Preferences & Utility

ECON 410

May 18, 2023

Introduction

- Economics is all about optimal decision making
- People, firms, and governments are often faced with a set of alternatives which they must decide between
  - Students decide which college major to pick
  - Firms decide how to price their goods
  - Governments decide what to set the corporate tax rate to
- Their ultimate decision will depend on their *preferences* over the set of alternatives
- In this section, we'll outline the mathematical foundations of modeling preferences

Why are Models Useful?

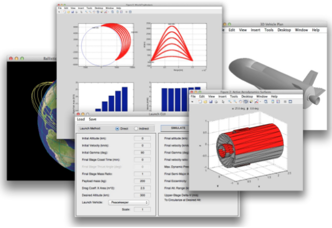

- Prior to launching a rocket, scientists model the rocket's dynamics
- Why?
  - Can assess the impact of the rocket's shape, weight, etc. on its performance
  - Can assess the impact of factors such as weather on the flight trajectory
- To ensure no surprises at time of launch, many simulations are run beforehand

Why are Models Useful?

- What we do in economics is very similar
- We write down models of individual & firm economic behavior
- Using these models, we can assess the impact of various economic policies on individual & firm decision making
  - Impact of minimum wage changes?
  - Impact of subsidized college?
  - Impact of fossil fuel taxes?
- This process helps companies, individuals, governments, etc. make informed decisions prior to adopting any new policies

Preferences

Alternatives

- We start with a set of alternatives X
  - Ex: $X = \{Economics, English, Biology\}$... what major to pick?
  - Ex: $X = \{Chicken, Pasta, Salad\}$... what to have for lunch?
  - Ex: $X = [0, \infty)$... how many dollars to invest in the stock market?
- The set of alternatives contains all things which a decision maker (DM) is choosing between
- We refer to X often as the DM's *choice set*
- X can be discrete (ex. $X = \{a, b, c\}$) or continuous (ex. $X = [0, \infty)$)
  - Usually, we'll work with continuous choice sets

Preferences

- We need a way to describe a DM's *preferences* over the set of alternatives
- Take two elements x and y in the set of alternatives X
- If x is preferred to y, we use the notation:

$$x \succ y$$

- Examples:
  - Apple $\succ$ Banana
  - Chocolate $\succ$ Vanilla
  - \$10 $\succ$ \$5

Preferences

- The preference relation $\succ$ orders the elements of X based upon the DM's preferences
- $\succ$ completely characterizes the DM's preferences over X
- If $X = \{apple, orange, banana\}$, a preference over X could be:
  - apple $\succ$ banana, banana $\succ$ orange, apple $\succ$ orange

Indifference

- A DM may prefer x to y ($x \succ y$) or prefer y to x ($y \succ x$)
- Another possibility is that the DM may be *indifferent* between x and y
- If the DM is indifferent between x and y, they think x is exactly as good as y
- If a DM is indifferent between x and y, we use the notation:

$$x \sim y$$

Preferences: Summary

- Suppose we are face with choice set $X = \{a, b\}$
- If a is preferred to b:
  - $a \succ b$
- If b is preferred to a:
  - $b \succ a$
- If we are indifferent between a and b:
  - $a \sim b$

Rationality

- With only two alternatives a and b, writing out a preference ordering is simple
- But with many alternatives, writing out a preference ordering can quickly become very messy
  - For example, how could I possibly write my preference ordering over dollars in $[0, \infty)$?
- Thankfully, there is a much easier way to represent preferences than explicitly writing out the preference ordering
- We can encode a preference into a *utility function*

Rationality

- Before discussing utility functions, we'll need to define *rationality*
  - It will soon be clear why
- Rationality is a very important concept in economics
- Many people misunderstand what rationality means in economics
- Before defining rationality, we need to introduce two concepts:
  1. Completeness
  2. Transitivity

Completeness

- Given a set of alternatives X, a preference $\succ$ over X is complete if:

$$x \succ x' \,,\; x' \succ x \,,\; \text{or} \;\; x \sim x' \quad \text{for all} \;\; x, x' \in X$$

- In words: a preference is complete if all of the relevant alternatives are comparable

- For example, let $X = \{a, b, c\}$

- An example of a complete preference ranking is:
  - $a \succ b$, $b \succ c$, $a \succ c$

- An example of an incomplete preference is:
  - $a \succ^* b$, $a \succ^* c$
  - $\succ^*$ takes no stand on the ranking of b and c

Transitivity

- Given set of alternatives $X = \{a, b, c\}$, a preference $\succ$ over X is transitive if:

$$a \succ b, b \succ c \rightarrow a \succ c$$

- In words: a preference is transitive if $a \succ b$ and $b \succ c$ imply that $a \succ c$
- An example of a transitive preference is:
    - $a \succ b$, $b \succ c$, $a \succ c$
- An example of an intransitive preference is:
    - $a \succ b$, $b \succ c$, $c \succ a$
    - "Cyclical" preferences like this are never transitive

Rationality

Definition

A preference relation $\succ$ is *rational* if it is complete and transitive.

- Completeness and transitivity are two very mild restrictions on preferences
  - Just require everything to be rankable, and for the ranking to not be cyclical
- Why do we bother with the notion of rationality?

Utility Functions

Theorem

A preference relation $\succ$ can be represented by a utility function only if it is rational.

- Preference relations are complicated
- We can take a DM's preference $\succ$ and encode it in a *utility function*, which is much easier to work with
- A utility function $u : X \rightarrow \mathbb{R}$ is a function mapping from the set of alternatives X to the real numbers $\mathbb{R}$
  - Simply assigns a "utility value" to each alternative in X
- Let's explore utility functions in more detail

Utility Functions

- A utility function $u : X \to \mathbb{R}$ is a function mapping from the set of alternatives X to the real numbers $\mathbb{R}$
- Importantly, utility functions satisfy the condition:

$$u(x) > u(y) \iff x \succ y$$

- Utility from x is larger than utility from y if and only if x if preferred to y
- Rather than keeping track of the explicit preference ordering, we assign numeric values to all alternatives, where better alternatives get higher values

Utility Functions (Example)

- For $X = \{a, b, c\}$, my preferences may be: $a \succ b$, $b \succ c$, $a \succ c$
- A utility function representing these preferences could be:

$$u(x) = \begin{cases} 6 & \text{if } x = a \\ 4 & \text{if } x = b \\ 2 & \text{if } x = c \end{cases}$$

- $u(x)$ preserves my preference ordering over a, b, and c
- Importantly, for a given preference $\succ$, there is not a *unique* utility function which represents it

Utility Functions (Example)

- Another utility function representing my preferences could be:

$$u_2(x) = \begin{cases} 3 & \text{if } x = a \\ 2 & \text{if } x = b \\ 1 & \text{if } x = c \end{cases}$$

- $u_2(x)$ also preserves my preference ordering over a, b, and c, so is a perfectly valid representation of my preferences
- The actual value of utility is meaningless, only the ordering of alternatives matters
  - Ordinal interpretation, but no cardinal interpretation

Monotone Transformations of Utility

- In the previous example, $u(x) = 2u_2(x)$
  - $u(x)$ and $u_2(x)$ represent the same preferences
- In fact, if $u(x)$ represents preferences $\succ$, any *monotone* transformation of $u(x)$ also represents $\succ$
- If $g(x)$ is a monotone transformation of $u(x)$, then:

$$g(a) > g(b) \iff u(a) > u(b)$$

- Monotone transformations are order-preserving transformations
- For example, the two utility functions represent the same preferences:

$$u_1(x) = x$$
$$u_2(x) = \alpha + \beta x \quad \beta > 0$$

Utility Functions (Example #2)

- Let's denote the amount of money I have by x
- I prefer more money to less, for example $\$10 \succ \$5$
- Some utility functions representing my preference over money could be:

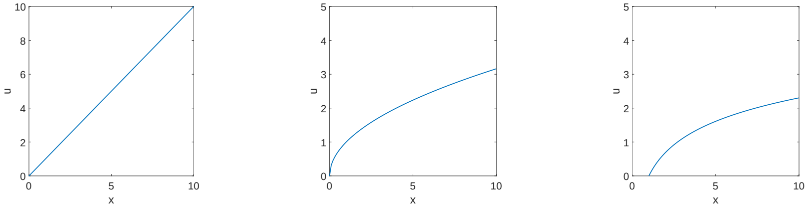

(a) $u(x) = x$ (b) $u(x) = \sqrt{x}$ (c) $u(x) = log(x)$

Marginal Utility

- If I were to consume one more unit of x, how would my utility change?
- To answer this, we revisit the concept of a derivative from calculus
- For a given function $y = f(x)$, we denote its derivative by $\frac{dy}{dx}$
- $\frac{dy}{dx}$ gives the change in y following a 1 unit increase in x
- Example: $y = 2x$
  - $\frac{dy}{dx} = 2$
  - In words: if x increases by 1, y increases by 2

Marginal Utility

- We can do the exact same thing with utility functions
- Suppose my preferences over x are represented by: $u(x) = 3 + 4x$
- If I increase my consumption of x by 1, the resulting change in my utility is $\frac{du}{dx} = 4$
  - In words: consuming 1 more x increases my utility by 4
- We refer to the quantity $\frac{du}{dx}$ as the *marginal utility* of x
- Frequently, we'll write $MU_x = \frac{du}{dx}$ to simplify notation

Marginal Utility

- For example, if $u(x) = 4x + 1$, then:

$$MU_x = \frac{du}{dx} = 4$$

- If x increases by 1, $u(x)$ increases by 4
- Alternatively, if $u(x) = 3x^{\frac{1}{3}}$, then:

$$MU_x = \frac{du}{dx} = x^{-\frac{2}{3}}$$

- If x increases by 1, $u(x)$ increases by $x^{-\frac{2}{3}}$
- Notice that MU_x can depend on the actual value of x

"Goods" and "Bads"

- If a good increases utility, we refer to it as a *good* good
- If a good decreases utility, we refer to is as a *bad* good
- Suppose we are given a utility function $u(x)$. How do we tell if x is a good or a bad?
- Answer: check the sign of the marginal utility
  - If $MU_x > 0$, consuming more x increases utility $\rightarrow$ x is a good good
  - If $MU_x < 0$, consuming more x decreases utility $\rightarrow$ x is a bad good

“Goods” and “Bads”

- Suppose my utility from consumption of pickles is $u(p) = -p^2$
- The marginal utility of an additional pickle is:

$$MU_p = \frac{du}{dp} = -2p < 0$$

- $MU_p < 0$, so pickles are “bads” in this case
- Consuming pickles decreases my utility
  - The less pickles I consume, the better

"Goods" and "Bads"

- Suppose my utility from consumption of chocolate is $u(c) = \sqrt{c}$
- The marginal utility of an additional chocolate is:

$$MU_c = \frac{du}{dc} = \frac{1}{2}c^{-\frac{1}{2}} > 0$$

- $MU_c > 0$, so chocolate is a "good" good
- Consuming chocolate increases my utility
  - The more chocolate I consume, the better
- Notice that as I mentioned before, MU_c depends on c
  - Marginal utility differs depending on how much c I am already consuming

Diminishing Marginal Utility

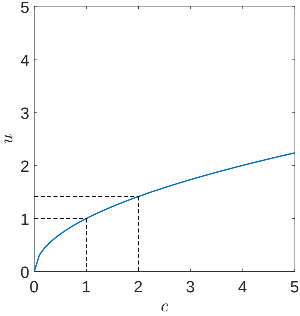

- Given the utility function $u(c) = \sqrt{c}$:
  - The first unit of c increases u by 1 ($MU_c = 1$)
  - But, the second unit of c only increases u by $\approx .5$ ($MU_c \approx .5$)
- $u(c) = \sqrt{c}$ exhibits *diminishing marginal utility*
- I always like c, but I like each additional unit $<$ the previous one

Diminishing Marginal Utility

- Diminishing marginal utility captures the natural case when a DM likes something, but likes it less and less the more they consume
  - Ex: first slice of pizza is great, but after 5 I start getting tired of it
- How do we tell if a utility function $u(x)$ exhibits diminishing marginal utility?
  - Compute MU_x
  - See if it diminishes with x: Does $\frac{dMU_x}{dx} < 0$?
- If $\frac{dMU_x}{dx} < 0$, then $u(x)$ exhibits diminishing marginal utility
- Equivalently, if $\frac{d^2u}{dx^2} < 0$, then $u(x)$ exhibits diminishing marginal utility

Increasing Marginal Utility

- Alternatively, there may be cases where a DM likes something more and more the more they consume
- Their marginal utility *increases* in this case
- Similar to before, if $\frac{dMU_x}{dx} > 0$, then $u(x)$ exhibits *increasing marginal utility*
  - I like x, and I like x more and more the more I consume
- Equivalently, if $\frac{d^2u}{dx^2} > 0$, then $u(x)$ exhibits increasing marginal utility

Constant Marginal Utility

- Marginal utility can also be constant
- In particular, utility functions of the form: $u = a + bx$ will always have constant marginal utility
- Notice that $MU_x = b$, so $\frac{dMU_x}{dx} = 0$
- Linear utility functions exhibit constant marginal utility

Marginal Utility - Summary

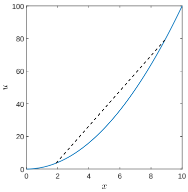

(d) Increasing MU

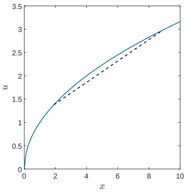

(e) Decreasing MU

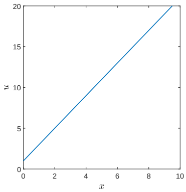

(f) Constant MU

- Convex utility functions exhibit increasing MU
- Concave utility functions exhibit decreasing MU
- Linear utility functions exhibit constant MU
- Diminishing MU will be the case we see most often

Multivariable Utility Functions

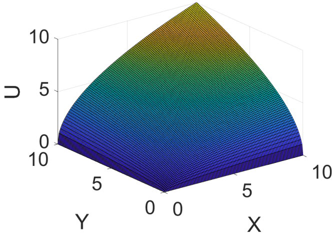

- Most of the time in this class, we will focus on utility functions which depend on two variables: $u(x, y)$
- *Trade offs* are an important consideration in economics, and are difficult to talk about with only one good
- With two goods, we can say quite a bit about optimal behavior in the face of trade offs

Marginal Utility with Two Goods

- Given a utility function $u(x, y)$, computing the marginal utility of x and y works very much the same as before
- Rather than taking "normal" derivatives, we take *partial* derivatives:

$$MU_y = \frac{\partial u}{\partial y}$$
$$MU_x = \frac{\partial u}{\partial x}$$

- Increasing y by 1 increases u by $\frac{\partial u}{\partial y}$, all else equal (ceteris paribus)
- Increasing x by 1 increases u by $\frac{\partial u}{\partial x}$, all else equal (ceteris paribus)

Indifference Curves

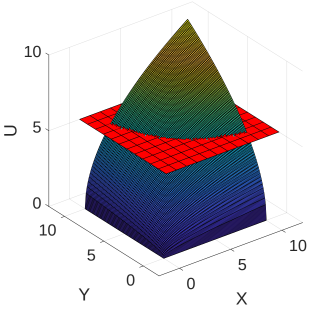

- If we fix the level of utility to $u(x, y) = 5$, we see that there is a curve of (x, y) combinations along this level of $u(x, y)$
- Any (x, y) combination on this curve yields $u(x, y) = 5$

Indifference Curves

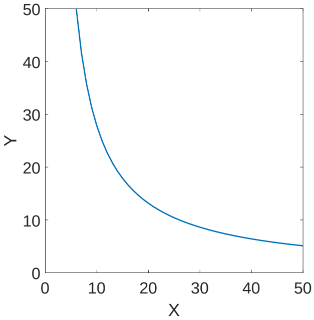

- This gives a *level curve* of $u(x, y)$
  - All combinations of (x, y) yielding the same *level* of $u(x, y)$
- Each point on the curve is a potential "bundle" of (x, y) we can consume
- Each (x, y) combination along this curve gives us the same utility

Indifference Curves

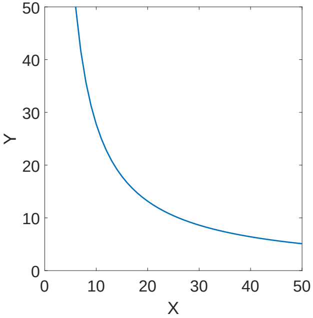

- If (x, y) and (x', y') are two points on the curve above:

$$(x, y) \sim (x', y')$$

- We refer to these curves as *indifference curves*

Indifference Curves

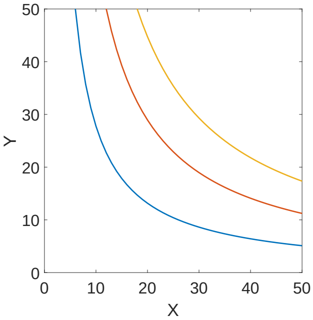

- Each indifference curve corresponds to a different level of utility u
- Importantly, indifference curves can never cross

Slope of Indifference Curves

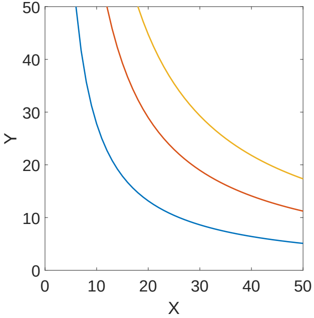

- What is the slope of an indifference curve?
- The slope of an indifference curve is given by the ratio of marginal utilities (times -1):

$$slope(IC) = -\frac{MU_x}{MU_y} = -MRS$$

Marginal Rate of Substitution

$$MRS = \frac{MU_x}{MU_y} = \frac{\frac{\partial u}{\partial x}}{\frac{\partial u}{\partial y}} = \frac{\partial u}{\partial x}\frac{\partial y}{\partial u} = \frac{\partial y}{\partial x}$$

- The MRS tells us how y changes with x, holding u constant
- In other words: if I increased my consumption of x, how would I have to change my consumption of y in order to leave utility unchanged?
- The MRS reveals a lot about the trade off DMs face between x and y
- To see this, let's continue talking about the slope of ICs

Slope of Indifference Curves

$$slope(IC) = -\frac{MU_x}{MU_y}$$

- Back to the slope of ICs
- Notice that the sign of the slope is determined by the signs of MU_x and MU_y
- In other words: whether the ICs slope up or down depends on whether x and y are goods or bads

Slope of Indifference Curves (Two Goods)

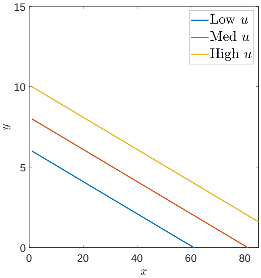

- If $MU_x > 0$ and $MU_y > 0$, MRS will be > 0 and $slope(IC)$ will be < 0
- In words: if x and y are both good goods, ICs will slope downwards
- Downward slope represents a trade off:
  - In order to keep utility fixed, if I gain some x I must give up some y
- Utility increases as we move to the top right

Slope of Indifference Curves (Two Bads)

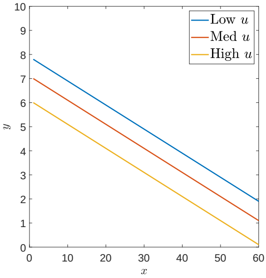

- If $MU_x < 0$ and $MU_y < 0$ (two bads), $slope(IC)$ will again be < 0
- Utility increases as we move towards the origin
  - The less I consume, the better
- Won't really consider this case much

Slope of Indifference Curves (y Good, x Bad)

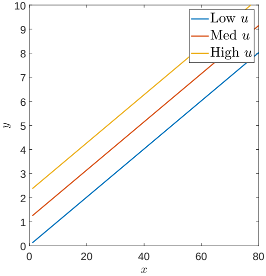

- If $MU_x < 0$ and $MU_y > 0$, $slope(IC) > 0$
- With one good and one bad, ICs slope upwards
- Intuition: if you force me to consume more x, you must compensate me with y in order to leave my utility unchanged

Slope of Indifference Curves (y Bad, x Good)

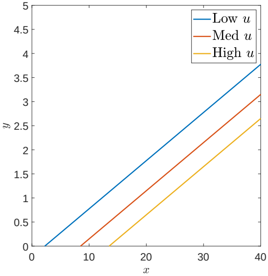

- If $MU_x > 0$ and $MU_y < 0$, $slope(IC) > 0$
- Mirror case of previous slide: we want to increase x and decrease y as much as possible

Curvature of Indifference Curves

- So far, we've looked at the slope of ICs
  - Whether ICs slope up or down depends on whether x and y are good/bad
- What about the curvature of ICs?
  - i.e. when are ICs convex/concave/linear?
- For a given utility function, the curvature of its ICs will have important implications in the sections to come
- When discussing the curvature of ICs, let's restrict attention to the case when $MU_x > 0$ and $MU_y > 0$ (so $MRS > 0$ and $slope(IC) < 0$)

Curvature of Indifference Curves

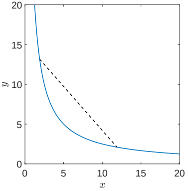

(g) Convex IC

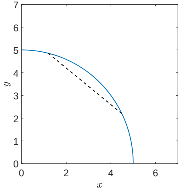

(h) Concave IC

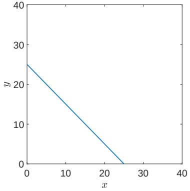

(i) Linear IC

- Convex ICs $\rightarrow$ IC gets flatter as x increases
- Concave ICs $\rightarrow$ IC gets steeper as x increases
- Linear ICs $\rightarrow$ constant slope as x increases

Curvature of Indifference Curves

$$slope(IC) = -MRS$$

- Recall the IC slope formula
- And recall that we are assuming $MRS > 0$ (ICs slope down)
- The smaller the magnitude of the MRS, the flatter the IC

Convex IC

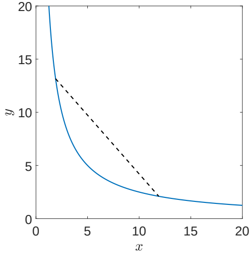

- For a convex IC, as x increases, magnitude of $slope(IC)$ decreases
- This requires:

$$\frac{\partial MRS}{\partial x} < 0$$

- IC is convex if MRS decreases with x

Concave IC

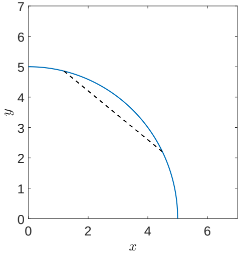

- For a convex IC, as x increases, magnitude of $slope(IC)$ increases
- This requires:
$$\frac{\partial MRS}{\partial x} > 0$$
- IC is concave if MRS increases with x

Linear IC

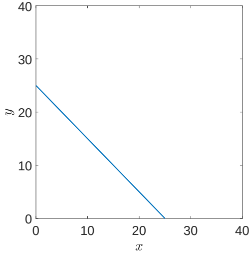

- For a convex IC, as x increases, magnitude of $slope(IC)$ remains constant
- This requires:

$$\frac{\partial MRS}{\partial x} = 0$$

- IC is linear if MRS does not change with x

Curvature of Indifference Curves (Summary)

- For a given utility function $u(x, y)$, to determine whether its ICs are convex, concave, or linear:
  - Derive the *MRS*
  - Check the sign of $\frac{\partial MRS}{\partial x}$
- If $\frac{\partial MRS}{\partial x} < 0 \rightarrow$ IC is convex
- If $\frac{\partial MRS}{\partial x} > 0 \rightarrow$ IC is concave
- If $\frac{\partial MRS}{\partial x} = 0 \rightarrow$ IC is linear

Common Utility Functions

- We'll frequently encounter the following utility functions:
  - Perfect substitutes: $u(x, y) = ax + by$
  - Perfect complements: $u(x, y) = min\{ax, by\}$
  - Cobb Douglas: $u(x, y) = cx^a y^b$
  - Quasi-linear: $u(x, y) = f(x) + by$ or $u(x, y) = ax + f(y)$
- Let's take a quick look at each
- We'll again assume x and y are both good goods in the slides that follow

Perfect Substitutes

$$u(x, y) = ax + by$$

- Perfect substitute utility is useful when the DM views x and y as perfectly substitutable
  - Butter vs. margarine
  - iPhone vs. Android
  - HP laptop vs. Dell laptop
- x and y do not complement each other, DM is fine consuming just x or just y
- Linear in both x and y
  - An implication: marginal utility of both goods is constant

Perfect Substitutes

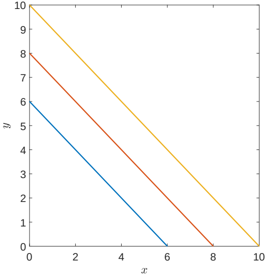

- For perfect substitutes:

$$MRS = \frac{a}{b}$$

- $\frac{\partial MRS}{\partial x} = 0$, so slope of IC does not change as we move from left to right
- ICs have constant slope as a result

Perfect Complements

$$u(x, y) = min\{ax, by\}$$

- Perfect complements utility is useful when x and y are to be consumed in **fixed proportion**. For example:
  - One left shoe per right shoe
  - One scoop of peanut butter per one scoop of jelly
  - 4 wheels per one car
- Notice that if $x = 0$ or $y = 0$, $u(x, y) = 0$
  - Ex: if I have one left shoe and zero right shoes, I effectively have no shoes
- x and y must be consumer together in order to generate any utility

Perfect Complements

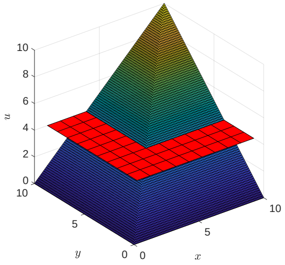

- A perfect complements utility function is shaped somewhat like a pyramid
- Fixing a level of utility, we see the resulting indifference curve looks "L-shaped"

Perfect Complements

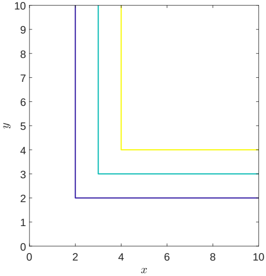

- Perfect complements utility functions have L-shaped ICs
- **Interpretation:** utility only increases when x and y increase together. Increasing one but not the other has no impact on utility.

Cobb-Douglas Utility

$$u(x, y) = cx^a y^b$$

- Above is the general form of the Cobb-Douglas utility function
- This one is very commonly used, and we'll see it over and over again in class
- One thing worth mentioning at this point: Cobb-Douglas utility always has strictly convex ICs

Cobb Douglas Utility

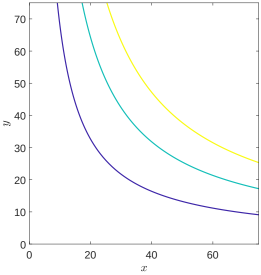

- Cobb-Douglas ICs are always strictly convex
  - This fact will be useful later
- Cobb-Douglas is in a sense “in between” the two extreme cases of perfect substitutes and perfect complements

Quasi-Linear Utility

$$u(x, y) = f(x) + by \quad \text{or} \quad u(x, y) = ax + f(y)$$

- Quasi-linear utility is linear in one good, and non-linear in the other
- Examples:

$$u(x, y) = \sqrt{x} + y$$
$$u(x, y) = 3x + log(y)$$

- Not easy to see at this point, but QL utility is useful when a DM consumes the same amount of a good no matter their income
    - Ex: insulin

Quasi-Linear Utility

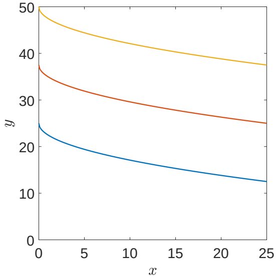

- Suppose that $u(x, y) = f(x) + by$
- For a given level of utility $\bar{u}$, quasi-linear ICs follow the formula:

$$y = \frac{1}{b}\left(\bar{u} - f(x)\right)$$

- ICs will be convex as long as $f(x)$ is concave