Common Utility Specifications: Properties, KKT Implications, and Expected Optimal Solutions

FRD

February 11, 2026

Introduction

This note provides a compact roadmap to the utility functions most commonly used in consumer theory. The objective is to understand how the *shape* of a utility function determines the *shape* of the optimal solution and what the Karush–Kuhn–Tucker (KKT) conditions can deliver.

$$u(x) \quad \text{shapes preferences} \quad \Rightarrow \quad \text{shapes the budget–constrained optimum.}$$

We classify utilities using three core properties:

1. Monotonicity. We say that preferences are (weakly) monotone if "more of any good, holding the others fixed, is never worse":

$$x_i' > x_i,\ x_{-i}' = x_{-i} \ \Rightarrow\ u(x') \geq u(x).$$

In this case we expect the consumer to spend the whole budget. Formally, it is enough to assume local non–satiation[1]: whenever $p \gg 0$, $y > 0$ and a solution x^* to the consumer problem exists, local non–satiation implies

$$p \cdot x^* = y.$$

2. Curvature. Throughout, assume u is twice continuously differentiable and let $D^2u(x)$ denote its Hessian.

- *Concavity.* $u(x)$ is concave if the Hessian is negative semidefinite (NSD) at every x. In *one dimension* this reduces to $u''(x) \leq 0$. Economically, each marginal utility $\partial u/\partial x_i$ is weakly decreasing in its own argument; if $D^2u(x)$ is negative *definite* (ND) everywhere (strict

[1] Local non–satiation means that around any bundle x and for any $\varepsilon > 0$ there is some x' with $\|x' - x\| < \varepsilon$ and $u(x') > u(x)$. It is weaker than strict monotonicity but still guarantees that any optimum cannot leave unspent income when all prices are strictly positive.

concavity), marginal utilities are strictly decreasing along any nontrivial direction.[2] With a convex, compact budget set, concavity ensures that any local maximizer is global; strict concavity implies the maximizer (if it exists) is unique (though still possibly a corner when preferences are monotone).

- *Quasi–concavity.* $u(x)$ is quasi–concave if all upper contour sets $\{x : u(x) \geq c\}$ are **convex**. Quasi–concavity is weaker than concavity but enough to guarantee "single-peakedness" along any budget line: the set of maximizers is convex. Strict quasi–concavity[3] then delivers a *unique* optimal bundle for each (p, y) whenever a solution exists. Geometrically, strictly quasi–concave utilities have "nicely curved" indifference curves; in practice they are hard to distinguish by eye from concave utilities. The difference is technical: there exist nonconcave functions that are still quasi–concave, so the maximization problem is still well behaved. Every concave utility is quasi–concave, and any strictly increasing affine transformation $v(x) = a\,u(x) + b$ (with $a > 0$) preserves quasi–concavity. This is why quasi–concavity is the natural minimal shape assumption in consumer theory.

- *Convexity.* $u(x)$ is convex if the Hessian is positive semidefinite (PSD) everywhere. Then marginal utilities are (weakly) increasing, and in a maximization problem over a convex budget set, any global maximizer (if it exists) typically lies at extreme points (corners); interior stationary points given by first–order conditions tend to be minima or saddle points rather than maxima.

3. Implications for optimization. In what follows, we mostly care about the coarse distinction between "effectively concave" and "genuinely non–concave" problems.

- *Concave program (broad sense).* We say the consumer problem is (effectively) concave if the budget set is convex and preferences admit a concave representation: either u itself is concave, or there exists a strictly increasing function ϕ such that $v = \phi \circ u$ is concave. Many standard specifications (Cobb–Douglas, CES with $\rho < 1$, quasi–linear with concave v) fall into this class after a monotone transformation (e.g. taking logs). In this case we can treat the problem as a concave maximization problem: under a mild constraint qualification, KKT conditions are *necessary and sufficient* for global optimality, and strict concavity of the representative yields a unique solution (which may still be a corner if preferences are monotone).

[2] The usual phrase "diminishing marginal utility" refers to $\partial^2 u/\partial x_i^2 < 0$ for each i. Full (strict) concavity is stronger: besides negative diagonal terms, the cross–partials must be such that the whole Hessian is negative (semi)definite.

[3] A utility u is strictly quasi–concave if for any distinct x, y and any $\theta \in (0, 1)$,

$$u(\theta x + (1 - \theta)y) > \min\{u(x), u(y)\}.$$

Equivalently, all upper contour sets $\{x : u(x) \geq c\}$ are **strictly convex** (pay attention to the difference in the degree of convexity of the UCS).

- *Quasi–concave but not (effectively) concave.* If u is only quasi–concave and no monotone concave representative is available, the problem is not formally a concave program, but preferences are still convex and the set of maximizers over a convex budget set is convex; strict quasi–concavity delivers a unique maximizer whenever a solution exists. First–order and KKT conditions remain *necessary* for interior candidates, but are not by themselves sufficient for global optimality; in practice they must be complemented with boundary and corner checks.

- *Non–quasi–concave.* When preferences fail quasi–concavity (e.g. convex or multi–peaked utilities), KKT conditions only characterize stationary points, which can be minima or saddle points. Global optima, if they exist, typically lie at corners or along the boundary and must be identified by explicit comparison of feasible bundles.

1 Cobb–Douglas utilities

Definition and basic properties

A Cobb–Douglas utility has the form

$$u(x) \;=\; \prod_{i=1}^{k} x_i^{\alpha_i}, \qquad \alpha_i > 0.$$

This function is continuous on $\mathbb{R}^k_+$ and strictly increasing in each good x_i as long as $\alpha_i > 0$. Monotonicity is immediate from the partial derivatives:

$$\frac{\partial u}{\partial x_i}(x) \;=\; \alpha_i x_i^{\alpha_i - 1} \prod_{j \neq i} x_j^{\alpha_j} \;>\; 0 \quad \text{for all } x \in \mathbb{R}^k_{++}.$$

The key shape property is *strict quasi-concavity*[4], which implies convex, "nicely curved" upper contour sets and uniqueness of the demanded bundle (I mean, given by all the properties of preferences that *strict quasi-concavity* implies). A convenient way to see this is to use a monotonic transformation.

Define

$$v(x) \;:=\; \log u(x) \;=\; \sum_{i=1}^{k} \alpha_i \log x_i, \qquad x \in \mathbb{R}^k_{++}.$$

The logarithm is strictly increasing, so u and v represent the *same* preferences: for any $x, y \in \mathbb{R}^k_{++}$,

$$u(x) \geq u(y) \iff v(x) = \log u(x) \;\geq\; \log u(y) = v(y).$$

Thus, u is (strictly) quasi-concave if and only if v is (strictly) quasi-concave.

[4] I was wrong assuming that Cobb-Douglas is concave. But I have an excusse: it is concave if the sum of exponents is less than 1. My confusion comes from the fact that concave function presents decreasing returns to scale (in a production function approach but extrapolated). But thats only true in that case. A simple proof can be found here.

Now observe that v is actually *strictly concave* on $\mathbb{R}^k_{++}$:

- Each function $x_i \mapsto \log x_i$ is strictly concave on $(0, \infty)$.
- A positive linear combination of strictly concave functions is strictly concave (here it is simply a sum of strictly concave).
- Therefore $v(x) = \sum_i \alpha_i \log x_i$ is strictly concave on $\mathbb{R}^k_{++}$.

Formally, the Hessian of v is diagonal with negative entries on $\mathbb{R}^k_{++}$:

$$\frac{\partial^2 v}{\partial x_i^2}(x) \;=\; -\frac{\alpha_i}{x_i^2} \;<\; 0, \quad \frac{\partial^2 v}{\partial x_i \partial x_j}(x) = 0 \text{ for } i \neq j,$$

so the Hessian is negative definite and v is strictly concave. Strict concavity of v implies strict quasi-concavity of v, and because u is just a strictly increasing transformation of v, u is strictly quasi-concave as well.

Role of the exponents α_i. The parameters α_i affect both how much the consumer values each good and how the utility surface bends in each direction. It is useful to separate three ideas:

- **One dimension vs. many dimensions.** In one dimension, $f(x) = x^\alpha$ is:
  - concave on $(0, \infty)$ if $0 < \alpha \leq 1$,
  - linear if $\alpha = 1$,
  - convex if $\alpha > 1$.

  However, in one dimension *any* strictly increasing f generates strictly quasi–concave preferences, because every upper contour set $\{x : f(x) \geq c\}$ is just an interval $[x_c, \infty)$, which is convex. So "convex utility" in 1D does *not* destroy quasi–concavity.

- **What happens when some $\alpha_i > 1$ in Cobb–Douglas?** The raw Cobb–Douglas

  $$u(x) = \prod_{i=1}^{k} x_i^{\alpha_i}$$

  need not be concave as a function on $\mathbb{R}^k_{++}$ when some $\alpha_i > 1$. Looking at u directly can therefore be misleading. For the maximization problem, we instead work with the monotone transformation

  $$v(x) = \log u(x) = \sum_{i=1}^{k} \alpha_i \log x_i.$$

  This v is *strictly concave for any* $\alpha_i > 0$, because each $\log x_i$ is strictly concave and we are taking a positive linear combination. The Hessian of v is diagonal with entries

  $$\frac{\partial^2 v}{\partial x_i^2}(x) = -\frac{\alpha_i}{x_i^2} < 0,$$

and zeros off the diagonal, so the Hessian is negative definite on $\mathbb{R}^k_{++}$. Since $\log(\cdot)$ is strictly increasing, u and v have the *same* argmax on any budget set and represent the same preferences. Thus:

- preferences are strictly quasi–concave (upper contour sets are strictly convex),
- equivalently, there *exists* a strictly concave representation of these preferences (namely v),
- and this remains true even if some $\alpha_i > 1$ and the original u is not concave.

This is exactly what fails in "bad" convex examples such as $u(x_1, x_2) = x_1^2 + x_2^2$: in that case no monotone concave transform exists, upper contour sets are not convex, and preferences are not quasi–concave. For Cobb–Douglas, by contrast, the log-transformed v rescues concavity and convex preferences, regardless of whether some exponents exceed one.

- **Economic meaning of α_i.** The parameters α_i control how strongly the consumer values each good. In the common normalization $\sum_i \alpha_i = 1$, the Marshallian demand takes the form
$$x_i(p, y) = \frac{\alpha_i y}{p_i},$$
so the consumer spends a fraction α_i of income on good i. Larger α_i means a systematically larger budget share for good i. The strict quasi-concavity and the "nice" indifference curves come from the concavity of $v(x) = \sum_i \alpha_i \log x_i$, not from concavity of u itself.

Economic implications

Because preferences are strictly increasing and strictly quasi-concave, and the budget set

$$B(p, y) = \{x \geq 0 : p \cdot x \leq y\}$$

is convex and compact (for $p \gg 0$, $y > 0$), several consequences follow:

- **Existence:** continuity of u on a compact $B(p, y)$ implies an optimizer exists.
- **Uniqueness:** strict quasi-concavity implies a unique maximizer on a convex set. Intuitively, if two bundles on the budget line gave the same utility, their average would give strictly higher utility, contradicting optimality.
- **Interiority (no corners):** for each good i,
$$\frac{\partial u}{\partial x_i}(x) = \alpha_i x_i^{\alpha_i - 1} \prod_{j \neq i} x_j^{\alpha_j} \to +\infty \quad \text{as } x_i \to 0^+.$$
Marginal utility becomes arbitrarily large as x_i approaches zero. Since prices are finite, the consumer can always increase utility by buying a small amount of any good at zero

consumption. Therefore, with $p_i > 0$ and $y > 0$, it is never optimal to set $x_i = 0$: the optimal bundle lies in $\mathbb{R}^k_{++}$.

Intuitively: the consumer always wants a little of every good, and is willing to pay a lot (in terms of other goods) to avoid having any good at zero. This is why Cobb–Douglas demands are strictly positive.

KKT conditions and the log trick

From an optimization viewpoint, the crucial fact is that $v(x) = \sum_i \alpha_i \log x_i$ is strictly concave and the constraint set is convex. Hence, for the problem

$$\max_{x \geq 0} v(x) \quad \text{s.t. } p \cdot x \leq y,$$

the KKT conditions are:

$$\begin{aligned}
&\frac{\alpha_i}{x_i} \leq \lambda p_i \quad \text{for all } i \quad (\text{equality if } x_i > 0), \\
&y - p \cdot x \geq 0 \quad (\text{equality if } \lambda > 0), \\
&x_i \geq 0, \ \lambda \geq 0.
\end{aligned}$$

Because we are maximizing a strictly concave function over a convex set, these KKT conditions are *necessary and sufficient*: any point solving them is the unique global maximizer of v.

Finally, since u and v induce the same preferences and have the same maximizers on $B(p, y)$, the unique solution to the KKT system for v is also the unique solution to the original problem with u. In practice, one can safely:

(i) take logs to pass from u to v,

(ii) solve the concave problem using KKT,

(iii) and interpret the resulting $x^*(p, y)$ as the Marshallian demand for the Cobb–Douglas utility.

This explains both why the FOCs are enough and why the solution is interior and unique[5].

[5] *Note for Joaquin*: Here is why we can assume that optimum will be interior in the case of CD functions (even in the non-transformed case). Because we can find a monotonous transformation of that function that is strictly concave and then the solution is interior and unique. Since is a monotonous transformation, the solution should be the same of the original function. Then, all $x_i > 0$.

2 CES utilities

Definition and basic properties

A CES (constant elasticity of substitution) utility takes the form

$$u(x) = \left(\sum_{i=1}^{k} \beta_i x_i^{\rho} \right)^{1/\rho}, \qquad \beta_i > 0, \qquad \rho \in \mathbb{R}.$$

It is convenient to consider also the monotone transformation

$$v(x) := \log u(x) = \frac{1}{\rho} \log \left(\sum_{i=1}^{k} \beta_i x_i^{\rho} \right),$$

because u and v represent the same preferences and have the same maximizers.

Concavity and quasi-concavity. The CES class exhibits three regimes depending on ρ:

- **$\rho < 1$ (including $\rho < 0$):** The function u is strictly quasi–concave, and in fact *strictly concave* whenever $\rho < 1$. Intuitively, x_i^{ρ} is concave for $\rho \leq 1$, the weighted sum remains concave, and the outer transformation preserves concavity in this range.

- **$\rho = 1$:** u becomes linear

  $$u(x) = \sum_{i=1}^{k} \beta_i x_i,$$

  i.e., perfect substitutes.

- **$\rho > 1$:** The function becomes convex, preferences cease to be quasi–concave, and indifference curves bend "the wrong way" (similar a $u = x_1^2 + x_2^2$). There is no monotone transformation that restores concavity, which means the consumer tends to choose extreme or corner bundles.

Elasticity of substitution. CES preferences are parameterized so that the elasticity of substitution is

$$\sigma = \frac{1}{1-\rho}.$$

Hence:

- $\rho \to 1 \Rightarrow \sigma \to \infty$ (perfect substitutes),

- $\rho \to 0 \Rightarrow \sigma = 1$ (Cobb–Douglas),

- $\rho \to -\infty \Rightarrow \sigma \to 0$ (perfect complements).

Economic implications

When $\rho < 1$, marginal utility diminishes as a function of each x_i, and the MRS varies smoothly. As a consequence:

- **Strict quasi–concavity** guarantees unique demand.
- **Strict concavity** ($\rho < 1$) ensures global interior solutions whenever $p_i > 0$ and $y > 0$.
- The consumer always demands a positive amount of every good, except in degenerate price configurations.

As ρ approaches the boundaries:

- $\rho \to 1$: substitution is so easy that any good can replace any other; corners arise just like with perfect substitutes.
- $\rho \to -\infty$: substitution becomes impossible; consumption fixes proportions (Leontief case).

Shape of the solution

Under $\rho < 1$, the first-order conditions of the concave problem

$$\max_{x \geq 0} \left(\sum_{i=1}^{k} \beta_i x_i^{\rho} \right)^{1/\rho} \qquad \text{s.t. } p \cdot x \leq y$$

imply that the optimal bundle satisfies, for all i, j,

$$\frac{\beta_i x_i^{\rho-1}}{p_i} = \frac{\beta_j x_j^{\rho-1}}{p_j}.$$

This yields a proportionality rule:

$$x_i \propto \left(\frac{\beta_i}{p_i} \right)^{\frac{1}{1-\rho}}.$$

Imposing the budget constraint determines the unique scalar of proportionality.

KKT status

- **If $\rho < 1$:** u is concave on the relevant domain; the problem is a concave maximization with a convex constraint set. Therefore, the KKT conditions are *necessary and sufficient* for optimality, and yield a unique interior bundle (for $p \gg 0$, $y > 0$).
- **If $\rho = 1$:** u becomes linear,

$$u(x) = \sum_{i=1}^{k} \beta_i x_i,$$

i.e. preferences are those of perfect substitutes. Linear functions are both concave and convex, so the maximization problem is still a concave program (a concave objective over a convex set). Hence the KKT conditions are again *necessary and sufficient* for global optimality. The difference with the $\rho < 1$ case is not the validity of KKT, but the *shape* of the solution:

  - generically, the optimum is at a corner (the consumer spends all income on any good with maximal β_i / p_i);
  - when there are price-adjusted ties, there is a whole *continuum* of optimal bundles (the entire budget line segment spanned by tied goods), so KKT admits many solutions and demand is set-valued.

- **If $\rho > 1$:** u is convex, preferences are not quasi–concave, and the problem is no longer concave. In this regime KKT conditions are *necessary only*: they can identify stationary points that are not global maxima (and may fail to rule out interior minima or saddle points). As with other convex utilities (e.g. $u = x_1^2 + x_2^2$), the true optimal choices are typically at extreme points of the budget set, i.e. corners such as "all income spent on the cheapest good". Verification requires explicitly comparing utilities across candidate bundles.

3 Quasi–linear utilities

Definition and basic properties

A quasi–linear utility takes the form

$$u(x) = x_1 + v(x_2, \dots, x_k),$$

where x_1 is the *numeraire* good and $v : \mathbb{R}_+^{k-1} \to \mathbb{R}$ is increasing (in each argument) and typically concave.

Basic properties on $\mathbb{R}_+^k$:

- **Continuity and monotonicity.** If v is continuous and increasing in each x_i for $i \geq 2$, then u is continuous and strictly increasing in all goods x_i (for $i = 1$ marginal utility is constant and equal to 1).

- **Concavity and quasi–concavity.** If v is concave in $(x_2, \dots, x_k)$, then u is concave in x (sum of a linear function in x_1 and a concave function in $x_{2..k}$). If v is strictly concave in $(x_2, \dots, x_k)$, then u is strictly concave in $(x_2, \dots, x_k)$ and strictly quasi–concave in the full vector x.

- **Shape of indifference curves.** Indifference curves are "parallel" in the direction of x_1: adding one unit of the numeraire always increases utility by exactly 1, independently of the consumption of other goods. This implies a very simple income effect structure (see below).

Economic implications

It is convenient to rewrite the budget constraint as

$$p_1 x_1 + \sum_{i=2}^{k} p_i x_i \le y \quad \iff \quad x_1 \le \frac{y}{p_1} - \sum_{i=2}^{k} \frac{p_i}{p_1} x_i,$$

and view x_1 as "money left over" after choosing $(x_2, \ldots, x_k)$. Substituting this into u,

$$u(x_1, x_2, \ldots, x_k) = x_1 + v(x_2, \ldots, x_k) \le \frac{y}{p_1} - \sum_{i=2}^{k} \frac{p_i}{p_1} x_i + v(x_2, \ldots, x_k),$$

with equality when the budget binds.

Thus the consumer's problem can be decomposed as:

1. Choose $(x_2, \ldots, x_k)$ to maximize

$$\tilde{v}(x_2, \ldots, x_k) := v(x_2, \ldots, x_k) - \sum_{i=2}^{k} \frac{p_i}{p_1} x_i,$$

   subject to $x_i \ge 0$ and the implicit constraint that x_1 computed as residual is nonnegative.

2. Set x_1 as the residual money:

$$x_1 = \frac{y}{p_1} - \sum_{i=2}^{k} \frac{p_i}{p_1} x_i.$$

Two key implications:

- **No income effects for non–numeraire goods (away from the $x_1 = 0$ corner).** The objective in step (1) does not depend on y. If the optimum of step (1) yields a residual $x_1 > 0$, then a small change in y alters only x_1, not $(x_2, \ldots, x_k)$. Thus Marshallian demand for goods $2, \ldots, k$ is *independent of income* as long as the numeraire is not at a corner. This is the hallmark of quasi–linearity.
- **All income effects absorbed by the numeraire.** Changes in y simply translate into changes in x_1 at rate $1/p_1$, holding the optimal $(x_2, \ldots, x_k)$ fixed (until x_1 hits zero). Graphically, indifference curves are vertically shifted copies of each other in the x_1 direction.

KKT status

If v is concave in $(x_2, \ldots, x_k)$, then $u(x) = x_1 + v(x_2, \ldots, x_k)$ is concave in x, and strictly concave in $(x_2, \ldots, x_k)$ if v is strictly concave. Combined with the linear budget constraint, this implies:

- The consumer problem is a concave maximization over a convex set; under mild constraint qualifications (which hold here for $p \gg 0, y > 0$), the KKT conditions are *necessary and sufficient* for optimality.
- If v is strictly concave, the solution for the non–numeraire goods $(x_2^*, \ldots, x_k^*)$ is unique. The numeraire x_1^* is then uniquely pinned down by the budget, so the entire bundle is unique whenever the budget binds and $x_1^* > 0$.
- If v is concave but not strict (e.g. has linear segments), KKT still characterize the set of global maximizers, but demand may be set–valued in those directions (the consumer is indifferent among some variations of $(x_2, \ldots, x_k)$).

In summary, quasi–linear utilities with concave v combine:

- a clean decomposition of the problem into a "nonlinear block" and a residual numeraire,
- exact KKT characterizations (necessary and sufficient),
- absence of income effects for non–numeraire goods away from the $x_1 = 0$ corner.

4 Summary of Other Common Utility Forms

In addition to Cobb–Douglas, CES, and quasi–linear utilities, three further specifications frequently appear in exercises. They fit naturally within the general framework already developed (concavity, quasi–concavity, interior vs. corner solutions, and the role of KKT). This section provides a compact synthesis.

4.1 Perfect Substitutes

Definition and properties.

$$u(x) = \sum_{i=1}^{k} a_i x_i, \qquad a_i > 0.$$

Linear utilities are continuous, monotone, concave, convex, but *not* strictly quasi–concave. Indifference curves are straight lines with constant slope.

Economic meaning. Goods are completely interchangeable: the consumer cares only about utility per dollar a_i / p_i.

Shape of the solution. Optimal consumption is a **corner**:

$$x^* \in \arg\max_i \frac{a_i}{p_i}.$$

If several goods tie, the whole budget segment between those goods is optimal (set–valued demand).

KKT status. Because the maximization of a linear function over a convex set is a concave program, *KKT conditions are necessary and sufficient*. Non-uniqueness arises from the geometry (lack of strict quasi–concavity), not from failure of KKT.

4.2 Perfect Complements (Leontief)

Definition and properties.

$$u(x) = \min_i \left\{ \frac{x_i}{\alpha_i} \right\}, \qquad \alpha_i > 0.$$

Being the minimum of linear functions, u is concave, continuous, and strictly increasing along rays that preserve proportions. Indifference curves exhibit right-angle kinks and zero substitutability. Usually we see these preferences with $\alpha_i = 1$ but this is more generally applicable.

Economic meaning. Goods must be consumed in fixed proportions; substituting one for another has no value.

Shape of the solution. The optimizer sits exactly at the kink:

$$x_i^* = \alpha_i \frac{y}{\sum_i p_i \alpha_i}.$$

Thus the optimal bundle is **unique** even though u is not differentiable.

KKT status. u is concave but not differentiable; KKT conditions are more complicated to satisfy, but apparently they are **necessary and sufficient**. The structure is simple because the kink identifies the unique optimal ratio of goods.

4.3 Convex Utilities (Non–Quasi–Concave)

Definition and properties. Example:

$$u(x) = x_1^2 + x_2^2.$$

Such utilities are convex and their upper contour sets are not convex; preferences violate quasi–concavity. Indifference curves bend "outward," encouraging extremal choices.

Economic meaning. Consumers strictly prefer extreme bundles to interior mixes: they "love" specialization.

Shape of the solution. True maximizers lie at **corners**. Typically the consumer spends all income on the cheapest good. If multiple goods share the lowest price, the entire budget line is optimal.

KKT status. KKT conditions are *only necessary*. They may produce interior stationary points that are minima or saddles. Therefore, solving requires:

1. listing interior and boundary candidates via KKT,
2. evaluating u at each feasible candidate,
3. selecting the global maximizer manually.