indexed by a parameter. We hold the parameter fixed while solving a single problem, but keep it in the notation so that the solution set and value can later be compared across problems.

Definition 1.1 (Optimization problem). Let Θ be a parameter set, let $U \subseteq \mathbb{R}^n$, and let $f : U \times \Theta \rightarrow \mathbb{R}$. For a fixed $\theta \in \Theta$, the problem

$$\max_{x \in D(\theta)} f(x; \theta), \quad D(\theta) \subseteq U,$$

has the following parts.

- The vector x is the *choice variable* and θ is the *parameter*.
- The set $D(\theta)$ is the *feasible set* and $f(\cdot; \theta)$ is the *objective*.
- The *solution set* is

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

and when it is nonempty its elements all share one objective value, the *value* $V(\theta) = f(x^*; \theta)$ for any $x^* \in X^*(\theta)$.

Example 1.2 (Utility maximization). A consumer with utility $u : \mathbb{R}_+^n \rightarrow \mathbb{R}$ faces prices $p \in \mathbb{R}_{++}^n$ and wealth $w > 0$:

$$\max_{x \in B(p, w)} u(x), \quad B(p, w) = \{x \in \mathbb{R}_+^n : p \cdot x \leq w\}.$$

Here the bundle x is the choice variable, (p, w) is the parameter, the budget set is the feasible set, demand is the solution set, and indirect utility is the value.

Note that writing $\max$ instead of $\sup$ presumes that the supremum is attained. A supremum can be finite without being attained, and then $X^*(\theta)$ is empty and $V(\theta)$ is undefined. Before using the notation, we must establish that a solution exists.

Definition 1.3 (Local and global maximizers). Let $D \subseteq \mathbb{R}^n$ and $f : D \rightarrow \mathbb{R}$, and let $x^* \in D$.

- x^* is a *global maximizer* if $f(x^*) \geq f(x)$ for every $x \in D$.
- x^* is a *local maximizer* if there is $\varepsilon > 0$ with

$$f(x^*) \geq f(x) \quad \text{for every } x \in D \cap B(x^*, \varepsilon),$$

and a *strict local maximizer* if that inequality is strict for every such $x \neq x^*$.

- x^* is *interior* if $x^* \in \text{int } D$, meaning that $B(x^*, \varepsilon) \subseteq D$ for some $\varepsilon > 0$, and a *boundary point* if every ball $B(x^*, \varepsilon)$ meets both D and its complement.

Minimizers are defined by reversing every inequality.

Every global maximizer is a local maximizer, since $D \cap B(x^*, \varepsilon) \subseteq D$; the converse holds under concavity, as we show in Section 3. The interior–boundary distinction determines which first-order condition applies: at an interior point every small perturbation is feasible, so the objective can be compared with x^* in both directions along every line, while at a boundary point some directions leave the feasible set and the comparison is one-sided.

Recall the existence result of Lecture 2, which we use without proof.

Theorem 1.4 (Weierstrass). *If $D \subseteq \mathbb{R}^n$ is nonempty and compact and $f : D \rightarrow \mathbb{R}$ is continuous, then f attains both a maximum and a minimum on D .*

In the next example, we first establish existence and then generate candidates.

1.2 Existence on an unbounded set

The following condition is sufficient for existence on an unbounded feasible set.

Definition 1.5 (Coercivity for maximization). Let $D \subseteq \mathbb{R}^n$ be unbounded. A function $f : D \rightarrow \mathbb{R}$ is *coercive for maximization* on D if

$$x_k \in D, \quad \|x_k\| \rightarrow \infty \quad \implies \quad f(x_k) \rightarrow -\infty$$

for every sequence (x_k) ; equivalently, if for every $M \in \mathbb{R}$ there is $R > 0$ such that $f(x) < M$ whenever $x \in D$ and $\|x\| > R$.

Theorem 1.6 (Coercive existence). *Let $D \subseteq \mathbb{R}^n$ be nonempty, closed, and unbounded. If $f : D \rightarrow \mathbb{R}$ is continuous and coercive for maximization, then f attains a maximum on D .*

Proof (optional). Choose any $x_0 \in D$. Applying Definition 1.5 with $M = f(x_0)$ gives $R > 0$, which we may take larger than $\|x_0\|$, such that

$$x \in D, \quad \|x\| > R \quad \implies \quad f(x) < f(x_0).$$

Put $K = D \cap \{x : \|x\| \leq R\}$. It contains x_0 , so it is nonempty; it is bounded by R ; and it is closed, being the intersection of two closed sets. By the Heine–Borel theorem it is therefore compact, so Theorem 1.4 supplies a maximizer x^* of f on K . Any $x \in D \setminus K$ has $\|x\| > R$ and hence $f(x) < f(x_0) \leq f(x^*)$, the second inequality because $x_0 \in K$. So $f(x) \leq f(x^*)$ for every $x \in D$, which is Definition 1.3. $\square$

$$\forall M \in \mathbb{R} \exists R > 0 : |x| > R \implies f(x) < M$$




Figure 1.1: Coercivity reduces maximization on an unbounded closed feasible set to maximization on a compact subset that contains every candidate with value at least $f(x_0)$.

Notice the structure of the proof: coercivity allows us to replace the unbounded feasible set with a compact one that contains a maximizer, and then apply the Weierstrass theorem. Like compactness, coercivity is sufficient for existence but not necessary.

Example 1.7 (Coercivity is not necessary). On $D = [0, \infty)$, the function $f(x) = xe^{-x}$ is not coercive for maximization, since $f(x) \rightarrow 0$ rather than $-\infty$; yet $f'(x) = (1 - x)e^{-x}$ is positive before 1 and negative after 1, so $x = 1$ is the unique global maximizer. On $D = \mathbb{R}$, the function $f(x) = \arctan x$ is likewise not coercive, but here there is no maximizer: f increases toward the supremum $\pi/2$, which it never attains.

2 Necessary local conditions

2.1 The interior first-order condition

Both first-order conditions below are proved the same way: we restrict f to a line through the candidate point and apply one-variable arguments to the restriction. We collect the one-variable facts first.

Lemma 2.1 (One-variable optimality). *Let $\delta > 0$ and let $g : (-\delta, \delta) \rightarrow \mathbb{R}$ have a local maximum at 0.*

1. *If g is differentiable at 0, then $g'(0) = 0$.*
2. *If g is twice continuously differentiable, then $g''(0) \leq 0$.*

Proof. For part (1), take t small enough that $g(t) \leq g(0)$. When $t > 0$, dividing $g(t) - g(0) \leq 0$ by t preserves the inequality, so

$$\frac{g(t) - g(0)}{t} \leq 0 \quad \text{and letting } t \downarrow 0, \quad g'(0) \leq 0.$$

When $t < 0$, dividing by t reverses it, so the same quotient is nonnegative and letting $t \uparrow 0$ gives $g'(0) \geq 0$. The derivative exists, so both one-sided limits equal it, and $g'(0) = 0$.

For part (2), part (1) gives $g'(0) = 0$. The one-dimensional case of the second-order Taylor expansion from Lecture 5 gives

$$g(t) - g(0) = \frac{1}{2}g''(0)t^2 + r(t), \quad \frac{|r(t)|}{t^2} \longrightarrow 0.$$

If $g''(0) > 0$, then for every sufficiently small nonzero t the positive quadratic term dominates the remainder, so $g(t) > g(0)$. This contradicts the local maximum. Hence $g''(0) \leq 0$. $\square$

Theorem 2.2 (Interior first-order condition). *Let $U \subseteq \mathbb{R}^n$ be open, let $D \subseteq U$, and let $f : U \rightarrow \mathbb{R}$ be differentiable at $x^* \in \text{int } D$. If x^* is a local maximizer or a local minimizer of f on D , then $\nabla f(x^*) = 0$.*

Proof. Fix $v \in \mathbb{R}^n$ and put $g(t) = f(x^* + tv)$. Since $x^* \in \text{int } D$, there is $\varepsilon > 0$ with $B(x^*, \varepsilon) \subseteq D$, so g is defined for $|t| < \varepsilon/\|v\|$ when $v \neq 0$. If x^* is a local maximizer of f on D , then $g(t) \leq g(0)$ for all small $|t|$ by Definition 1.3, so 0 is a local maximum of g . By part (1) of Lemma 2.1 and the chain rule for a line restriction,

$$0 = g'(0) = \nabla f(x^*) \cdot v.$$

This holds for every $v \in \mathbb{R}^n$; taking $v = \nabla f(x^*)$ gives $\|\nabla f(x^*)\|^2 = 0$, so $\nabla f(x^*) = 0$. For a local minimizer, apply the same argument to $-f$. $\square$

Definition 2.3 (Critical point). A point x at which f is differentiable and $\nabla f(x) = 0$ is a *critical point*, or *stationary point*, of f .

Theorem 2.2 is a necessary condition, not a sufficient one. Its converse fails even at a differentiable interior point. Moreover, the condition itself fails if we drop interiority, and it does not apply without differentiability.

Example 2.4 (Three counterexamples). 1. Criticality does not imply optimality. The origin is critical for $f(x) = x^3$ on $\mathbb{R}$ by Definition 2.3, yet $f(t) > 0 > f(-t)$ for every $t > 0$, so it is neither a local maximizer nor a local minimizer.

2. Optimality at a non-interior point does not imply criticality. For $f(x) = x$ on $D = [0, 1]$, the maximizer is $x^* = 1$, which is a boundary point of D , and $f'(1) = 1 \neq 0$.

3. The condition does not apply without differentiability. The function $f(x) = -|x|$ on $\mathbb{R}$ has a strict global maximizer at 0, an interior point, but is not differentiable there, so Theorem 2.2 does not apply.

2.2 Second-order conditions

A critical point can be a local maximizer, a local minimizer, or neither. The Hessian classifies it whenever the associated quadratic form is definite. The tool is the second-order Taylor expansion of Lecture 5: if f is C^2 on an open set containing x^* , then

$$f(x^* + h) - f(x^*) = \nabla f(x^*) \cdot h + \frac{1}{2}h^\top H_f(x^*)h + r(h), \quad \frac{|r(h)|}{\|h\|^2} \rightarrow 0, \quad (1)$$

and the second derivative of a line restriction $g(t) = f(x^* + tv)$ is the quadratic form $g''(t) = v^\top H_f(x^* + tv)v$.

Theorem 2.5 (Second-order conditions). *Let f be C^2 on an open set containing x^* .*

1. *If x^* is a local maximizer, then $H_f(x^*)$ is negative semidefinite.*
2. *If $\nabla f(x^*) = 0$ and $H_f(x^*)$ is negative definite, then x^* is a strict local maximizer.*
3. *If $\nabla f(x^*) = 0$ and $H_f(x^*)$ is indefinite, then x^* is a saddle point, which is neither a local maximizer nor a local minimizer.*

For local minimizers, reverse the signs in the first two statements.

Proof (optional). (1) *Necessity.* Let $v \in \mathbb{R}^n$ and put $g(t) = f(x^* + tv)$, which has a local maximum at 0. Part (2) of Lemma 2.1 and the line-restriction formula give $v^\top H_f(x^*)v = g''(0) \leq 0$. Since v was arbitrary, $H_f(x^*)$ is negative semidefinite.

(2) *Sufficiency.* The function $u \mapsto u^\top H_f(x^*)u$ is continuous and the unit sphere $\{u : \|u\| = 1\}$ is nonempty and compact, so Theorem 1.4 gives a maximum value $-c$ on it, and $c > 0$ because negative definiteness makes every such value negative. For $h \neq 0$, applying this to $u = h/\|h\|$ and multiplying by $\|h\|^2$ gives

$$h^\top H_f(x^*)h \leq -c\|h\|^2.$$

Since $|r(h)|/\|h\|^2 \rightarrow 0$, there is $\rho > 0$ with $|r(h)| \leq (c/4)\|h\|^2$ whenever $0 < \|h\| < \rho$. For such h , (1) with $\nabla f(x^*) = 0$ gives

$$f(x^* + h) - f(x^*) \leq -\frac{c}{2}\|h\|^2 + \frac{c}{4}\|h\|^2 = -\frac{c}{4}\|h\|^2 < 0,$$

so x^* is a strict local maximizer.

(3) *Indefiniteness.* Choose v and w with $v^\top H_f(x^*)v > 0$ and $w^\top H_f(x^*)w < 0$. Because the gradient vanishes, substituting $h = tv$ into (1) gives

$$f(x^* + tv) - f(x^*) = \frac{t^2}{2}v^\top H_f(x^*)v + r(tv) > 0$$

for every sufficiently small nonzero t : after division by t^2 , the remainder tends to zero while the quadratic coefficient is positive. The same argument with $h = tw$ makes the difference negative for every sufficiently small nonzero t . Thus every ball around x^* contains points with larger and smaller values, so x^* is neither a local maximizer nor a local minimizer.

The minimum statements follow by applying the above to $-f$, whose Hessian is $-H_f$. $\square$

When the Hessian is semidefinite but not definite, neither part (2) nor part (3) applies, and no conclusion follows. The next example shows that every outcome is possible in that case.

Example 2.6 (A semidefinite Hessian is inconclusive). Each of $-x^4$, x^4 , and x^3 has $f'(0) = f''(0) = 0$, so their first- and second-order data at the origin are identical. Yet the origin is a strict maximizer of the first, a strict minimizer of the second, and neither for the third. A zero Hessian therefore supplies no classification.

In practice, we check definiteness of the Hessian with the tests of Lecture 5 instead of the definition.

Example 2.7 (A saddle point classified). Let $f(x_1, x_2) = x_1^2 - x_2^2$. Then $\nabla f(0) = 0$ and $H_f = \text{diag}(2, -2)$, whose eigenvalues 2 and -2 have opposite signs, so the Hessian is indefinite and part (3) of Theorem 2.5 makes the origin neither a local maximizer nor a local minimizer.




Figure 2.1: The graph of $f(x_1, x_2) = x_1^2 - x_2^2$ near its saddle point at the origin.

2.3 Boundary first-order conditions

At a boundary maximizer the gradient need not vanish, as Example 2.4 showed. A one-sided statement remains true, and to state it we first name the directions that stay inside the feasible set.

Definition 2.8 (Feasible direction). Let $x \in D \subseteq \mathbb{R}^n$. A vector $v \in \mathbb{R}^n$ is a *feasible direction* from x if there is $\delta > 0$ with

$$x + tv \in D \quad \text{for every } t \in (0, \delta).$$

If D is convex, then $y - x$ is a feasible direction from x for every $y \in D$, since Definition 2.8 is satisfied with $\delta = 1$.

Theorem 2.9 (Boundary first-order condition). Let $U \subseteq \mathbb{R}^n$ be open, let $D \subseteq U$, and let $f : U \rightarrow \mathbb{R}$ be differentiable. If $x^* \in D$ is a local maximizer of f on D , then

$$\nabla f(x^*) \cdot v \leq 0 \quad \text{for every feasible direction } v \text{ from } x^*.$$

In particular, if D is convex, then

$$\nabla f(x^*) \cdot (y - x^*) \leq 0 \quad \text{for every } y \in D. \quad (2)$$

Proof (optional). Fix a feasible direction v and put $g(t) = f(x^* + tv)$. By Definition 2.8, this restriction is feasible for all small $t > 0$. Since x^* is a local maximizer, $g(t) \leq g(0)$ for all sufficiently small $t > 0$, so the difference quotients satisfy

$$\frac{g(t) - g(0)}{t} \leq 0 \quad \text{for all small } t > 0.$$

Letting $t \downarrow 0$ and evaluating the derivative of the line restriction by the chain rule,

$$0 \geq g'(0) = \nabla f(x^*) \cdot v.$$

If D is convex and $y \in D$, then $v = y - x^*$ is feasible, which gives (2). $\square$

At an interior point the condition collapses back to Theorem 2.2: for any v , both $y = x^* + \varepsilon v$ and $y = x^* - \varepsilon v$ lie in D for small $\varepsilon > 0$, so (2) gives $\nabla f(x^*) \cdot v \leq 0$ and $-\nabla f(x^*) \cdot v \leq 0$, forcing $\nabla f(x^*) \cdot v = 0$ for every v .

If x^* is a boundary local maximizer, then the gradient is allowed to not be zero, and if it's not, then it must point out of the feasible set. To see this, note that $\nabla f(x^*) \cdot v \leq 0$ means that the angle between the gradient and any feasible direction v is at least ninety degrees. That is equivalent to saying that $f(x^* + tv) \leq f(x^*)$, meaning that moving in the direction v decreases the value of f relative to $f(x^*)$. That is, at a boundary point, the function is allowed to increase in some directions (those that point out of the feasible set), but it cannot increase in any direction that points into the feasible set.




Figure 2.2: Left: the directions in which f increases and decreases relative to $f(x^*)$. f cannot rise if the angle between the gradient and a feasible direction is at least ninety degrees (the red area). Right: the first-order condition (2) at a local maximizer x^* at the boundary of D . The gradient points out of D and makes an angle of at least ninety degrees with every feasible direction v . Moving **in** a feasible direction therefore cannot increase f .

On a nonnegative orthant the condition becomes one sign restriction and one complementarity relation per coordinate.

Theorem 2.10 (Componentwise form on the orthant). *Let $U \subseteq \mathbb{R}^n$ be open with $\mathbb{R}_+^n \subseteq U$, let $f : U \rightarrow \mathbb{R}$ be differentiable, and let $x^* \in \mathbb{R}_+^n$. Then (2) with $D = \mathbb{R}_+^n$ holds if and only if*

$$\frac{\partial f}{\partial x_i}(x^*) \leq 0, \quad x_i^* \geq 0, \quad x_i^* \frac{\partial f}{\partial x_i}(x^*) = 0 \quad (i = 1, \dots, n). \quad (3)$$

Proof (optional). Necessity. Taking $y = x^* + e_i \in D$ in (2) gives $\partial f / \partial x_i(x^*) \leq 0$. If in addition $x_i^* > 0$, then $y = x^* - x_i^* e_i$ is also in D , and (2) applied to it gives $-x_i^* \partial f / \partial x_i(x^*) \leq 0$, hence $\partial f / \partial x_i(x^*) \geq 0$ and therefore $= 0$. So the product $x_i^* \partial f / \partial x_i(x^*)$ vanishes whether x_i^* is zero or positive.

Sufficiency. Let $y \in D$, so $y_i \geq 0$ for every i . Expanding the inner product and using (3) twice,

$$\nabla f(x^*) \cdot (y - x^*) = \sum_{i=1}^n \frac{\partial f}{\partial x_i}(x^*) y_i - \sum_{i=1}^n \frac{\partial f}{\partial x_i}(x^*) x_i^* = \sum_{i=1}^n \frac{\partial f}{\partial x_i}(x^*) y_i \leq 0,$$

the second sum vanishing term by term and each remaining term being a product of a nonpositive number with a nonnegative one. $\square$

The three conditions in (3) say that increasing any one coordinate cannot raise f to first order, that feasibility holds, and that a partial derivative can be strictly negative only where the corresponding coordinate is already zero and cannot be reduced. The last identity is a complementarity relation; in the parameterized example of Section 3.2, it decides between an interior point and a corner.

3 Global certification

3.1 Concavity, globality, and uniqueness

Everything so far has been about local maximizers, and a local maximizer can fail to be a global one. However, if a function is concave, then every local maximizer is also a global maximizer.

Theorem 3.1 (Concave maximization). *Let $D \subseteq \mathbb{R}^n$ be convex and let $f : D \rightarrow \mathbb{R}$ be concave.*

1. *Every local maximizer of f on D is a global maximizer.*
2. *The set $X^* = \arg \max_{x \in D} f(x)$ is either empty or convex.*
3. *If f is strictly concave, then X^* contains at most one point.*

Proof (optional). Local implies global. Suppose x^* is a local maximizer that is not global, so some $y \in D$ has $f(y) > f(x^*)$. By Definition 1.3 there is $\varepsilon > 0$ with $f(x^*) \geq f(x)$ for every feasible $x \in B(x^*, \varepsilon)$. The points $x_t = (1 - t)x^* + ty$ are feasible by convexity of D , and $\|x_t - x^*\| = t\|y - x^*\|$, so $x_t \in B(x^*, \varepsilon)$ once t is small enough. For such $t \in (0, 1)$, concavity gives

$$f(x_t) \geq (1 - t)f(x^*) + tf(y) > (1 - t)f(x^*) + tf(x^*) = f(x^*),$$

contradicting local maximality.

Convexity of X^ .* Suppose X^* is nonempty and let $x_0, x_1 \in X^*$ have common maximum value M . For $t \in [0, 1]$, the point $(1 - t)x_0 + tx_1$ is feasible, and concavity gives

$$f((1 - t)x_0 + tx_1) \geq (1 - t)M + tM = M.$$

Since M is the maximum, the value is M and the point lies in X^* . So X^* contains the segment joining any two of its points, and is therefore convex.

Strict-concavity uniqueness. If distinct $x, y \in X^*$ had common value M , their midpoint would be feasible and strict concavity would give

$$f\left(\frac{x + y}{2}\right) > \frac{1}{2}f(x) + \frac{1}{2}f(y) = M,$$

contradicting maximality. □

Under concavity, the necessary condition of Theorem 2.9 is also sufficient, so a single first-order check certifies a global maximizer.

Theorem 3.2 (First-order characterization). *Let $D \subseteq \mathbb{R}^n$ be convex and let f be differentiable and concave on an open convex set containing D . A point $x^* \in D$ is a global maximizer of f on D if and only if*

$$\nabla f(x^*) \cdot (y - x^*) \leq 0 \quad \text{for every } y \in D.$$

Proof. Necessity. A global maximizer is a local maximizer, so Theorem 2.9 applies.

Sufficiency. For $y \in D$, the tangent inequality from Lecture 5 at x^* and then the hypothesis give

$$f(y) \leq f(x^*) + \nabla f(x^*) \cdot (y - x^*) \leq f(x^*),$$

so x^* is a global maximizer by Definition 1.3. □

Part (3) of Theorem 3.1 bounds the number of maximizers without producing one. The function $f(x) = -e^x$ is strictly concave on $\mathbb{R}$, and its supremum 0 is approached as $x \rightarrow -\infty$ but never attained, so the solution set is empty. Existence remains a separate step. However, if there exists a point satisfying the first-order condition of Theorem 3.2, then existence is guaranteed by Theorem 3.2.

3.2 A complete parameterized problem

We now apply the three steps, in order, to a parameterized problem.

Example 3.3. For $\theta \in \mathbb{R}$, consider

$$\max_{x \in \mathbb{R}_+} f(x; \theta), \quad f(x; \theta) = \theta x - \frac{1}{2}x^2.$$

Find the solution set $X^*(\theta)$ and value $V(\theta)$ for every θ .

Existence. For fixed θ the objective is continuous, and $D = \mathbb{R}_+$ is nonempty, closed, and unbounded. Since $f(x; \theta) \rightarrow -\infty$ as $x \rightarrow \infty$, the objective is coercive for maximization in the sense of Definition 1.5, so Theorem 1.6 gives a maximizer.

Candidates. Here $n = 1$ and

$$\frac{\partial f}{\partial x}(x; \theta) = \theta - x.$$

By Theorem 2.10 the condition (3) reads

$$\theta - x^* \leq 0, \quad x^* \geq 0, \quad x^*(\theta - x^*) = 0.$$

If $x^* > 0$ the last equation forces $x^* = \theta$, which is admissible only when $\theta > 0$; if $x^* = 0$ the first inequality requires $\theta \leq 0$. Thus exactly one point satisfies the condition: $x^* = \max\{\theta, 0\}$.

Globality and uniqueness. Since $\partial^2 f / \partial x^2 = -1 < 0$, the function is strictly concave on the convex set $\mathbb{R}_+$. The first-order condition is therefore sufficient by Theorem 3.2, and uniqueness follows from part (3) of Theorem 3.1. Hence

$$X^*(\theta) = \{\max\{\theta, 0\}\}, \quad V(\theta) = \begin{cases} 0, & \theta \leq 0, \\ \frac{1}{2}\theta^2, & \theta > 0. \end{cases}$$

The two regimes are the two ways (3) can hold: at $\theta \leq 0$ the solution is the corner 0, whereas at $\theta > 0$ it is the interior point θ .

Removing the quadratic term separates the roles of concavity, strict concavity, and coercivity.

Example 3.4 (Removing the curvature). For $\beta \geq 0$, let $f_\beta(x; \theta) = \theta x - (\beta/2)x^2$ on $\mathbb{R}_+$. When $\beta > 0$, the argument above gives the unique maximizer $x_\beta^*(\theta) = \max\{\theta/\beta, 0\}$. When $\beta = 0$, the objective is the affine function θx , and the three regimes are

$$X^*(\theta) = \{0\} \text{ if } \theta < 0, \quad X^*(\theta) = \mathbb{R}_+ \text{ if } \theta = 0, \quad X^*(\theta) = \emptyset \text{ if } \theta > 0.$$

At $\theta = 0$, existence holds but uniqueness fails because the objective is no longer strictly concave. At $\theta > 0$, existence fails because the affine objective is unbounded above.

Example 3.5 (A non-concave problem in $\mathbb{R}^2$). Find the solution set X^* and value V of

$$\max_{x \in \mathbb{R}^2} \{-x_1^4 - x_2^4 + 4x_1^2 + 2x_2^2\}.$$

Example 3.6 (A concave problem in $\mathbb{R}^2$). Find the solution set X^* and value V of

$$\max_{x \in \mathbb{R}^2} \{3x_1 + 3x_2 - x_1^2 - x_1x_2 - x_2^2\}.$$

3.3 A workflow

The steps below summarize the workflow for solving a parameterized maximization problem.

1. Specify the objective, its domain, the feasible set, and the parameter, as in Definition 1.1.
2. Establish existence, by compactness (Theorem 1.4), coercivity (Theorem 1.6), or a direct argument.
3. Generate every candidate: interior critical points (Theorem 2.2), boundary points satisfying the feasible-direction condition of Theorem 2.9, and points where f fails to be differentiable, where no first-order condition applies.
4. Classify interior candidates locally with Theorem 2.5 when the classification is needed and the Hessian is definite or indefinite. Discard the minimizers and saddle points, and keep the local maximizers and the inconclusive points.
5. Certify globality, either by concavity (Theorem 3.1 and Theorem 3.2) or by comparing the values of f across all remaining candidate points. If f is concave, then any critical point in D is a global maximizer, and under strict concavity the unique one.
6. Report $X^*(\theta)$ and $V(\theta)$.