Convex sets and concave functions

Robert J. Carroll

Last revised: 6 May 2026

1 Motivation

Convexity is the hidden structural property under nearly every well-behaved optimization problem in formal political-economy theory, and it is the property that makes the next two handouts' static and dynamic optimization machinery actually work. Two facts together make this so. First, when the feasible set in a maximization problem is *convex* — the set of mixed strategies in a game, the simplex of probability distributions over outcomes, the budget set of a coalition allocating resources subject to a linear constraint, the upper-contour set of a quasi-concave utility above a level — the set is well-behaved enough for the standard optimization arguments to go through. Second, when the objective function is *concave* — a candidate's quadratic-loss vote-share function, a voter's risk-averse expected-utility, a regulator's decreasing-marginal-returns welfare function — every local maximum is a global maximum, and the first-order condition is sufficient (not just necessary) for an optimum. Together, "convex feasible set + concave objective" is the structural setup in which optimization is tractable, and the upcoming static-optimization handout's KKT framework presumes both conditions.

This handout develops the vocabulary. Convex sets and their algebra come first (§2), followed by the separating-hyperplane theorem (§3) — a structural result that underlies existence-of-equilibrium-prices, supporting-hyperplane characterizations of policy preferences, and the duality theory of convex optimization. Concave and convex functions are introduced in §4, with the global-optimum-from-concavity result that motivates the whole cluster. Two differential characterizations of concavity follow in §5: the first-order (gradient-based) one, which says the function lies below its tangent hyperplane at every point, and the second-order (Hessian-based) one, which says the Hessian is negative semidefinite everywhere. The Hessian-definiteness diagnostic is exactly the structural payoff of §5 of the eigenvalue-theory handout (definiteness via eigenvalues, Sylvester's criterion via leading principal minors), now linked to a substantive PE-relevant property of the function. Finally, §6 takes up quasi-concavity — a weaker condition that nonetheless preserves the "every local max is global" structure for many purposes, and that is the operative condition for most utility functions in microeconomic theory and for the spatial-voting models with single-peaked preferences. The bordered-Hessian diagnostic for quasi-concavity is the one the constrained-optimization SOC of the static-optimization handout will lean on.

2 Convex sets

Many of the feasible sets a working political economist computes optima over are convex — not by accident, but because the natural feasibility conditions in formal modeling tend to produce convex sets. The simplex of probability distributions over a finite outcome space (the set of lotteries from the choice-under-risk handout) is convex. The set of mixed strategies in a normal-form game is convex (a product of simplices). The set of policies a coalition can implement subject to a linear budget constraint is convex. The upper-contour set of a quasi-concave utility above a fixed level is convex. The intersection of any number of convex constraints is convex. Convexity of feasible sets is the structural reason these settings admit clean existence and uniqueness theorems for optima.

Definition 1. A set $C \subseteq \mathbb{R}^n$ is *convex* if for every $\boldsymbol{x}, \boldsymbol{y} \in C$ and every $t \in [0, 1]$,

$$t\boldsymbol{x} + (1-t)\boldsymbol{y} \in C.$$

That is, the line segment between any two points of C lies entirely in C.

Example 2 (Standard convex sets in PE applications)**.**

- $\mathbb{R}^n$ itself, any affine subspace, any hyperplane, and any half-space $\{\boldsymbol{x} : \boldsymbol{a}^\top \boldsymbol{x} \le b\}$.
- The closed ball $\{\boldsymbol{x} : \|\boldsymbol{x} - \boldsymbol{x}_0\| \le r\}$ and the open ball.
- The probability simplex $\Delta^{n-1} = \{\boldsymbol{p} \in \mathbb{R}^n_+ : \sum_i p_i = 1\}$ — the home of the lotteries from the choice-under-risk handout.
- Any product of convex sets is convex (so a mixed-strategy profile in a game lives in a product of simplices, which is convex).
- The intersection of any (possibly infinite) family of convex sets is convex.

The intersection-closure property is what makes convex constraints compose: a coalition that must satisfy a budget constraint, a non-negativity constraint, and a minimum-allocation constraint faces a feasible set that is the intersection of three convex sets, hence convex. The structural reason most of optimization theory is set up under convex-feasible-set hypotheses is that this kind of compositional structure is generic.

Definition 3. The *convex hull* of a set $S \subseteq \mathbb{R}^n$, written $\operatorname{conv}(S)$, is the smallest convex set containing S. Equivalently, $\operatorname{conv}(S)$ is the set of all finite convex combinations of points in S:

$$\operatorname{conv}(S) = \left\{ \sum_{i=1}^{k} t_i \boldsymbol{x}_i : k \ge 1,\ \boldsymbol{x}_i \in S,\ t_i \ge 0,\ \sum_i t_i = 1 \right\}.$$

The convex hull is the natural "filling-in" of a discrete set. The convex hull of the vertices of the unit cube is the cube. The convex hull of $n+1$ generic points in $\mathbb{R}^n$ is an n-simplex. The convex hull of a finite set of policy alternatives is the set of all probabilistic mixtures of those alternatives — which is exactly the lottery space of the choice-under-risk handout.

3 The separating hyperplane theorem

A structural property of convex sets that recurs throughout political-economy applications — in equilibrium existence proofs, in supporting-price characterizations, in the duality theory of optimization — is that two disjoint convex sets can be separated by a hyperplane. The geometric content is intuitive in low dimensions: two convex blobs that don't touch can be split by a line (in $\mathbb{R}^2$) or a plane (in $\mathbb{R}^3$). The structural content in higher dimensions is the same statement, with the line / plane generalized to a hyperplane.

Theorem 4 (Separating hyperplane)**.** *Let $C_1, C_2 \subseteq \mathbb{R}^n$ be non-empty disjoint convex sets, with at least one of them open. Then there exists $\boldsymbol{a} \in \mathbb{R}^n \setminus \{\boldsymbol{0}\}$ and $b \in \mathbb{R}$ such that*

$$\boldsymbol{a}^\top \boldsymbol{x} \le b \text{ for all } \boldsymbol{x} \in C_1, \qquad \boldsymbol{a}^\top \boldsymbol{y} \ge b \text{ for all } \boldsymbol{y} \in C_2.$$

*If C_1, C_2 are both closed and at least one is bounded, the inequalities can be made strict (*strict separation*).*

The proof uses the closest-point characterization of the projection onto a closed convex set: pick the closest point of C_1 to C_2 (which exists by compactness when one is bounded), and the perpendicular bisector of the segment connecting that point to its counterpart on C_2 is the separating hyperplane. Boyd and Vandenberghe (2004, Ch. 2) works through the proof carefully; Rockafellar (1970) treats the technical edge cases (when the sets are not bounded and the relevant separation is non-strict).

Example 5 (Supporting hyperplane to a convex set)**.** A direct corollary: if $C \subseteq \mathbb{R}^n$ is a closed convex set and $\boldsymbol{x}_0$ is a boundary point, there exists a non-zero $\boldsymbol{a} \in \mathbb{R}^n$ with $\boldsymbol{a}^\top \boldsymbol{x}_0 = b$ and $\boldsymbol{a}^\top \boldsymbol{x} \le b$ for every $\boldsymbol{x} \in C$. The hyperplane $\{\boldsymbol{x} : \boldsymbol{a}^\top \boldsymbol{x} = b\}$ *supports* C at $\boldsymbol{x}_0$. The political-economy reading: at any boundary point of a convex feasible set, there is a half-space description that contains the whole set — the linear functional $\boldsymbol{a}^\top(\cdot)$ is the "shadow price" of the boundary at that point. This is the structural reason Lagrange multipliers exist for convex constrained-optimization problems, a connection the static-optimization handout will pursue.

4 Concave and convex functions

The right structural property of an objective function for optimization to be tractable is *concavity*: the function curves down (or stays flat) along any line segment, so it has no internal "valleys" that could trap a hill-climbing search at a local maximum that isn't global. Concavity is the formal expression of decreasing marginal returns — the analyst's substantive intuition about voter utility, candidate vote share, or coalition payoff functions in many settings. The structural payoff of a concave-and-convex setup is global tractability: every local maximum is global, and a first-order condition is sufficient (not just necessary) for an optimum.

Definition 6. Let $C \subseteq \mathbb{R}^n$ be convex. A function $f : C \to \mathbb{R}$ is *concave* if for every $\boldsymbol{x}, \boldsymbol{y} \in C$ and every $t \in [0, 1]$,

$$f(t\boldsymbol{x} + (1-t)\boldsymbol{y}) \ge t f(\boldsymbol{x}) + (1-t) f(\boldsymbol{y}).$$

f is *strictly concave* if the inequality is strict for $\boldsymbol{x} \ne \boldsymbol{y}$ and $t \in (0, 1)$. f is *convex* if $-f$ is concave; *strictly convex* similarly.

The geometric reading: concavity says the chord between any two points on the graph lies *below* the graph, and the function curves up to meet the chord (or sits flat with it, in the linear case). Convexity is the dual — the chord lies above the graph, and the function curves down to meet the chord.

Example 7 (Standard concave functions in PE)**.**

- Linear functions $\boldsymbol{a}^\top \boldsymbol{x} + b$ are both concave and convex (the chord coincides with the graph).
- Quadratic-loss utility $f(\boldsymbol{x}) = -(\boldsymbol{x} - \boldsymbol{x}_v)^\top A(\boldsymbol{x} - \boldsymbol{x}_v)$ for symmetric positive-definite A is strictly concave on $\mathbb{R}^n$ (the spatial-voting workhorse).

- Logarithmic utility $f(x) = \log x$ on $\mathbb{R}_{++}$ is strictly concave (the canonical risk-averse utility from the choice-under-risk handout).
- Concave utility = risk aversion: a vNM utility $u : \mathbb{R} \to \mathbb{R}$ representing a risk-averse agent is concave (Jensen's inequality, from the random-variables-and-expectations handout).
- Norms $\|\cdot\|$ are convex (and not concave); negative norms are concave.

The substantive payoff is the global-maximum result.

Theorem 8 (Local-implies-global maximum for concave functions)**.** *Let $C \subseteq \mathbb{R}^n$ be convex and let $f : C \to \mathbb{R}$ be concave. Then any local maximum of f on C is also a global maximum. If f is strictly concave, the global maximum is unique.*

Proof. Suppose $\boldsymbol{x}^*$ is a local but not global maximum, so there exists $\boldsymbol{y} \in C$ with $f(\boldsymbol{y}) > f(\boldsymbol{x}^*)$. By concavity, for $t \in (0, 1)$:

$$f((1-t)\boldsymbol{x}^* + t\boldsymbol{y}) \geq (1-t)f(\boldsymbol{x}^*) + tf(\boldsymbol{y}) > (1-t)f(\boldsymbol{x}^*) + tf(\boldsymbol{x}^*) = f(\boldsymbol{x}^*).$$

For t small enough, the convex combination $(1-t)\boldsymbol{x}^* + t\boldsymbol{y}$ is arbitrarily close to $\boldsymbol{x}^*$, so $\boldsymbol{x}^*$ has neighbors with strictly higher value — contradicting local maximality. The strict-concavity uniqueness argument is similar: if $\boldsymbol{x}^*$ and $\boldsymbol{y}^*$ are both global maxima, the midpoint achieves strictly more by strict concavity, a contradiction. □

The theorem is the structural reason concavity matters for optimization. Without it, finding a local maximum (which the FOC delivers) is no guarantee of finding the optimum the modeler actually cares about. With it, the FOC is enough.

5 Differential characterizations of concavity

Checking concavity from the definition — verifying the inequality on every pair of points and every $t \in [0, 1]$ — is intractable for any function more complicated than a textbook example. The two differential characterizations of this section are how concavity actually gets checked in practice. The first-order characterization says the function lies below its tangent hyperplane at every point. The second-order characterization says the Hessian is negative semidefinite everywhere. The second is the most-used in applied work, because the Hessian-definiteness check reduces to the eigenvalue / Sylvester-criterion machinery of the previous handout.

Theorem 9 (First-order characterization)**.** *Let $C \subseteq \mathbb{R}^n$ be open convex and let $f : C \to \mathbb{R}$ be continuously differentiable. Then f is concave iff for every $\boldsymbol{x}, \boldsymbol{y} \in C$,*

$$f(\boldsymbol{y}) \leq f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top (\boldsymbol{y} - \boldsymbol{x}).$$

f is strictly concave iff the inequality is strict for $\boldsymbol{x} \neq \boldsymbol{y}$.

The reading: the linear approximation to f at $\boldsymbol{x}$ — the right-hand side, geometrically the tangent hyperplane to the graph at $(\boldsymbol{x}, f(\boldsymbol{x}))$ — lies above the function everywhere. Equivalently, the tangent hyperplane is a *supporting* hyperplane to the hypograph $\{(\boldsymbol{x}, t) : t \leq f(\boldsymbol{x})\}$. The connection to the separating-hyperplane theorem of §3 is direct: the hypograph of a concave function is a convex set, and the tangent hyperplanes are supporting hyperplanes to that set.[1]

Theorem 10 (Second-order characterization)**.** *Let $C \subseteq \mathbb{R}^n$ be open convex and let $f : C \to \mathbb{R}$ be twice continuously differentiable. Then:*

- *f is concave iff $H_f(\boldsymbol{x})$ is negative semidefinite for every $\boldsymbol{x} \in C$.*
- *f is convex iff $H_f(\boldsymbol{x})$ is positive semidefinite for every $\boldsymbol{x} \in C$.*
- *If $H_f(\boldsymbol{x})$ is negative definite for every $\boldsymbol{x} \in C$, then f is strictly concave.*

Proof sketch. Apply Taylor's theorem to second order at any $\boldsymbol{x} \in C$: $f(\boldsymbol{x}+\boldsymbol{h}) = f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top \boldsymbol{h} + \frac{1}{2}\boldsymbol{h}^\top H_f(\boldsymbol{x})\boldsymbol{h} + o(\|\boldsymbol{h}\|^2)$. The first-order characterization says $f(\boldsymbol{x}+\boldsymbol{h}) \leq f(\boldsymbol{x}) + \nabla f(\boldsymbol{x})^\top \boldsymbol{h}$, which forces the quadratic term $\frac{1}{2}\boldsymbol{h}^\top H_f(\boldsymbol{x})\boldsymbol{h}$ to be non-positive for every $\boldsymbol{h}$ — i.e., $H_f(\boldsymbol{x})$ is negative semidefinite. The converse argument runs through integrating the inequality along line segments. Sundaram (1996, Ch. 7) and Boyd and Vandenberghe (2004, Ch. 3) give complete proofs. □

The note about strict concavity is structurally important: negative-definite Hessian is sufficient for strict concavity but not necessary. The function $f(x) = -x^4$ is strictly concave on $\mathbb{R}$ but has Hessian $-12x^2$, which is zero at $x = 0$ (only negative semidefinite there). Strict concavity is a global condition and the Hessian gives a local one; the local-to-global gap is what produces the asymmetry.

Example 11 (Spatial-voting utility revisited)**.** Continuing the running spatial-voting example: $u_v(\boldsymbol{p}) = -(\boldsymbol{p} - \boldsymbol{x}_v)^\top A(\boldsymbol{p} - \boldsymbol{x}_v)$ for symmetric positive-definite A. The Hessian is $H_{u_v}(\boldsymbol{p}) = -2A$, negative definite by the positive-definiteness of A (its eigenvalues are positive, so $-2A$'s eigenvalues are negative, and the eigenvalue characterization of definiteness from the eigenvalue-theory handout certifies $-2A$ as negative definite). So u_v is strictly concave, by Theorem 10. By Theorem 8, every local max is the unique global max, and the FOC ($\nabla u_v = -2A(\boldsymbol{p} - \boldsymbol{x}_v) = \boldsymbol{0}$, i.e., $\boldsymbol{p} = \boldsymbol{x}_v$) pins it down. The voter's optimum is her ideal point, certified globally.

6 Quasi-concavity

A great many utility functions in microeconomic theory are not concave but *are* quasi-concave — a weaker condition that nonetheless preserves the most useful structural property of concavity (every local maximum on a convex set is global, under weak technical conditions) and is the right one for spatial-voting models with single-peaked preferences. Cobb–Douglas utility $u(x_1, x_2) = x_1^{1/2} x_2^{1/2}$ on $\mathbb{R}^2_{++}$ is quasi-concave but not concave (its Hessian has mixed signs); the upper-contour sets are the convex regions above the indifference hyperbolas. A spatial voter with single-peaked preferences along a single dimension has a quasi-concave utility on $\mathbb{R}$ — and the multidimensional generalization (utility quasi-concave on $\mathbb{R}^k$) is the operative condition for the median-voter and Plott-style results.

[1]The first-order characterization gives an immediate sufficient condition for an unconstrained maximum that the static-optimization handout's FOC theorem will state: if $\nabla f(\boldsymbol{x}^*) = \boldsymbol{0}$ at an interior critical point and f is concave, then $f(\boldsymbol{y}) \leq f(\boldsymbol{x}^*) + \boldsymbol{0}^\top(\boldsymbol{y} - \boldsymbol{x}^*) = f(\boldsymbol{x}^*)$ for every $\boldsymbol{y}$, so $\boldsymbol{x}^*$ is a global maximum. The structural moral: under concavity, the FOC is sufficient as well as necessary, and the analyst's hill-climbing is global. The complementary structural reading runs through the supporting-hyperplane theorem: the tangent hyperplane at a critical point is horizontal (gradient zero), and "the function lies below the tangent hyperplane everywhere" becomes "the function value at every other point is at most the function value at the critical point." Boyd and Vandenberghe (2004, Ch. 3) works through the equivalence in the cleanest way; Rockafellar (1970) is the classical reference for the deeper convex-analytic theorems (Fenchel duality, polar cones, conjugate functions) that this handout does not develop.

Definition 12. Let $C \subseteq \mathbb{R}^n$ be convex. A function $f : C \to \mathbb{R}$ is *quasi-concave* if for every $\alpha \in \mathbb{R}$, the upper-contour set $\{\boldsymbol{x} \in C : f(\boldsymbol{x}) \geq \alpha\}$ is convex. Equivalently, f is quasi-concave iff

$$f(t\boldsymbol{x} + (1-t)\boldsymbol{y}) \geq \min\{f(\boldsymbol{x}), f(\boldsymbol{y})\}$$

for every $\boldsymbol{x}, \boldsymbol{y} \in C$ and every $t \in [0, 1]$.

The equivalence of the two formulations is a standard exercise. The upper-contour-set formulation is geometrically transparent: a function is quasi-concave iff its level sets bound convex regions on the high-value side. The voter's indifference curves ($\{u = \alpha\}$) bound her upper-contour sets ($\{u \geq \alpha\}$), and the voter's preferences are quasi-concave iff every upper-contour set is convex. This is the substantive content of *convex preferences* in microeconomic theory: an agent's preferences are convex iff her utility is quasi-concave.

Proposition 13. *Every concave function is quasi-concave; the converse fails.*

Proof. The upper-contour set $\{f \geq \alpha\}$ of a concave f is convex: if $f(\boldsymbol{x}), f(\boldsymbol{y}) \geq \alpha$, then $f(t\boldsymbol{x} + (1-t)\boldsymbol{y}) \geq tf(\boldsymbol{x}) + (1-t)f(\boldsymbol{y}) \geq \alpha$. The converse fails: $f(x) = x^3$ on $\mathbb{R}$ is quasi-concave (upper-contour sets are half-lines, hence convex) but not concave (Hessian is $6x$, positive for $x > 0$). □

The structural payoff of quasi-concavity for optimization is similar to concavity but slightly weaker: every local maximum is a global maximum, but the FOC is generally not sufficient (additional regularity is needed). The differential diagnostic is the bordered Hessian.

Definition 14. For a twice-continuously-differentiable $f : \mathbb{R}^n \to \mathbb{R}$ at $\boldsymbol{x}$, the *bordered Hessian* is the $(n+1) \times (n+1)$ symmetric matrix

$$B_f(\boldsymbol{x}) = \begin{pmatrix} 0 & \nabla f(\boldsymbol{x})^\top \\ \nabla f(\boldsymbol{x}) & H_f(\boldsymbol{x}) \end{pmatrix}.$$

Theorem 15 (Bordered-Hessian sufficient condition for quasi-concavity)**.** *Let $C \subseteq \mathbb{R}^n_{++}$ be open convex and let $f : C \to \mathbb{R}$ be twice continuously differentiable with $\nabla f(\boldsymbol{x}) > \boldsymbol{0}$ for every $\boldsymbol{x} \in C$. If at every $\boldsymbol{x} \in C$ the leading principal minors of $B_f(\boldsymbol{x})$ alternate in sign with $D_2 < 0$, $D_3 > 0$, $D_4 < 0, \ldots$ — equivalently, $(-1)^k D_{k+1}(\boldsymbol{x}) > 0$ for $k = 1, \ldots, n$ — then f is strictly quasi-concave on C.*

The condition is awkward to state cleanly, and the proof is technical. Sundaram (1996, Ch. 7) treats it carefully, and Mas-Colell, Whinston, and Green (1995, App. M) presents it in the microeconomic-theory context where it is most often invoked. The political-economy take-away: when an applied-modeler colleague writes "the bordered Hessian satisfies the standard sign condition," she is invoking Theorem 15 to certify quasi-concavity and (by the local-implies-global theorem analogous to Theorem 8 for quasi-concavity) the global-tractability of the corresponding constrained-optimization problem. The static-optimization handout's KKT-with-SOC framework will use this directly.

Example 16 (Single-peaked preferences and quasi-concavity)**.** A voter with single-peaked preferences over a one-dimensional policy space $\mathbb{R}$ has utility $u : \mathbb{R} \to \mathbb{R}$ that increases up to her ideal point x_v and decreases past it. Her upper-contour set $\{x : u(x) \geq \alpha\}$ is the (possibly empty) interval $[x_v - r^L_\alpha, x_v + r^R_\alpha]$ for some non-negative half-widths r^L_α, r^R_α depending on the level α — a convex set. So u is quasi-concave on $\mathbb{R}$. The multidimensional generalization to $\mathbb{R}^k$ keeps the structure: utility is quasi-concave iff every upper-contour set is convex. This is exactly the structural condition under which the median-voter theorem and the Plott chaos results are stated, and it is the right structural reading of the modeling assumption "voters have convex preferences."

7 What's next

The next handout in the cluster is static optimization: unconstrained problems (FOC and SOC, with the differential characterizations of concavity from §5 as the diagnostic), constrained problems via Lagrange multipliers (with the supporting-hyperplane theorem of §3 as the structural backbone), inequality-constrained problems via KKT (with quasi-concavity as the operative SOC condition), existence theorems via Weierstrass and Berge, and the envelope theorem as the engine of comparative statics. The dynamic-optimization handout closes the cluster with the Bellman equation and contraction-mapping arguments.

For graduate-level treatments at this handout's level of abstraction, Sundaram (1996, Ch. 7) is the standard PE-flavored reference. Boyd and Vandenberghe (2004, Ch. 2–3) treats convex sets and convex functions with the modern computational perspective and is freely available online. Rockafellar (1970) is the classical and definitive reference for convex analysis, including the technical edge cases this handout has skirted (relative interiors, recession cones, Fenchel duality). Mas-Colell, Whinston, and Green (1995, App. M) treats quasi-concavity and the bordered-Hessian condition in the microeconomic-theory context where these tools are most often invoked.

8 Exercises

Exercise 17. Show that the intersection of any (possibly infinite) family of convex sets is convex. Use this to argue that the feasible set of a coalition with m linear inequality constraints (each defining a half-space) is convex.

Exercise 18. Verify that the probability simplex $\Delta^{n-1} = \{\boldsymbol{p} \in \mathbb{R}^n_+ : \sum_i p_i = 1\}$ is a convex set in $\mathbb{R}^n$. Identify Δ^{n-1} as the convex hull of the standard basis vectors $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n$.

Exercise 19. *Supporting hyperplane in spatial voting.* Let $C = \{\boldsymbol{p} \in \mathbb{R}^k : \|\boldsymbol{p}\| \leq 1\}$ be the unit ball, viewed as the set of feasible policy positions. Pick a boundary point $\boldsymbol{p}_0$ with $\|\boldsymbol{p}_0\| = 1$. Compute the supporting hyperplane to C at $\boldsymbol{p}_0$. (Hint: it is the hyperplane through $\boldsymbol{p}_0$ perpendicular to $\boldsymbol{p}_0$.)

Exercise 20. Use the definition of concavity to verify directly that $f(x) = -x^2$ on $\mathbb{R}$ is strictly concave. Then use the definition to verify that $f(x_1, x_2) = -x_1^2 - x_2^2$ on $\mathbb{R}^2$ is strictly concave.

Exercise 21. Compute the Hessian of $f(x_1, x_2) = -x_1^2 - 4x_1x_2 - 4x_2^2 + 5x_1$ and use Theorem 10 to determine whether f is concave, strictly concave, or neither.

Exercise 22. *Risk aversion as concavity.* Continuing the choice-under-risk handout: a vNM utility $u : \mathbb{R} \to \mathbb{R}$ representing a risk-averse agent satisfies $\mathbb{E}[u(X)] \leq u(\mathbb{E}[X])$ for every random variable X (Jensen's inequality). Argue from the inequality that risk aversion is equivalent to concavity of u. Discuss in one or two sentences how the strict-concavity case relates to strict risk aversion.

Exercise 23. *Voter's concave utility.* For a voter with quadratic-loss utility $u_v(\boldsymbol{p}) = -(\boldsymbol{p} - \boldsymbol{x}_v)^\top A(\boldsymbol{p} - \boldsymbol{x}_v)$ on $\mathbb{R}^k$, with A symmetric positive-definite: (a) verify that u_v is strictly concave on $\mathbb{R}^k$ via the Hessian. (b) Conclude that the voter's ideal point $\boldsymbol{x}_v$ is the unique global maximum of u_v, with the FOC $\nabla u_v = \mathbf{0}$ both necessary and sufficient.

Exercise 24. Show that $f(x_1, x_2) = x_1x_2$ on $\mathbb{R}^2_{++}$ is quasi-concave but not concave. (Hints: for quasi-concavity, the upper-contour sets $\{x_1x_2 \geq \alpha\}$ for $\alpha > 0$ are the regions above the rectangular hyperbola $x_1x_2 = \alpha$; verify these are convex. For non-concavity, compute the Hessian and check it is not negative semidefinite.)

Exercise 25. *Single-peaked preferences are quasi-concave.* A voter with utility $u : \mathbb{R} \to \mathbb{R}$ has *single-peaked* preferences with peak at x_v if u is strictly increasing on $(-\infty, x_v]$ and strictly decreasing on $[x_v, \infty)$. Show that any single-peaked utility is quasi-concave on $\mathbb{R}$. (Hint: identify the upper-contour sets and check they are intervals.)

Exercise 26. *Cobb–Douglas via bordered Hessian.* Let $f(x_1, x_2) = x_1^{1/2} x_2^{1/2}$ on $\mathbb{R}^2_{++}$. (a) Compute ∇f and H_f. (b) Show f is not concave (the Hessian is not negative semidefinite). (c) Verify Theorem 15: compute the leading principal minors D_2, D_3 of the bordered Hessian and check the sign-alternation condition.

References

Boyd, Stephen and Lieven Vandenberghe (2004). *Convex Optimization.* Cambridge: Cambridge University Press.

Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995). *Microeconomic Theory.* New York: Oxford University Press.

Rockafellar, R. Tyrrell (1970). *Convex Analysis.* Princeton: Princeton University Press.

Sundaram, Rangarajan K. (1996). *A First Course in Optimization Theory.* Cambridge: Cambridge University Press.