Expected utility

There is a fundamental difference between a cup of coffee and *maybe* having a cup of coffee; that is to say, there is an important distinction between being given a cup of coffee and someone telling you, "If this coin lands heads-up, I'll give you a cup of coffee; otherwise, I'll give you nothing." While this may seem obvious, it gets at something deeper that we've been ignoring until this point. We have been considering the value for explicit objects — the utility from receiving an apple, or fifty cents and two t-shirts — but we've been ignoring a fundamental feature of the real world: nothing is certain!

It is, of course, a bit of a stretch to look at a world where there is literally *no* determinism, but uncertainty is a concept which has great economic implications.[1] Although it is possible to model uncertainty as given in the problem sets, a basic understanding is not too difficult and will make the intuition behind many models much clearer. If you have taken Econ 41, none of this will be news to you and you may as well skip to the next section.

Probability

Underpinning our mathematical notion of randomness is *probability*: a probability is a number between 0 and 1 that describes the likelihood of an event occurring; an event with probability 1 will happen for sure, and an event with probability 0 will not happen.[2] We can translate a percentage into a probability by dividing it by 100. The standard example is that a fair coin will land heads-up 50% of the time; in probability, we say that the probability that it lands heads-up is $\frac{1}{2}$. This would be denoted

$$\mathrm{P}\left(\text{heads-up}\right) = \frac{1}{2}.$$

An important consequence of this view of the world is that since, for sure, something either happens or doesn't happen, the probability that something happens plus the probability that it doesn't happen must equal 1. In the coin example, this is

$$\mathrm{P}\left(\text{heads-up}\right) + \mathrm{P}\left(\text{not heads-up}\right) = \mathrm{P}\left(\text{heads-up}\right) + \mathrm{P}\left(\text{tails-up}\right) = 1.$$

As such, when talk about the probability of something happening, generally denoted by π or α, we know without stating that the probability that is *does not* happen is $1 - \pi$ or $1 - \alpha$, respectively. That is to say, if I know that there is a 30% chance of rain today, I also know that there is a $1 - 30\% = 70\%$ chance of no rain today.

A fairly deep question is, "What is a probability?" Since, in the real world, things either happen or do not[3] it can be off-putting to think about a world in which things *maybe* happen. But consider this: before something either happens or does not, all you can say is that it *might* happen or might not. A probability can be thought to represent the concept that, if we watch a particular setting enough times, the number of times the event we're interested in happens out of the total number of times we watched should roughly equal the probability that it happens. That is, if you flip a coin sufficiently many times it will have landed heads-up roughly 1 time out of 2, for a probability of $\frac{1}{2}$; if the forecast for today is a 30% chance of rain, then if I repeat *today* enough times[4] roughly 30 times out of 100 it should rain, for a probability of $\frac{30}{100} = 30\%$.

[1] For your weekly TMI, right now I am sitting in my parents' kitchen in Pittsburgh, PA. I will be here for three days, and the weather forecast has variously called for 90°F, rain, sunshine, and snow. The optimal choice of what to put in my suitcase was a function of how I evaluated my utility for having long sleeves in muggy weather versus having short sleeves in snow.

[2] There is an important caveat here (Econ 41 students: think about the difference between a PMF and a PDF), however the intuition is on track.

[3] Unless, of course, you are versed in quantum mechanics; economics doesn't particularly care about this.

[4] A la Bill Murray in *Groundhog Day*.

Expectation

Probability gives us the tools for describing randomness in the world, but it does not give us a direct method of *evaluating* this randomness. Consider, for example, the following three options:

(a) I pay you a dollar if a coin lands heads-up, and nothing otherwise.

(b) I pay you a dollar if *two* coins land heads-up, and nothing otherwise.

(c) I pay you *two* dollars if a coin lands heads-up, and nothing otherwise.

Without resorting to math, we should all intuitively know that option (c) is the best for us, provided you don't care about how much money I lose. To understand choice in the real world, we need to quantify *why* what gives rise to this intuition.

A *random variable* X is a real number which randomly takes values in its *support*, S. Suppose that I pay you a dollar if a coin lands heads-up (with probability $\frac{1}{2}$) and zero otherwise — situation (a) above; the support of the amount I pay you (X) is $\{0, 1\}$, and the probability that the payment takes either value is $\mathrm{P}(X = 0) = \frac{1}{2} = \mathrm{P}(X = 1)$.

The *expectation* of a random variable is the sum over its support of all its values times the probability that they occur,

$$\mathbb{E}[X] = \sum_{x \in S} x\mathrm{P}(X = x).$$

In the coin example, this is

$$\mathbb{E}[X] = (0)\mathrm{P}(X = 0) + (1)\mathrm{P}(X = 1) = \mathrm{P}(X = 1) = \frac{1}{2};$$

changing to situation (c) above, the expectation is

$$\mathbb{E}[X] = (0)\mathrm{P}(X = 0) + (2)\mathrm{P}(X = 1) = 2\mathrm{P}(X = 1) = 1.$$

Roughly speaking, the expectation corresponds to the average outcome of the random variable if it is observed enough times.

Expected utility

The basis for our understanding of choice under uncertainty is the notion of *expected utility*, defined in terms of expectation (hence the name); alternatively, it is referred to as von Neumann-Morgenstern utility after its progenitors. In particular, given a utility function u we assume that the utility for a random outcome Y is

$$U(Y) = \mathbb{E}[u(Y)].$$

Suppose that I am given a $\frac{1}{4}$ chance at winning a Porsche, and a $\frac{3}{4}$ chance of winning a Geo. My expected utility for this set of chances[5] is

$$\mathbb{E}[u(Y)] = \frac{1}{4}u(\text{Porsche}) + \frac{3}{4}u(\text{Geo}).$$

If we specified my explicit utility over makes of cars, we could then evaluate how much I like this gamble.

The expected utility hypothesis puts a simple framework on the evaluation of nondeterministic outcomes. Once we've assumed this form, all the other rules of utility apply; in particular, a random outcome which

[5] We call these "lotteries."

yields more *expected* utility is better than one that yields less, regardless of its particular outcomes (its particular support).

It should be mentioned that there are *many* experimental contradictions of the expected utility hypothesis; however, as analytical tools go it is indispensible.

Insurance

A common setting with uncertainty involves insurance against risk. Although modeling an insurance market is extremely difficult and the underlying risks can be arbitrarily complex — think about your health insurance policy, if you've been so brave as to read through any part of it — we can consider a far simpler world in which there is a single risk. We often consider a world with states "high" (H) and "low" (L), in which the high state of the world indicates some good outcome and the low state of the world indicates some bad outcome.

To make this concrete, suppose that in the high state of the world you have \$100 of income and in the low state of the world you have \$0 of income. Clearly, to avoid starving in the event that the low state of the world arises you'd like to have some mechanism for shifting income from the high state to the low state. We consider insurance as providing this mechanism. The real-world analogy is a policy where, if nothing happens, you pay your monthly premium but, if something [terrible] happens, the money your policy pays out more than outweighs the premium you pay in.

Actuarially-fair pricing

The first issue that arises is the cost of insurance; in particular, while we often consider the prices for iPods and other commodities to be fairly arbitrary, there is a way in which insurance prices are derived from the same underlying processes no matter which insurer you choose. This concept defines the notion of *actuarially-fair* prices.

Suppose that a state of the world happens with probability π, and that an insurance policy will pay out c in the event of this state of the world and 0 otherwise. The actuarially-fair price for this policy is πc. The root of this is in expected utility: the firm will pay out c with probability π and 0 with probability $1 - \pi$. When we represent the payout R as an expectation, we have

$$\mathbb{E}[R] = (c)\pi + (0)(1 - \pi) = c\pi.$$

If the insurer charges less than $c\pi$, it will lose money in the long run; if it charges more, there are positive profits and there is reason for another insurer to enter. Intuitively then the only long-run sustainable price for this policy is $c\pi$.

As an example, consider a model where there are three states of the world, H, M, and L. In state H, the agent receives \$100 of income; in state M, the agent receives \$50 of income; and in state L, the agent receives \$0 of income. State H occurs with probability $\frac{1}{6}$; state M occurs with probability $\frac{1}{2}$; and state L occurs with probability $\frac{1}{3}$. There are three insurance policies: policy A pays out 1 unit of income in state M and 0 otherwise; policy B pays out 1 unit of income in state L and 0 otherwise; and policy C is a combination of policies A and B. What are the actuarially-fair prices for these policies?

For policy A, we compute the expectation of its payout as

$$\mathbb{E}[A] = (0)\frac{1}{6} + (1)\frac{1}{2} + (0)\frac{1}{3} = \frac{1}{2}.$$

For policy B, we compute the expectation of its payout as

$$\mathbb{E}[B] = (0)\frac{1}{6} + (0)\frac{1}{2} + (1)\frac{1}{3} = \frac{1}{3}.$$

It follows that the actuarially-fair prices of A and B are $p_A = \frac{1}{2}$ and $p_B = \frac{1}{3}$, respectively.

How then can we determine the price of policy C? The same logic of *there should be no profit available* tells us that the price of policy C should equal the sum of the prices of policies A and B; in particular,

$$p_C = p_A + p_B = \frac{1}{2} + \frac{1}{3} = \frac{5}{6}.$$

We can also consider this in another way: a policy which is policies A and B together will pay out 1 in states M and L, and 0 in state H. To directly compute the actuarially-fair price of policy C, we calculate

$$\mathbb{E}[C] = (0)\frac{1}{6} + (1)\frac{1}{2} + (1)\frac{1}{3} = \frac{5}{6}.$$

Example

Solving insurance questions is a matter of solving the implicit constrained optimization problem. We will set up and solve such a problem below.

Three Little Pigs

A pig would like to build a house to protect itself from the elements, and to keep warm at night. There are three materials available for construction of a house: straw, wood, and bricks. Building a house from straw costs $p_s = 1$ per square foot; building a house from wood costs $p_w = 2$ per square foot; building a house from brick costs $p_b = 3$ per square foot. Each construction material protects the pig equally well from the ravages of nature, and his utility is defined only by the number of square feet x of his house, $u(x) = \ln x$. Suppose that the pig has wealth today of $w = 200$.

The pig will build his house and live in it for one day; in the middle of the night, the Big Bad Wolf will arrive and blow his house down with probability $\pi < 1$. If the house is made of straw, it is destroyed; if the house is made of wood, $\frac{1}{4}$ of its square footage survives; if the house is made of brick, there is no damage.

Fortunately for the pig, there is an insurer which will mitigate the risk of damage. The insurer prices its policies at an actuarially-fair rate, so the price of a unit of insurance to rebuild one square foot of a straw house is $\pi p_s = \pi$; the price of a unit of insurance to rebuild one square foot of a wood house is $\pi p_w = 2\pi$; the price of a unit of insurance to rebuild one square foot of a brick house is $\pi p_b = 3\pi$.

Let x denote the number of square feet of the pig's original house, and let y denote the number of units of insurance bought. The pig's overall utility is given by his utility for square footage today plus his expected utility for square footage tomorrow, after the Big Bad Wolf has *maybe* blown down his house and the insurer has repaired what it is obligated to. In particular, the utility for a straw house is

$$u_s(x, y) = \ln x + [\pi \ln y + (1 - \pi) \ln x];$$

the utility for a wooden house is

$$u_w(x, y) = \ln x + \left[\pi \ln \left(\frac{1}{4}x + y\right) + (1 - \pi) \ln x\right];$$

the utility for a brick house is

$$u_b(x, y) = \ln x + [\pi \ln(x + y) + (1 - \pi) \ln x].$$

Notice that utility is structured as today's utility, plus tomorrow utility in the event that there is damage done to the house times the probability of this happening (π), plus tomorrow's utility in the event that there is no damage done to the house times the probability of this happening $(1-\pi)$. Tomorrow's utility in the event of damage to the house is given by the amount of square footage which survives plus the amount of square footage the insurer will repair.

Question: *for each building material, what is the optimal amount of square footage to build today, and what is the optimal amount of insurance to buy?*

Solution: notice that each question has a setup of

$$u_t(x,y) = \ln x + \pi \ln(k_t x + y) + (1-\pi)\ln x = (2-\pi)\ln x + \pi \ln(k_t x + y),$$

where k_t represents the fraction of square footage which survives, given the the building material is type t. Further, given that the price of building material t is p_t and the price of insuring a square foot of this building material is πp_t, the pig's budget constraint is given by

$$p_t x + \pi p_t y \leq w \quad \rightsquigarrow \quad p_t x + \pi p_t y = 200.$$

We can then setup the pig's optimization problem as

$$\max_{x,y} (2-\pi)\ln x + \pi \ln(k_t x + y) \text{ s.t. } p_t x + \pi p_t y = 200.$$

Marginal utilities are given by

$$\begin{aligned} \text{MU}_x &= \frac{2-\pi}{x} + \frac{\pi k_t}{k_t x + y} \\ &= \frac{2k_t x + (2-\pi)y}{x(k_t x + y)}, \\ \text{MU}_y &= \frac{\pi}{k_t x + y}. \end{aligned}$$

The price ratio is given by

$$\frac{p_x}{p_y} = \frac{p_t}{\pi p_t} = \frac{1}{\pi}.$$

At the optimum, we should then have

$$\begin{aligned} && \frac{\text{MU}_x}{\text{MU}_y} &= \frac{p_x}{p_y} \\ \iff && \frac{\frac{2k_t x + (2-\pi)y}{x(k_t x + y)}}{\frac{\pi}{k_t x + y}} &= \frac{1}{\pi} \\ \iff && \frac{2k_t x + (2-\pi)y}{x} &= 1 \\ \iff && 2k_t x + (2-\pi)y &= x \\ \iff && (2k_t - 1)x + (2-\pi)y &= 0 \\ \iff && x &= \left(\frac{2-\pi}{1-2k_t}\right) y. \end{aligned}$$

Substituting into the budget constraint, we find

$$\begin{aligned} && p_t x + \pi p_t y &= 200 \\ \iff && \left(\frac{2-\pi}{1-2k_t}\right) y + \pi y &= \frac{200}{p_t} \\ \iff && \left(\frac{2-2\pi k_t}{1-2k_t}\right) y &= \frac{200}{p_t} \\ \iff && y &= \frac{200}{p_t}\left(\frac{1-2k_t}{2-2\pi k_t}\right). \end{aligned}$$

For a straw house, we know $p_s = 1$ and $k_s = 0$; the optimal amount of insurance is then

$$y_s^* = 200\left(\frac{1}{2}\right) = 100;$$

returning to the budget constraint, we know

$$x_t = \frac{200}{p_t} - \pi y_t \quad \Longrightarrow \quad x_s^* = (2-\pi)100.$$

For a wooden house, we know $p_w = 2$ and $k_s = \frac{1}{4}$; the optimal amount of insurance is then

$$y_w^* = \frac{200}{2}\left(\frac{1-\frac{1}{2}}{2-\frac{1}{2}\pi}\right) = \frac{100}{4-\pi};$$

returning to the budget constraint, we find

$$x_w^* = \frac{200}{2} - \pi\left(\frac{100}{4-\pi}\right) = \frac{400-200\pi}{4-\pi}.$$

For a brick house, we know $p_b = 3$ and $k_s = 1$; the optimal amount of insurance is then

$$y_b = \frac{200}{3}\left(\frac{1-2}{2-2\pi}\right).$$

Since $\pi < 1$, we know $2 - 2\pi > 0$; with $1 - 2 = -1 < 0$, it follows that $y_b < 0$. As consumption should be positive, we know that this indicates a corner solution. Hence the pig will spend all his income on bricks and none on insurance, so $x_b^* = \frac{200}{3}$ and $y_b^* = 0$. □

Question: *what is the pig's optimal utility from each building material?*

Solution: this is a matter of substituting optimal levels of x and y into the appropriate utility functions. We can see

$$\begin{aligned} u_s(x_s^*, y_s^*) &= (2-\pi)\ln x_s^* + \pi \ln y_s^* \\ &= (2-\pi)\ln[(2-\pi)100] + \pi \ln 100 \\ &= (2-\pi)\ln(2-\pi) + (2-\pi)\ln 100 + \pi \ln 100 \\ &= 2\ln 100 + (2-\pi)\ln(2-\pi); \end{aligned}$$

$$\begin{aligned} u_w(x_w^*, y_w^*) &= (2-\pi)\ln x_w^* + \pi \ln\left[\frac{1}{4}x_w^* + y_w^*\right] \\ &= (2-\pi)\ln\left[\frac{400-200\pi}{4-\pi}\right] + \pi \ln\left[\frac{100-50\pi}{4-\pi} + \frac{100}{4-\pi}\right] \\ &= (2-\pi)\ln\left[\frac{400-200\pi}{4-\pi}\right] + \pi \ln\left[\frac{200-50\pi}{4-\pi}\right] \\ &= (2-\pi)\ln\left[\frac{400-200\pi}{4-\pi}\right] + \pi \ln 50 \\ &= (2-\pi)\ln 50 + (2-\pi)\ln(8-4\pi) - (2-\pi)\ln(4-\pi) + \pi \ln 50 \\ &= 2\ln 50 + (2-\pi)\left[\ln(8-4\pi) - \ln(4-\pi)\right]; \end{aligned}$$

$$\begin{aligned} u_b(x_b^*, y_b^*) &= (2-\pi)\ln x_b^* + \pi \ln(x_b^* + y_b^*) \\ &= (2-\pi)\ln\frac{200}{3} + \pi \ln\frac{200}{3} \\ &= 2\ln 100 + 2\ln\frac{2}{3}. \end{aligned}$$

□

Question: *which building material will the pig choose? Does it depend on the probability that the Big Bad Wolf arrives?*

Solution: we check which building material yields the highest level of utility. Firstly, we ask if straw is better than wood.

$$
\begin{aligned}
& u_s(x_s^*, y_s^*) \overset{?}{>} u_w(x_w^*, y_w^*) \\
\iff\quad & 2\ln 100 + (2-\pi)\ln(2-\pi) \overset{?}{>} 2\ln 50 + (2-\pi)[\ln(8-4\pi) - \ln(4-\pi)] \\
\iff\quad & 2\ln 2 \overset{?}{>} (2-\pi)[\ln(8-4\pi) - \ln(4-\pi) - \ln(2-\pi)] \\
\iff\quad & \ln 4 \overset{?}{>} (2-\pi)[\ln 4 - \ln(4-\pi)] \\
\iff\quad & (2-\pi)\ln(4-\pi) \overset{?}{>} (1-\pi)\ln 4 \\
\iff\quad & (4-\pi)^{2-\pi} \overset{?}{>} 4^{1-\pi} \\
\iff\quad & 4 \overset{?}{>} \left(\frac{4}{4-\pi}\right)^{2-\pi}.
\end{aligned}
$$

Since $\pi < 1$, we know that the "interior" of the right-hand side is less than $\frac{4}{3}$, while its exponent is less than 2; hence the right-hand side is less than $(\frac{4}{3})^2 = \frac{16}{9} < 4$. Thus the right-hand side is less than the left-hand side and we that the optimal utility from straw is above the optimal utility from wood, so the building with straw is superior to building with wood.

Now we ask whether straw is better than brick.

$$
\begin{aligned}
& u_s(x_s^*, y_s^*) \overset{?}{>} u_b(x_b^*, y_b^*) \\
\iff\quad & 2\ln 100 + (2-\pi)\ln(2-\pi) \overset{?}{>} 2\ln 100 + 2\ln\frac{2}{3} \\
\iff\quad & (2-\pi)\ln(2-\pi) \overset{?}{>} \ln\frac{4}{9} \\
\iff\quad & (2-\pi)^{2-\pi} \overset{?}{>} \frac{4}{9}.
\end{aligned}
$$

Since $2 - \pi > 1$ we know that the left-hand side is greater than 1. Since the right-hand side is less than 1, it follows that there is more utility from building with straw then from building with brick.

The pig's optimal choice is then to build from straw. This is independent of the probability that the Big Bad Wolf arrives and blows the pig's house down. □

Demand

We are by now familiar with finding the consumer's optimal demand. If this was all there was to economics, we would be in a very good place. Further generalization aside, one useful exercise is to ask how changes in the underlying market structure — i.e., changes in prices or agent wealth — alter the agent's behavior. This we capture by solving the consumer's optimization problem in terms of arbitrary wealth w or prices p, then watching consumption change with the variable we are interested in.

Although we will not have corner solutions in the example below, you need to be careful about corner solutions when solving for consumption with arbitrary wealth and prices. While it's easy to see whether a particular number is positive or negative, it is generally less simple to see whether an algebraic expression is positive or negative. When solving these problems and the Inada conditions (see last week's notes) are not satisfied, be sure to check for corner consumption.

Consider a two-commodity system with CES agent utility, $u(x, y) = -(x^{-1} + y^{-1})$. The agent has wealth w and prices are $p = (p_x, p_y)$; her maximization problem is given by

$$\max_{x,y} -(x^{-1} + y^{-1}) \text{ s.t. } p_x x + p_y y = w.$$

Marginal utilities are then

$$\begin{aligned} \mathrm{MU}_x &= x^{-2}, \\ \mathrm{MU}_y &= y^{-2}. \end{aligned}$$

Equating marginal utilities per cost, we have

$$\frac{x^{-2}}{p_x} = \frac{y^{-2}}{p_y} \quad \Longrightarrow \quad y = x\sqrt{\frac{p_x}{p_y}}.$$

Returning to the budget constraint, we find

$$\begin{aligned} && w &= p_x x + p_y y \\ && &= p_x x + p_y x\sqrt{\frac{p_x}{p_y}} \\ && &= x\left(p_x + \sqrt{p_x p_y}\right) \\ \Longleftrightarrow && x^* &= \frac{w}{p_x + \sqrt{p_x p_y}} \\ \Longrightarrow && y^* &= \frac{w}{p_y} - \frac{x^* p_x}{p_y} \\ && &= \frac{w}{p_y} - \frac{w p_x}{p_y(p_x + \sqrt{p_x p_y})} \\ && &= \frac{w\sqrt{p_x p_y}}{p_y p_x + p_y\sqrt{p_x p_y}} \\ \Longleftrightarrow && y^* &= \frac{w}{p_y + \sqrt{p_x p_y}}. \end{aligned}$$

We will use these equations throughout the following sections.

Engel curves

Engel curves describe how consumption changes with wealth; as is custom in economics, we place w along the y-axis and consumption along the x-axis.[6] A good with a positive slope — consumption increasing with wealth — is called *normal* and a good with a negative slope — consumption decreasing with wealth — is called *inferior*.

To do this, we will of course need to know — or assume — the prices of goods x and y. In general a question will supply these parameters, so here we'll assume $p_x = p_y = 1$. Then the equations for optimal consumption in terms of wealth are

$$x^* = \frac{w}{2}, \quad y^* = \frac{w}{2}.$$

The corresponding Engel curves are shown in Figure 1.

[6] For what it's worth, it may help to remember this as being the opposite of what you think it should be.

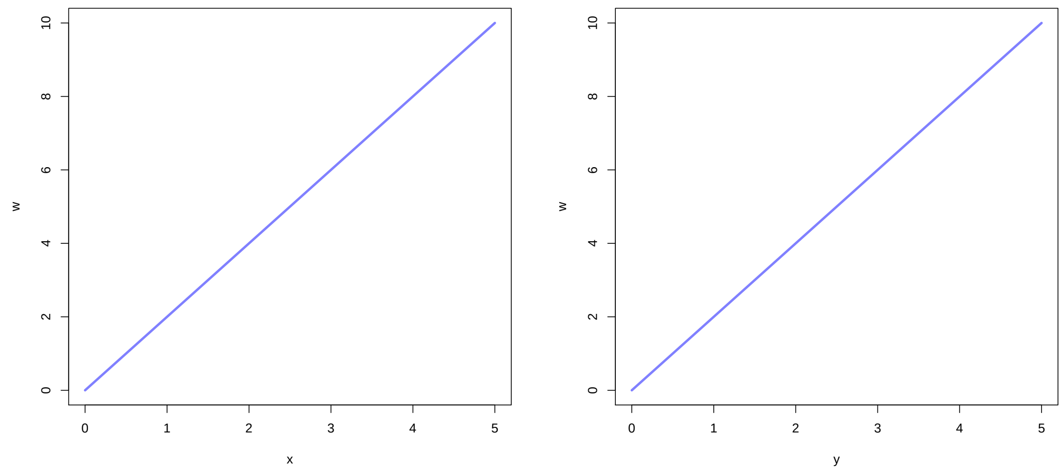

Figure 1: Engel curves for goods x (left) and y (right) when $(p_x, p_y) = (1, 1)$.

Demand curves

Demand curves describe how consumption changes with prices. From Econ 1, you are familiar with them describing the behavior of the economy as a whole; we are now looking at the same concept, but for a single agent. In general, consumption of a good should *decrease* as its price increases; a good in which the opposite happens — its consumption *increases* when its price increases — is referred to as a *Giffen good.* There is an open debate in economics as to whether Giffen goods exist in the real world, but microeconomic theory provides no particular reason that they are impossible.

As with Engel curves, we will need to fix some parameters of the problem to obtain results. We will assume $w = 1$, so demand is given by

$$x^* = \frac{1}{p_x + \sqrt{p_x p_y}}, \quad y^* = \frac{1}{p_y + \sqrt{p_x p_y}}.$$

Let's now consider two cases. In the first, the price of x changes while the price of y is held constant at $p_y = 1$. Demand is then given by

$$x^* = \frac{1}{p_x + \sqrt{p_x}}, \quad y^* = \frac{1}{1 + \sqrt{p_x}}.$$

The corresponding demand curves are shown in Figure 2.

In the second case, the price of y changes while the price of x is held constant at $p_x = 1$. Demand is then given by

$$x^* = \frac{1}{1 + \sqrt{p_y}}, \quad y^* = \frac{1}{p_y + \sqrt{p_y}}.$$

The corresponding demand curves are shown in Figure 3.

Plotting joint consumption

We can view the case of the Engel curve (and, of course, the demand curve) as describing a set of points $(x^*(w), y^*(w))$ which are traced out with changes in wealth.[7] We may be interested in watching as *overall*

[7] This structure is known as *parametric* and may bring back bad memories from high school.

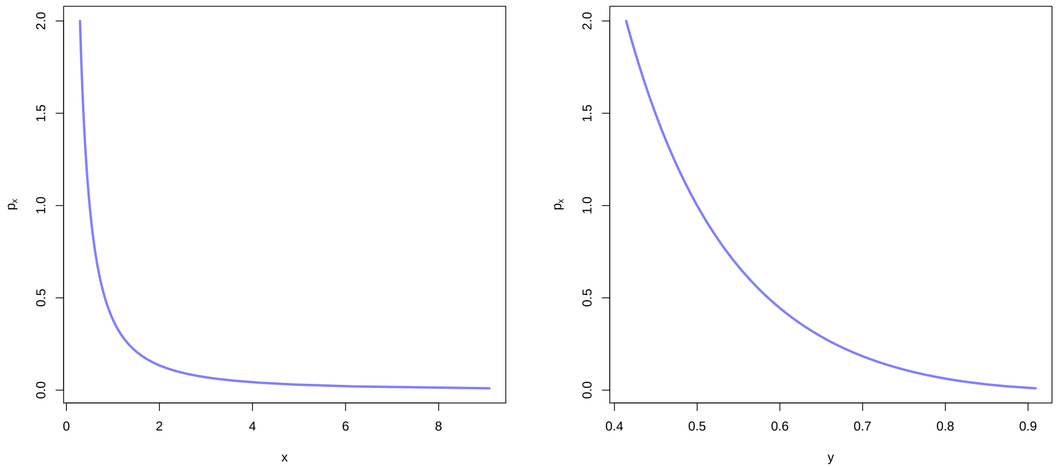

Figure 2: demand curves for goods x (left) and y (right) as a function of p_x, when $w = 1$ and $p_y = 1$.

consumption changes with wealth (or prices), rather than as consumption of a single good is changed; that is, we may want to plot the change in consumption directly rather than piecemeal-per-commodity. This is particularly useful in the two-commodity model. Using the above functions, this is shown in Figures 4 and 5.

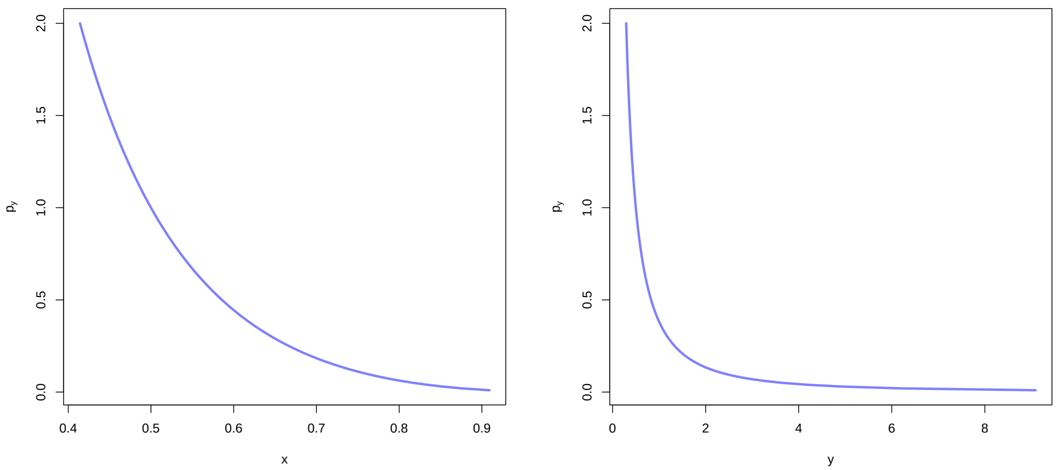

Figure 3: demand curves for goods x (left) and y (right) as a function of p_y, when $w = 1$ and $p_x = 1$.

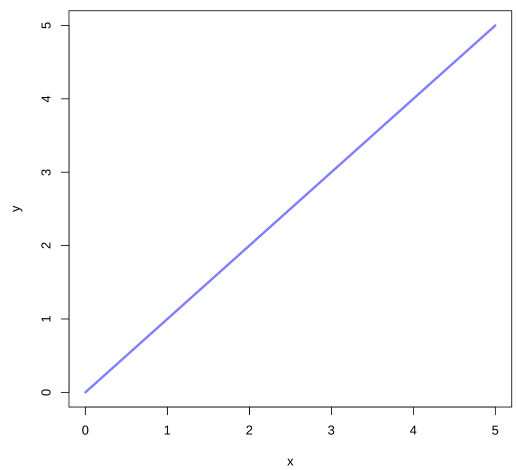

Figure 4: parametric consumption curve as a function of w, when $(p_x, p_y) = (1, 1)$.

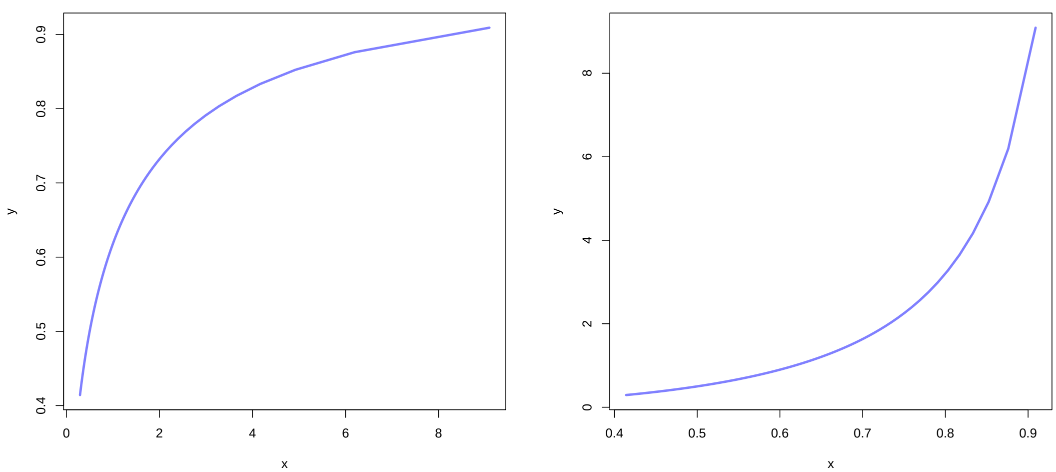

Figure 5: parametric consumption curve when $w = 1$, as a function of p_x with $p_y = 1$ (left) and as a function of p_y with $p_x = 1$ (right).