# 2 The Elements of a Canonical Model of Rational Consumer Choice

The most fundamental building block in microeconomic theory is the theory of the consumer. In essence, this theory simply says that the consumer picks whatever he or she likes the most among those options that are affordable. In the abstract, this may be considered a tautology: any observed consumer choice can in principle be rationalized by arguing that the basked was chosen because it was the one the consumer liked the most. Since (so far at least) preferences cannot be directly observed this may seem like (and some argue it is) a problem. However, much in the same way as in physics, where theories built on abstract notions as force, mass etc. (which cannot be directly observed) have implications that can be confronted with reality, economic theories built on rational choice have testable implications.

We will develop the model of a rational consumer in the context of a *market economy*, which means that consumers may purchase (or trade) goods and services at known prices (which are not affected by decisions by the individual consumer). Assuming, which we will, that goods are perfectly divisible this makes it very easy to describe the affordable bundles in terms of a *budget set*. The next step introduces some rational preferences, which we will assume are such that they can be described by a numerical *utility function*. The advantage of this is that the behavioral assumption, that the consumer picks his or her most preferred bundle among those that are affordable, can be reformulated as a maximization problem. This allows us to use standard calculus techniques when analyzing the consumer choice model.

We will for most of our discussion imagine a consumer who lives in a very simple world with two perfectly divisible goods.

- There are two goods. quantities will be denoted $x_1$ & $x_2$ respectively.
- The consumer has rational preferences represented by utility function $u(x_1, x_2)$

Recall from our discussion on preference relations that this means

$$(x_1', x_2') \succeq (x_1'', x_2'') \Leftrightarrow u(x_1', x_2') \geq u(x_1'', x_2'')$$

Observe that I construct indifference curves explicitly from the utility functions while Varian introduces indifference curves *before* even getting to utility functions. Both these approaches are equally valid, but I find the construction based on utility functions a bit more straightforward.

## 2.1 Budget Sets

*Optional Reading: Varian, Chapter 2.*

### 2.1.1 The Budget Constraint with Two Goods

For simplicity, we will usually consider decision problems in a simplified setting with only two goods. Obviously, this isn't a very good description of many real world problems, but, maybe surprisingly, it is a rich enough setup to analyze many trade-offs that are relevant for real world decision problems. The advantage of the restriction to two goods is that the budget constraint (and preferences) can be represented geometrically in a 2-dimensional graph.

A *consumption bundle* is denoted as a pair $x = (x_1, x_2)$ (the abbreviated symbol "$x$" instead of $(x_1, x_2)$ is useful sometimes simply to reduce the amount of ink needed, but I will warn you when I use $x$ to denote a vector (=pair if two goods).

The first restriction on a *feasible* consumption bundle is that consumption of both goods must be positive. This is an obvious "technological" restriction which simply says that it is impossible to consume negative quantities. Obviously a model that allows for consumption of, say, $-10$ Hershey bars would be borderline ridiculous, so we assume that $x_1 \geq 0$ and $x_2 \geq 0$ in order for consumption bundle $x = (x_1, x_2)$ to be feasible.

However, a theory that allows the consumer to pick just any consumption bundle (with positive consumptions) wouldn't be particularly interesting. Most people are probably painfully aware of the fact that there are things we would want (who doesn't want a $4 \times 4$ sports utility...) that we can't afford. Let $m$ be the available *income* for the consumer (for now we think of income as a given amount of dollars that the consumer is endowed with exogenously). To be concrete you may think of $m$ as "number of dollars" the consumer has in his or her pocket. Next let $(p_1, p_2)$ denote the prices of the two goods and assume $p_1 > 0$ and $p_2 > 0$ (zero or negative prices are simply not interesting). Given that $m$ is expressed in dollars $p_1$ is then the quantity dollars the consumer has to pay for one unit of good 1 and $p_2$ is the quantity dollars the consumer has to pay for one unit of good 2. Thus, if the consumer buys consumption bundle $(x_1, x_2)$ he/she has to give up

$$p_1 x_1 \text{ dollars in exchange for } x_1$$
$$p_2 x_2 \text{ dollars in exchange for } x_2,$$

so altogether the consumer needs to spend $p_1x_1 + p_2x_2$ dollars and since $m$ is the number of dollars available the *budget constraint* for the consumer is that $(x_1, x_2)$ needs to satisfy

$$p_1 x_1 + p_2 x_2 \leq m$$

Combining the budget constraint with the constraint that consumption bundles must be positive we have described all consumption bundles which are feasible for the consumer (=all bundles the consumer can choose from), which is what we call the *budget set:*

**Definition 1** *Given prices $(p_1, p_2)$ and income $m$ the Budget set of the consumer consists of all consumption bundles $(x_1, x_2)$ such that*

1. $x_1 \geq 0$ *and* $x_2 \geq 0$
2. $p_1 x_1 + p_2 x_2 \leq m$

### 2.1.2 The Budget Constraint with More than Two Goods

Not much changes with more goods and in a sense the model with $n$ goods is no more difficult (or informative) than the model with two goods. The only real drawback with more goods is the graphical representation of the two good case can't be used and that notation becomes a bit uglier. A consumption bundle is then a vector $x = (x_1, ...., x_n)$ and prices are represented by a vector $p = (p_1, ...., p_n)$. The budget constraint is now

$$p_1 x_1 + p_2 x_2 + ...... + p_n x_n \leq m.$$

Using standard summation notation we can write this briefly as $\sum_{i=1}^{n} p_i x_i \leq m$, or we could simply write $px \leq m$, with the understanding that $px$ means a scalar ("dot") product. Geometrically, $px = m$ is what is referred to as a hyperplane, and if we have more than three goods there is simply no way to visualize the budget set.

(You may safely ignore this discussion in small type if you feel lost) Note however that if $y$ is the number of dollars spent on goods $2, ..., n$ and if we let $y$ vary between 0 and $m$ we get a reduced form budget constraint of the form $p_1 x_1 + y \leq m$. Moreover (you will have to take on faith for now since the argument involves solving optimization problems), for each fixed $x_1$ and $y$ we can ask what the best way of spending the $y$ dollars on goods $2, ..., n$ is. From the solution to this problem we get a numerical utility level corresponding to consuming $x_1$ units of good one and spending $y$ dollars *in the best possible way.* Doing this for all $(x_1, y)$ satisfying the constraint $p_1 x_1 + y \leq m$ generates a reduced form utility function that described preferences over $x_1$ and $y$. In combination, this allows us to think of the case with two goods as an analysis of a model with more than two goods, but where all but a single good has been collapsed into an "aggregate" or a "composite good".

### 2.1.3 Graphical Representation of the Budget Set

Returning to the case with two goods, the *budget line* consists of the consumption bundles $(x_1, x_2)$ satisfying

$$p_1 x_1 + p_2 x_2 = m.$$

Bundles on the budget line are bundles that cost exactly $m$ (meaning that there is no money left when the bundle is paid). Some algebra can be used to solve the equation for $x_2$ :

$$\begin{aligned} p_1 x_1 + p_2 x_2 &= m \Leftrightarrow \text{ (subtracting } p_1 x_1 \text{ on both sides)} \\ p_2 x_2 &= m - p_1 x_1 \Leftrightarrow \text{ (dividing by } p_1\text{)} \\ x_2 &= \frac{m}{p_2} - \frac{p_1}{p_2} x_1 \end{aligned}$$

We can now draw the budget line in a two-dimensional picture as in Figure 1. To construct the picture, ask yourself how much the consumer would get if he/she spent everything on good 2. The budget line is constructed so that the whole income is exhausted along the line so we get the answer by setting $x_1 = 0$ in the equation above. Thus,

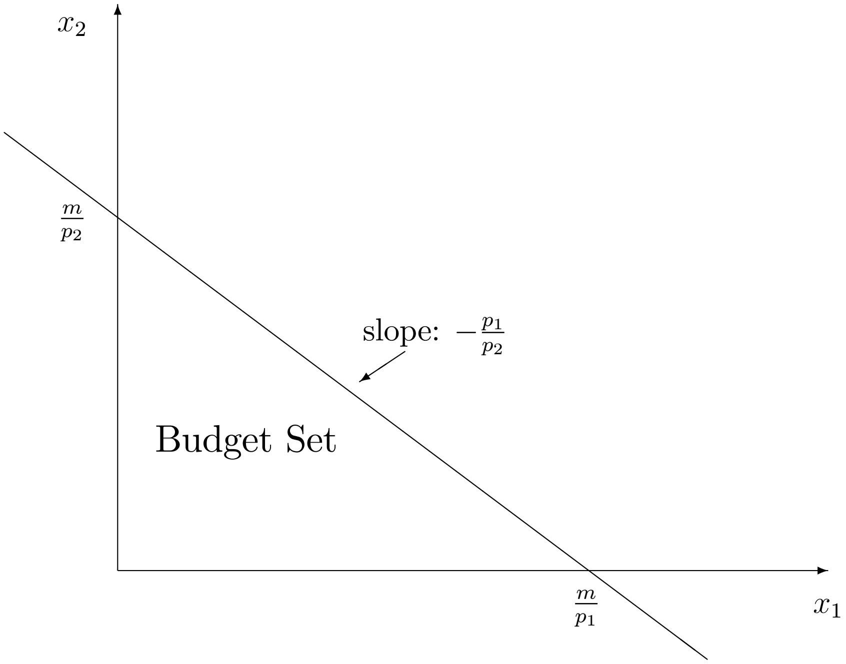

Figure 1: The Budget Set

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2} \cdot 0 = \frac{m}{p_2}$$

is the point where the budget line intersects the $x_2-$axis. Similarly, setting $x_2 = 0$ we get the intersection between the budget line and the $x_1$-axis, yielding

$$\begin{aligned} 0 &= \frac{m}{p_2} - \frac{p_1}{p_2} x_1 \Leftrightarrow \\ x_1 &= \frac{m}{p_1}. \end{aligned}$$

Once you got two points for a straight line you know what the line looks like: just draw a straight line such that both these points are on the line and I can assure you that you got the right line (i.e., there is one and only one straight line that passes though any given distinct pair of points)

It is intuitively clear that the budget set should consist of all bundles that are in the positive orthant and are **below** the budget line. The intuitive understanding should come from the fact that if you move upwards then the consumer buys more of good two (more expensive) and if you move towards the right then the consumer buys more of good 1 (also more expensive). That is,

$$\begin{aligned} p_1x_1 + p_2x_2 &\leq m \Leftrightarrow \text{ (subtracting } p_1x_1 \text{ on both sides)} \\ p_2x_2 &\leq m - p_1x_1 \Leftrightarrow \text{ (dividing by } p_1 > 0) \\ x_2 &\leq \frac{m}{p_2} - \frac{p_1}{p_2}x_1 \end{aligned}$$

Finally just note that everything to the left of the vertical axis are points were $x_1$ is negative and everything below the horizontal axis are points were $x_2$ is negative, so all these points are not included in the budget set for “technological common sense reasons”.

### 2.1.4 The Slope of the Budget Line

Consider two points $x' = (x_1', x_2')$ and $x'' = (x_1'', x_2'')$ that are both on the budget line given (same) prices $(p_1, p_2)$ and income $m$. Then

$$\begin{aligned} p_1x_1' + p_2x_2' &= m \text{ and} \\ p_1x_1'' + p_2x_2'' &= m. \end{aligned}$$

Hence

$$\begin{aligned} p_1x_1' + p_2x_2' &= m = p_1x_1'' + p_2x_2'' \Leftrightarrow \\ p_2x_2' - p_2x_2'' &= p_1x_1'' - p_1x_1' \Leftrightarrow \\ \frac{(x_2'' - x_2')}{(x_1'' - x_1')} &= -\frac{p_1}{p_2} \end{aligned}$$

That is

$$\frac{\text{Change in consumption of good 2}}{\text{Change in consumption of good 1}} = -\frac{p_1}{p_2}$$

By just staring at the budget line we see that this says that $-\frac{p_1}{p_2}$ is the slope coefficient of the budget line and this derivation gives a bit of an interpretation of how to think about this slope in economic terms. We note that:

1. $(x_2'' - x_2')$ and $(x_1'' - x_1')$ must have opposite signs (one negative and one positive)
2. For concreteness, say that $(x_1'' - x_1') > 0$ so that $(x_2'' - x_2')$ is a negative number. Then we see that $-\frac{p_1}{p_2}$ is the **quantity of good 2 that the consumer has to give up for each additional unit of good 1 he/she consumes.**
3. $\frac{p_1}{p_2}$ referred to as *the relative price* and we think of this as measuring the *opportunity cost* of consuming good 1 measured in good 2 units. Obviously we can do the same thing and measure opportunity cost of consuming good 2 in terms of good 1 units and this would be equivalent.

### 2.1.5 Changes in Income and Prices

**Experiment 1**: keep prices fix and increase income from $m$ to $m' > m$. Clearly, we compare

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2} x_1$$

with the line given by

$$x_2 = \frac{m'}{p_2} - \frac{p_1}{p_2} x_1,$$

so this causes a parallel shift as indicated by Figure 2.

**Experiment 2:** The next obvious experiment is to see what happens as one price goes up. Now, say that the price increases from $p_1$ to $p_1'$. Then we compare

$$x_2 = \frac{m}{p_2} - \frac{p_1}{p_2} x_1$$

with

$$x_2 = \frac{m}{p_2} - \frac{p_1'}{p_2} x_1,$$

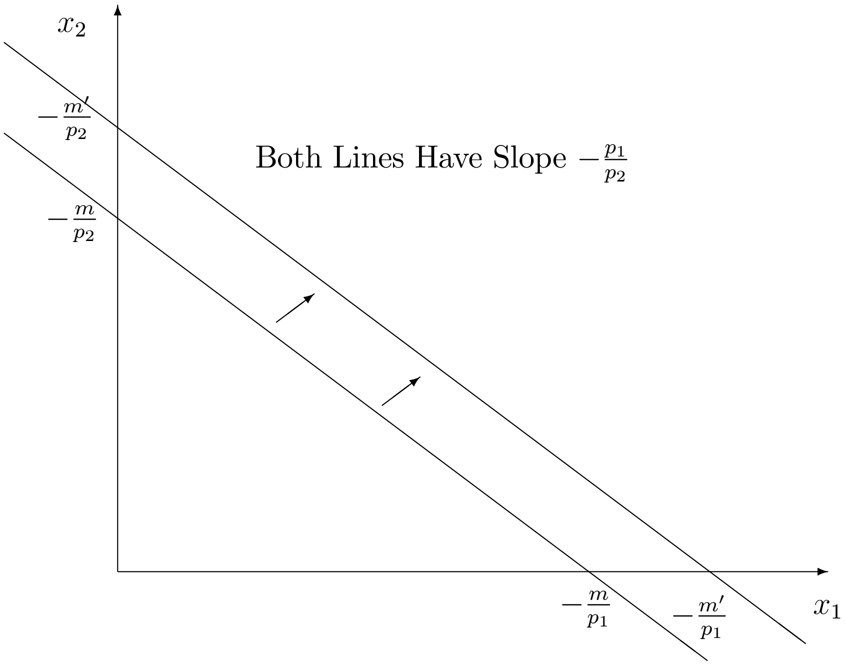

Figure 2: Increasing Income From $m$ to $m'$

where $\frac{p_1'}{p_2}$ is a larger number than $\frac{p_1}{p_2}$. Clearly, this means that the new line has the same intercept, but is steeper than the original line. The new intercept with the $x_1$-axis is $\frac{m}{p_1'} < \frac{m}{p_1}$. See Figure 3.

**Experiment 3**: what happens as $p_1$ and $p_2$ increases proportionally? Intuitively, if prices on all goods double this is just like decreasing income and we can see this by assuming that

$$\begin{aligned} p_1' &= tp_1 \\ p_2' &= tp_2, \end{aligned}$$

which means that the percentage change in the price is the same for both goods. Now, the new budget line is

$$\begin{aligned} m &= p_1'x_1 + p_2x_2' = tp_1x_1 + tp_2x_2 \Leftrightarrow \\ p_1x_1 + p_2x_2 &= \frac{m}{t} \end{aligned}$$

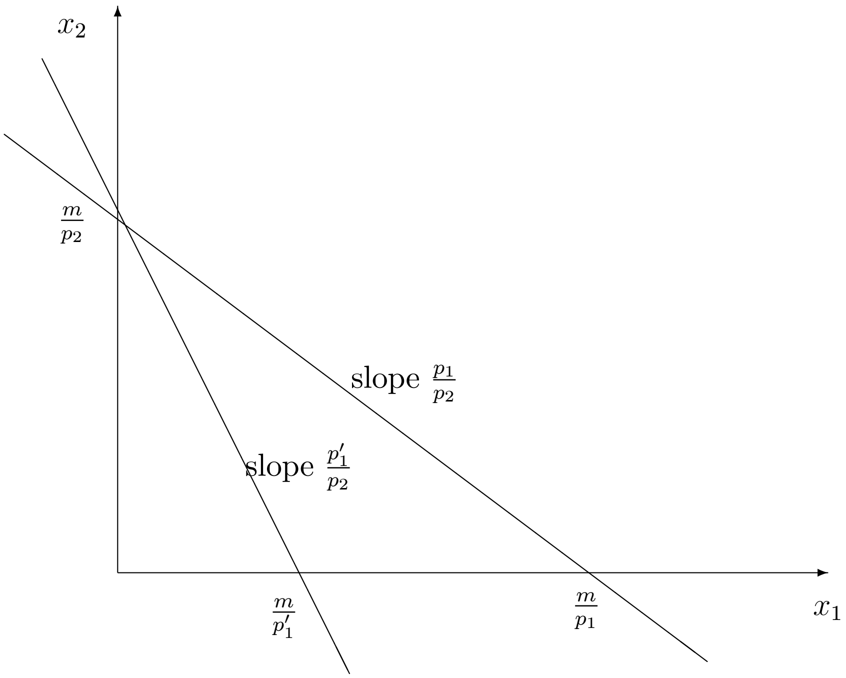

Figure 3: Increasing Price on good 1 from $p_1$ to $p_1'$

so if $t < 1$ (prices fall), then it is like an increase in income, and if $t > 1$, then it is like a decrease in income.

**Experiment 4**: increase (or decrease) $p_1, p_2$ and $m$ proportionally. I.e., compare

$$tp_1x_1 + tp_2x_2 = tm$$

with

$$p_1x_1 + p_2x_2 = m.$$

for some $t > 0$. Since both these equations describe the same budget line we conclude that the budget set does not change if prices and income change proportionally.

The interpretation of this is that multiplying all prices and income with the same factor $t$ is just changing the unit of account. That is, "completely balanced inflation" doesn't change anything else than the numbers on the bank notes.

### 2.1.6 More Exotic Budget Sets

In many real world situations, the relative prices between goods are not constant as in the simple budget constraints above. For example, utilities often use tariffs that generate quantity discounts and several government programs involve subsidies of the first $x$ units of particular goods (for example when dental coverage for, say, root canals is capped at some given amount). In this sort of situations, the budget constraints become more complex.

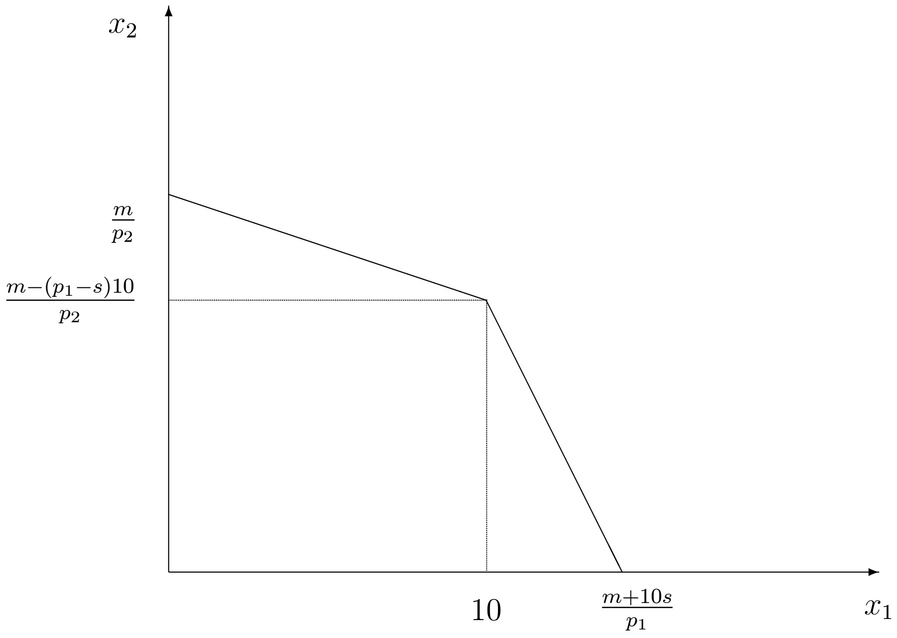

Figure 4: Kinked Budget Constraint due to Subsidy with maximum Usage

The simplest version of a budget constraint where relative prices are non-constant is often referred to as a “kinked budget constraint”. For example, imagine that two goods, $x_1$ and $x_2$ are sold by competitive firms at prices $p_1$ and $p_2$. However, say that the government thinks that good 1 is awfully important and decides to subsidize it. Specifically, for budget reasons (and maybe because the policy is mainly to help poor people consume good 1) the policy gives a per unit subsidy $s > 0$ for, say, the 10 first units of the good so that

- The per unit price for the first 10 units is $p_1 - s$
- The per unit price for any additional unit beyond 10 units is $p_1$

Assuming that $(p_1 - s)\,10 < m$ (that is: the consumer can afford to buy more than 10 units of the subsidized good) the budget set corresponding to this setup will look like in Figure 4. You should make sure you understand the derivation. One way to proceed is as follows:

1. If all resources are put on consumption of good 2, this means that the bundle consumed is $(x_1, x_2) = \left(0, \frac{m}{p_2}\right)$ just like in the standard case.
2. Next, the consumer consumes exactly 10 units of good 1 and the rest of the budget is used on good 2 we have that
$$(p_1 - s)\,10 + p_2 x_2 = m \Leftrightarrow x_2 = \frac{m - (p_1 - s)\,10}{p_2},$$
as is indicated in the graph.
3. Finally, if all money is spend on good 1 we have that
$$p_1 x_1 - s10 = m \Leftrightarrow x_1 = \frac{m + 10s}{p_1}$$
which gives the intercept on the $x_1-$axis.
4. The relative price when $x_1$ is between 0 and 10 is $-\frac{p_1 - s}{p_2}$ and $-\frac{p_1}{p_2}$ for $x_1 > 0$. The budget set is thus the area in between the kinked line in the graph and the two axes. You may note that if there had been no maximum quantity to enjoy the subsidy, then the intercept on the $x_1$ axis would have been $\frac{m}{p_1 - s}$. Some algebra shows that
$$\frac{m + 10s}{p_1} - \frac{m}{p_1 - s} = \frac{s\left[10\,(p_1 - s) - m\right]}{p_1\,(p_1 - s)}.$$
The point of this is that this should convince you that, as long as $m$ is larger than $10\,(p_1 - s)$, the slope gets steeper after the kink in the budget set (which of course also can be seen by comparing $-\frac{p_1 - s}{p_2}$ and $-\frac{p_1}{p_2}$).

## 2.2 Utility Functions and Indifference Curves

### 2.2.1 Utility Functions in the Model with two Goods

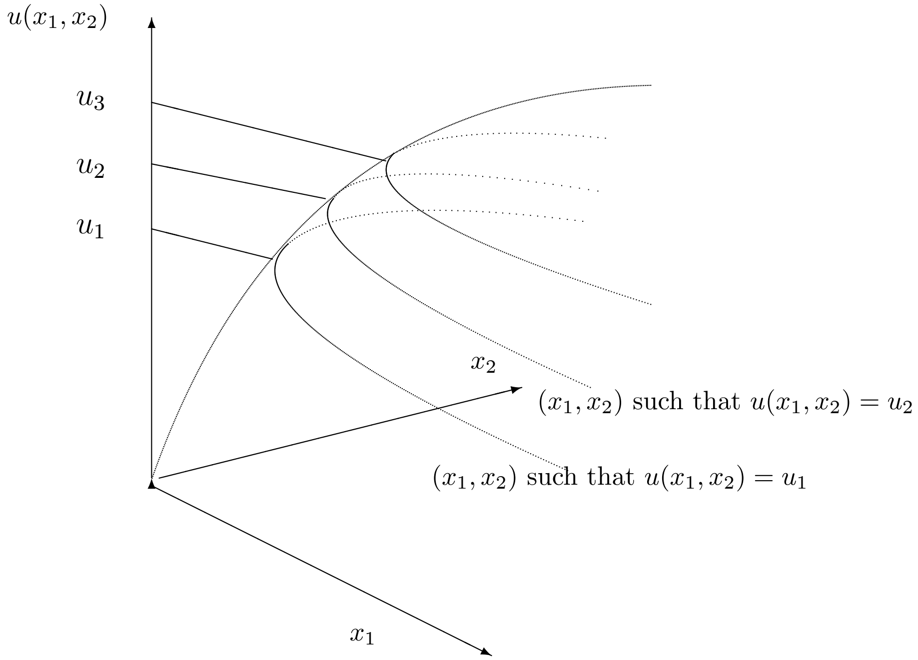

Figure 5: Utility Function with Two Arguments

A utility function assigns a number (the "happiness index") to each consumption bundle $(x_1, x_2)$.The advantage with the restriction to two goods is that we can visualize such a function in terms of a three-dimensional picture as in Figure 5.

The particular shape of the "mountain" depends on preferences and Figure 5 is drawn depicting the standard case (more is better and "convexity"). However, for now the important aspect is the construction of the pictures. The lines "on the mountain" is the picture can be thought of geometrically as the intersection between the "mountain" and a plane that is parallel to "the bottom" of the picture. That is, if we "cut a slice" through the "mountain" which is parallel to the bottom and at "height" $u_1$ from the bottom we find all combinations of $x_1$ and $x_2$ such that $u(x_1, x_2) = u_1$. We can obviously do the same for as many different utility levels as we like and the picture has three different levels of utility.

Now imagine yourself staring at the mountain straight from above so that what you see is

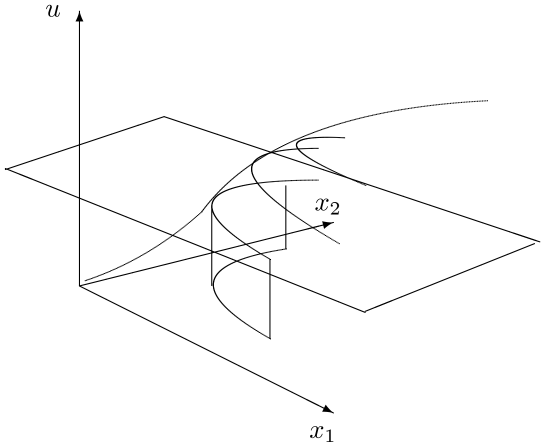

Figure 6: Projecting Utility Function to 2 Dimensions

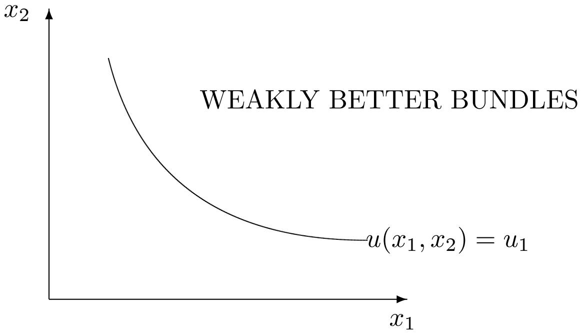

Figure 7: Indifference Curves are Projections of Level Curves

the two-dimensional surface with $x_1$ and $x_2$ on the axes (you'll have to have good binoculars and be far away). Also imagine that you've "sliced the mountain" with a plane parallel to the $x_1 - x_2$ plane so that you only see what is above the plane. Then, what you'd see is something looking like Figure 7.

Note that everything above the plane corresponds to bundles where the utility index is higher than $u_1$, so bundles to the north-east in the picture are bundles which are better than any bundle on the indifference curve. Obviously we can do the same thing for several choices of the utility index and then we generate the standard indifference curve diagram from your principles course.

### 2.2.2 Implications of Transitivity & Completeness

Rationality (transitivity & completeness) does not imply much about the shape of indifference curves: in Figure 8 are a few examples of "weird" preferences that are fully consistent with rational=complete and transitive preferences. We will typically rule out preferences that generate this type of indifference curves by making assumptions *in addition to* the basic assumption that the consumer is a rational decision maker.

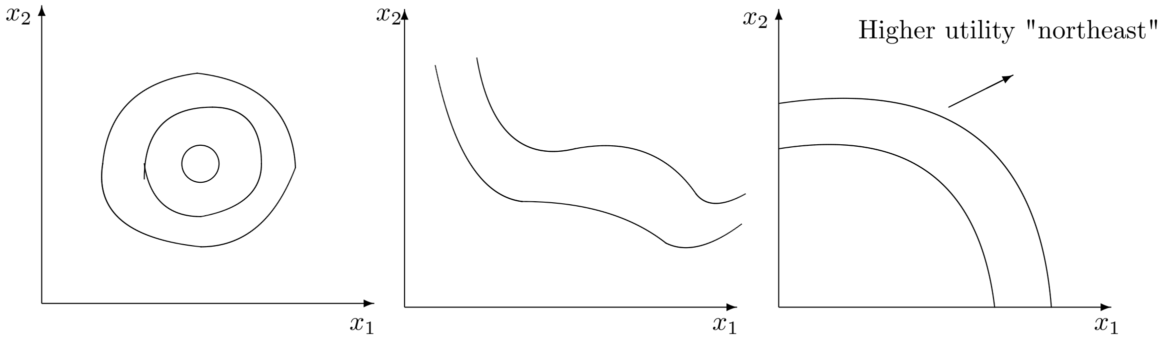

Figure 8: Examples of "Weird" Rational Preferences

The one and only restriction on indifference curves that come from rationality alone is that *indifference curves can not cross each other.* To see this, suppose they could, so that indifference curves could look as in Figure 9.

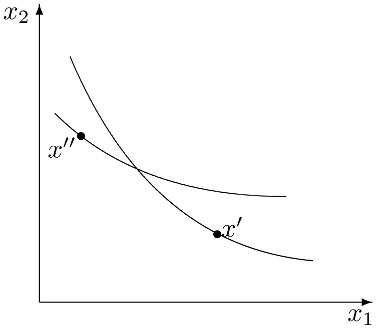

Figure 9: Indifference Curves can't Cross

If the crossing curves in Figure 9 are supposed to be distinct indifference curves, then

$$u\left(x_1', x_2'\right) \neq u\left(x_1'', x_2''\right).$$

But since they cross each other, there must be some point $x^*$ such that $x^*$ is on both indifference curves. Now:

1. $x^*$ being on the same indifference curve as $x'$ means that $u(x_1^*, x_2^*) = u(x_1', x_2')$ (recall that an indifference curve is *defined* as a set of bundles that give the same level of utility⇔bundles the consumer is equally happy with.
2. $x^*$ being on the same indifference curve as $x''$ means that $u(x_1^*, x_2^*) = u(x_1'', x_2'')$

Taken together this means that

$$u(x_1^*, x_2^*) = u(x_1', x_2') \neq u(x_1'', x_2'') = u(x_1^*, x_2^*),$$

which is an absurdity. Since the assumption that indifference curves *can* cross generates a conclusion that is false this is *proof* that indifference curves cannot cross (proof by contradiction).

## 2.3 Common Assumptions on Preferences

### 2.3.1 Monotonicity

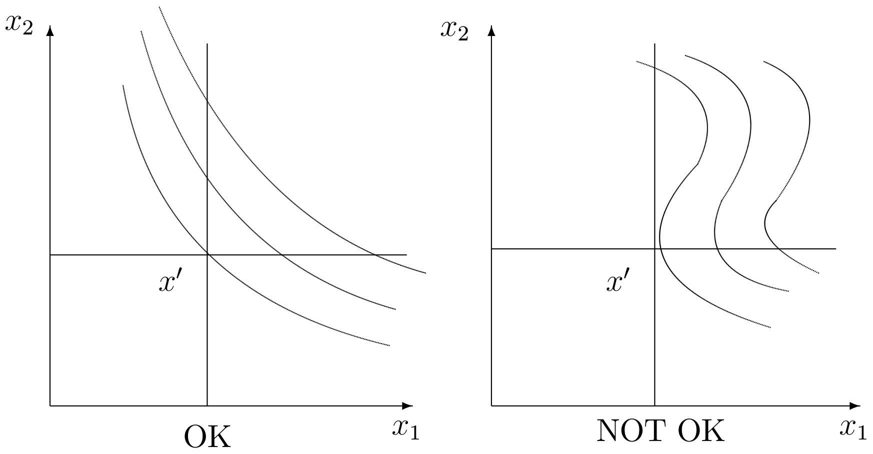

Figure 10: Monotonicity

This assumption simply says that the consumer always prefer more to less of any good given that the consumption of the other good is held fixed. In terms of the utility function, this means that we assume that $u\left(x_1, x_2\right)$ is strictly increasing in both arguments,

$$\begin{aligned} x_1' &< x_1'' \Rightarrow u(x_1', x_2') < u(x_1'', x_2') \\ x_2 &< x_2'' \Rightarrow u(x_1', x_2') < u(x_1', x_2'') \end{aligned}$$

Monotonicity can be checked easily in many cases by looking directly at the utility function or by taking the derivative if a quick glance isn't enough.

**Rationale for assumption:** Economics about how agents try to make the most out of limited resources. Natural in many applications. To see why this assumption makes sense it is useful to consider the "circular" indifference curves in Figure 8 and combine with a budget constraint. Figure 11 has combined two budget constraints with the indifference curves. Not that with the smaller budget, the point that puts the consumer on the highest indifference curve is at a tangency, while for the larger budget the best the consumer can do is to select the "bliss point". In the first case (with the smaller budget set) this is exactly as what would happen with monotonic preferences and you can note that as long as you are considering points that are to the "southeast" of the bliss point the preferences are *locally* monotonic. In the second case the consumer gets his or her most favorite bundle and (small) changes in prices or income will not matter. Now, what is interesting in economics is to study implications of agents that try to get the most out of *limited* resources and for that reason this second case is simply not that interesting.

### 2.3.2 Convex Preferences

Loosely speaking this assumption means that the consumer likes "mixing".

Graphical Criterion: Take two points on an indifference curve. Everything on a straight line in between should be at least as good (included in the set of better choices)

Saying the same thing in mathematical notation:

**Definition 2** *Preferences are convex if for any pair* $\left(x_1, x_2\right)$ *and* $\left(x_1', x_2'\right)$ *such that* $u\left(x_1, x_2\right) =$

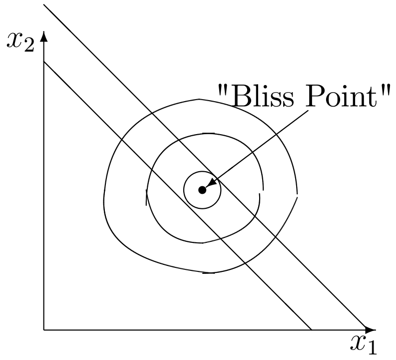

Figure 11: Non-Monotonic Preferences

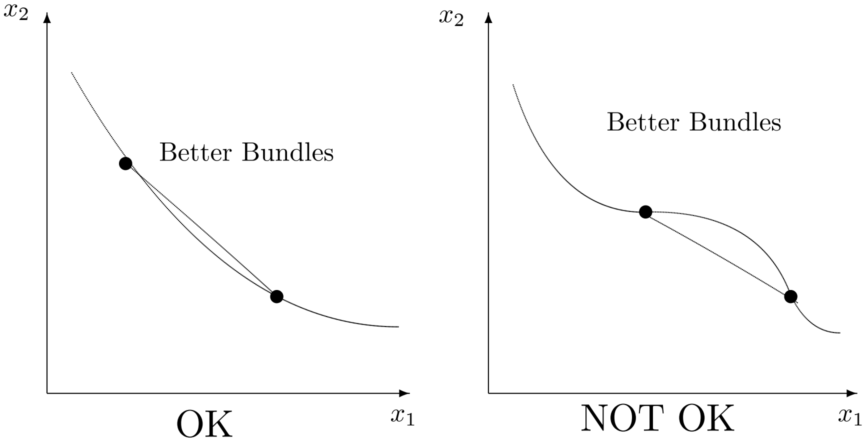

Figure 12: Convex and Non-Convex Preferences

$u(x_1', x_2')$ *and any* $\lambda$ *such that* $0 \leq \lambda \leq 1$

$$u\left(\lambda x_1 + (1-\lambda)\, x_1', \lambda x_2 + (1-\lambda)\, x_2'\right) \geq u\left(x_1, x_2\right) = u(x_1', x_2')$$

Note that this definition boils down to the graphical criterion for $x, x'$ such that $u\left(x\right) = u\left(x'\right)$. Actually, there is a plethora of different notions on convexity (and concavity) of functions. The definition above corresponds to one such notion (you don't want to know its name). However, when discussing choice under uncertainty later in the class we'll have to use a different notion of convexity/concavity. To avoid potential confusions, think of this as a definition of *convex preferences* and think of the notion of a *convex function* (which we will talk about later on) as something completely different.

Preferences are often called *strictly convex* if the inequality in the definition is strict whenever $0 < \lambda < 1$ (you'll get nothing but an equality for $\lambda = 0$ or $\lambda = 1$).

**Rationale for assumption:** This is an assumption that is mainly motivated by *tractability.* It makes some sense in some circumstances, in particular when we think about choice under uncertainty.

### 2.3.3 Example 1: "Perfect Substitutes"

Suppose the utility function is of the form

$$u(x_1, x_2) = ax_1 + bx_2,$$

where $a > 0$ and $b > 0$.

To construct indifference curves corresponding with this utility function we note that an indifference curve depicts

$$\text{All combinations of positive } (x_1, x_2) \text{ such that } u(x_1, x_2) = k$$

for some constant $k$. Thus, with these preferences an indifference curve are positive values of $(x_1, x_2)$ solving

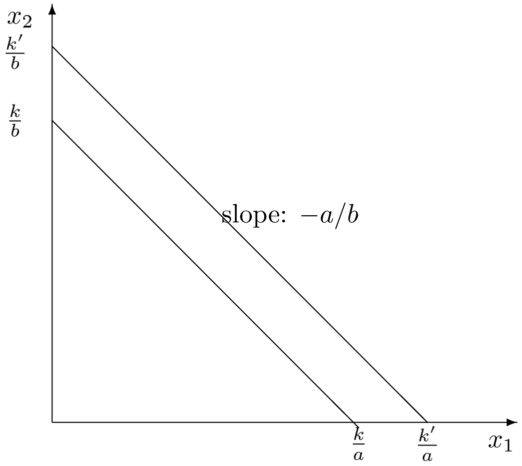

Figure 13: Indifference Curves for $u(x_1, x_2) = ax_1 + bx_2$

$$\begin{aligned} ax_1 + bx_2 &= k \Leftrightarrow \\ x_2 &= \frac{k}{b} - \frac{a}{b}x_1. \end{aligned}$$

These type of preferences are referred to as “perfect substitutes”. To see why it is useful to first consider the case with $a = b = 1$, which would mean that the consumer only would care about *the sum of the two goods.*

It may seem odd to some that we talk about perfect substitutes also when $a \neq b$. However, think of $x_1$ as number of 1\$ bills and $x_2$ as 5\$ or consider different sizes of packages.

### 2.3.4 Recall the Ordinal Nature of Preferences

Suppose instead that the utility function is

$$u(x_1, x_2) = c(ax_1 + bx_2),$$

where $c > 0 \Rightarrow$ Same system of indifference curves since an indifference curve now is given by

$$x_2 = \frac{k}{bc} - \frac{a}{b}x_1.$$

But $k$ is an arbitrary constant which can be varied from 0 to $\infty$ and for any given $k$ given the first utility function we get the same indifference curve by picking $k' = kc$ for the second utility function. Hence, the system of indifference curves is not affected by the multiplicative constant $c$.

Now let

$$u(x_1, x_2) = \sqrt{ax_1 + bx_2}$$

Then

$$\sqrt{ax_1 + bx_2} = k \Rightarrow ax_1 + bx_2 = k^2$$

and since $k^2$ is just some constant (which by appropriate choice of $k$ can be varied as freely as before), we again have the same system of indifference curves.

*The importance of this is that both the multiplication by a constant and changing the scale by taking the square root leaves the system of indifference curves unchanged. These are two EXAMPLES of the general principle that any change of the utility function that preserves the order leaves the indifference curves intact. Indeed, this is because the system of indifference curves has all information about the RANKING of bundles, so they fully describe preferences.*

### 2.3.5 General fact: Monotone Transformations Preserve Preferences

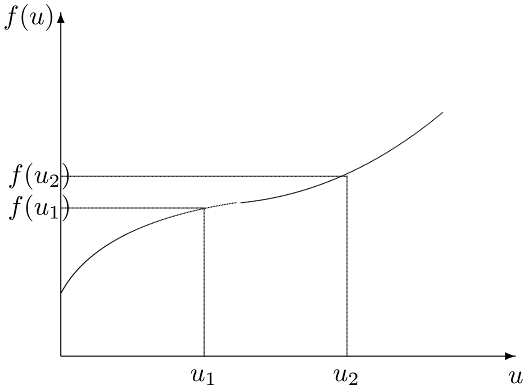

Figure 14: A Monotonically Increasing Function

Let $f(\cdot)$ be a strictly increasing function of a single variable (as the one in Figure 14 for example). Also, let $u(\cdot)$ be some given utility function and define the new utility function $\widetilde{u}(\cdot)$ as

$$\widetilde{u}(x_1, x_2) = f(u(x_1, x_2))$$

for every consumption bundle $(x_1, x_2)$. A utility function generated from $u$ in this way is called a *monotone transformation of* $u$.

**Claim** *$u(\cdot)$ and $\widetilde{u}(\cdot)$ represent the same preferences.*

The argument is maybe most easily understood from Figure 15 . The point is that if $u(x) > u(x')$ and $f$ is strictly increasing, then $f(u(x)) > f(u(x'))$, so the ranking between bundle $x$ and $x'$ is preserved (and this is true for any two bundles).

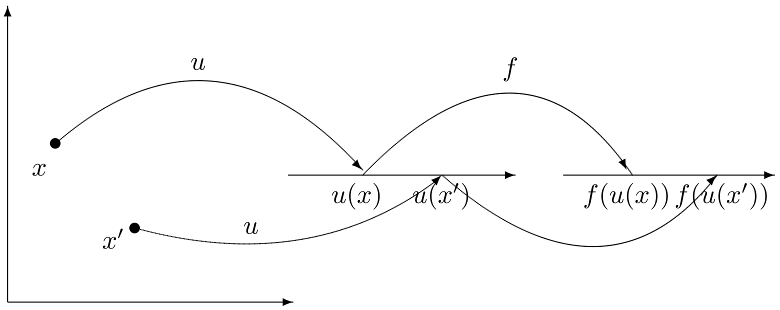

Figure 15: Transforming utility by Increasing Function Preserves Ranking

To actually prove the claim is easy once the notation is understood and it is clear what it means for $u$ to "represent preferences".

**Step 1** Note that $u$ represents $\succeq$ means that

$$u(x_1, x_2) \geq u(x_1', x_2') \Leftrightarrow (x_1, x_2) \succeq (x_1', x_2')$$

**Step 2** $f$ strictly increasing$\Rightarrow$

$$u(x_1, x_2) \geq u(x_1', x_2') \Leftrightarrow f(u(x_1, x_2)) \geq f(u(x_1', x_2'))$$

**Step 3** Combining this we get that

$$f(u(x_1, x_2)) \geq f(u(x_1', x_2')) \Leftrightarrow (x_1, x_2) \succeq (x_1', x_2'),$$

I.e., same preferences.

Notice that, for a given $k$, $u(x_1, x_2) = k$ and $f(u(x_1, x_2)) = k$ will in general correspond to two distinct indifference curves, whereas

$$u(x_1, x_2) = k \text{ and } f(u(x_1, x_2)) = f(k)$$

will generate the same indifference curve. In other words, the function $f$ changes the scale by which utility is measured, so $f$ needs to be applied to the original "utility index" in order to recapture the same indifference curve.

### 2.3.6 Example 2: Perfect Complements

Now, consider a utility function of the form

$$u(x_1, x_2) = \min\{ax_1, bx_2\},$$

where $a > 0$ and $b > 0$. To construct the indifference curves first draw the line of points $(x_1, x_2)$ such that

$$ax_1 = bx_2 \Leftrightarrow x_2 = \frac{a}{b}x_1,$$

as in Figure 16. Now consider a point on that line, i.e., some particular $(x_1', x_2')$ such that $ax_1' = bx_2'$. Now ask yourself, what happens to the value of the utility function if you increase good 1 and keep good 2 fix at $x_2'$. Then, for $x_1'' > x_1'$ we have that

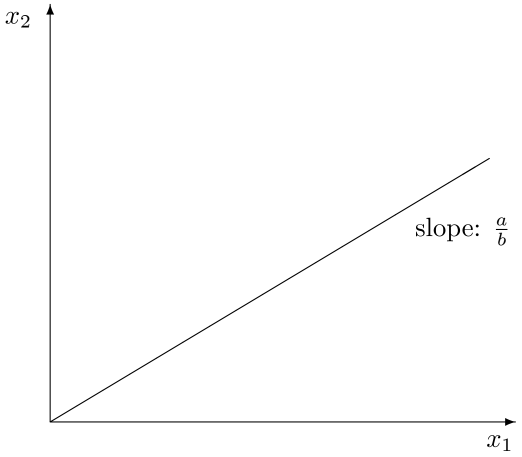

Figure 16: $(x_1, x_2)$ s.t. $ax_1 = bx_2$

$$\begin{aligned} ax_1'' &> ax_1' = bx_2' \Rightarrow \\ \min\{ax_1'', bx_2'\} &= bx_2' = ax_1' = \min\{ax_1', bx_2'\}, \end{aligned}$$

so the utility is unchanged. A symmetric argument shows that increasing good 2 to $x_2'' > x_2'$ also keeps the happiness constant, so the indifference curve through point $(x_1', x_2')$ can be depicted as in Figure 17. Now, by picking different points on the line with slope $\frac{a}{b}$ we generate a system of indifference curves as in the Figure where we have drawn 3 of the indifference curves.

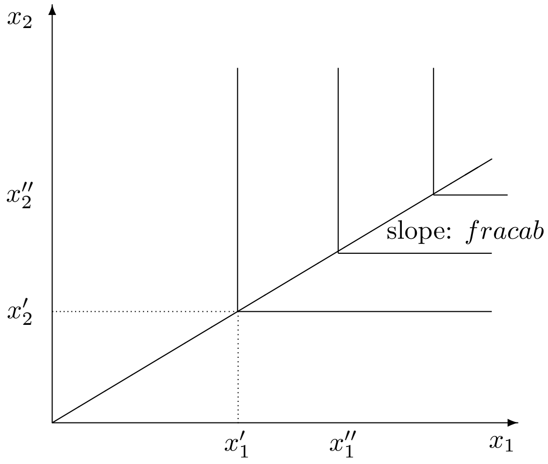

Figure 17: Indifference Curves for $u(x_1, x_2) = \min\{ax_1, bx_2\}$

This case is referred to as "perfect complements" since the consumer always want to keep consumption of the two goods in fixed proportion and you may think of these as natural preferences over left and right shoes, gin and dry vermouth and similar examples.

### 2.3.7 Example 3

Now let

$$u(x_1, x_2) = x_1 x_2$$

An indifference curve is then defined as the bundles $(x_1, x_2)$ satisfying

$$x_1 x_2 = k \Leftrightarrow x_2 = \frac{k}{x_1}.$$

To plot an indifference curve it is a good idea to start with some choice of $k$ that gives simple numbers, so let $k = 1$ and note that

$$\begin{aligned}
(1,1) \text{ satisfies } x_1 x_2 &= 1 \\
\left(2, \frac{1}{2}\right) \text{ satisfies } x_1 x_2 &= 1 \\
\left(\frac{1}{2}, 1\right) \text{ satisfies } x_1 x_2 &= 1 \\
\left(4, \frac{1}{4}\right) \text{ satisfies } x_1 x_2 &= 1 \\
\left(\frac{1}{4}, 4\right) \text{ satisfies } x_1 x_2 &= 1 \\
& \quad .... \\
\left(n, \frac{1}{n}\right) \text{ satisfies } x_1 x_2 &= 1
\end{aligned}$$

Once we have a few points and see the pattern it is kind of clear how to plot the indifference curve (one also realizes that multiplying all these bundles by the same constant we get a bunch of points that are all on the same higher indifference curve, so other indifference curves will have the same sort of shape). Observe, that the indifference curves will never reach the boundary (although they get closer and closer the further out an axis you go).

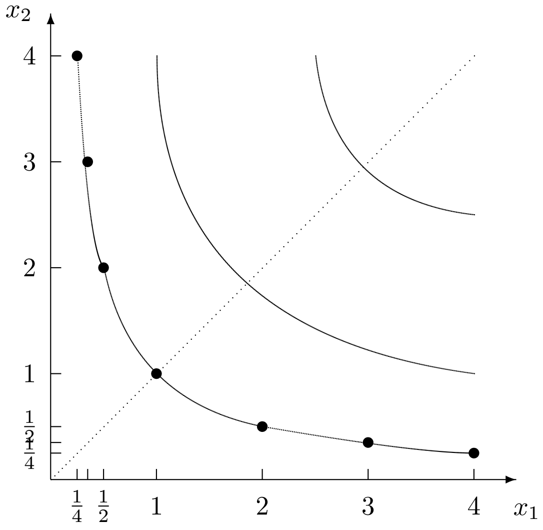

Figure 18: Indifference Curves for $u(x_1, x_2) = x_1 x_2$

This is the simplest example of a utility function representing preferences that satisfy:

1. Monotonic (this is satisfied by perfect substitutes & perfect complements as well)
2. (Strictly) Convex (perfect substitutes fails the *strict* convexity)
3. Smooth (perfect complements fails this due to kinks)

### 2.3.8 Example 4: Quasi-linear Preferences

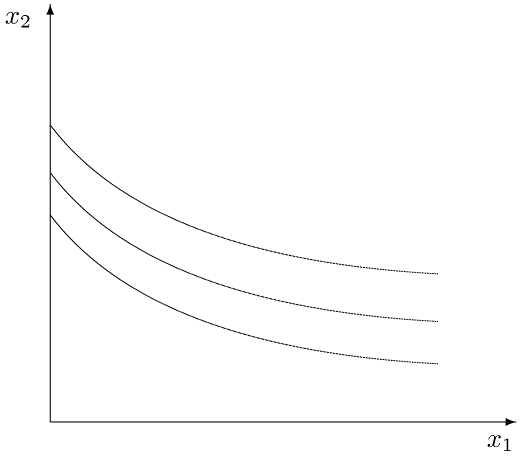

Figure 19: Indifference Curves for $u(x_1, x_2) = v(x_1) + x_2$

Now suppose that

$$u(x_1, x_2) = v(x_1) + x_2$$

An indifference curve is then given by solutions to

$$k = v(x_1) + x_2,$$

so the indifference curve will just shift vertically when changing $k$. When using quasi linear preferences we will typically let $v(x_1) = \ln x_1$ or $v(x_1) = \sqrt{x_1}$ and both of these specifications generate indifference curves similar to the ones in Figure 19.

### 2.3.9 Example 5: Cobb Douglas Preferences

Looking at applied work one may be led to believe that this is *the* utility function-this specification is by far the most popular and used in both applied theory and empirical work. These preferences are represented by

$$u\left(x_1, x_2\right) = x_1^{\alpha} x_2^{b}$$

for $a > 0, b > 0$. We've already depicted the special case with $a = b = 1$ and other choices of $a$ and $b$ produce similar pictures, but "skewed" rather than symmetric.