Some Microeconomic Theory[1]

1 In the beginning...

In the beginning, there were **consumer preferences**. Consider a consumer faced with a set S of consumption options. For any $x, y \in S$, we will write

- $x \sim y$ to mean that the consumer is indifferent between x and y;
- $x \succ y$ to mean that the consumer *strictly* prefers x to y; and
- $x \succsim y$ to mean that the consumer *weakly* prefers x to y.

We assume that the preference relation $\succsim$ satisfies the following conditions:

Completeness For all $x, y \in S$, either $x \succsim y$ or $y \succsim x$ (or both).

Transitivity For all $x, y, z \in S$, if $x \succsim y$ and $y \succsim z$ then $x \succsim z$.

Continuity If $x = \lim_{n\to\infty} \{x_n\}$, $y = \lim_{n\to\infty} \{y_n\}$, and $x_n \succsim y_n$ for all n, then $x \succsim y$.

These assumptions allow us to draw more-or-less familiar indifference curves. With two more assumptions, we can draw completely familiar indifference curves:

More is better than less If consumption bundle x has everything that bundle y has plus some more stuff, then $x \succ y$.

Convexity The set $\{y \in S : y \succsim x\}$ is convex.

These two additional assumptions allow us to draw indifference curves that are convex to the origin and increasing as we move away from the origin. Note also that the convexity assumption implies that there are *diminishing marginal rates of substitution*: if two of the goods in our consumption set are cakes and lattes, giving up lattes requires successively larger infusions of cake in order to remain on the same indifference curve.

Now that we have well-behaved indifference curves, we can do all the usual stuff. For example, the slope of an indifference curve measures the **marginal rate of substitution** between two goods. And we can draw budget constraints and all that stuff. More (much more) on all this later.

[1] These notes rely heavily on two books that are available in Penrose: Mas-Colell, Whinston, and Green's *Microeconomic Theory* (Oxford Univ. Press, 1995), which is basically the *Joy of Cooking* of microeconomics; and Jean-Jacques Laffont's *The Economics of Uncertainty and Information* (MIT Press, 1995).

2 What about utility functions?

Given the first three assumptions described above (completeness, transitivity, and continuity), we can show that there exists a continuous utility function $U(\cdot)$ that represents these preferences, i.e., such that $U(x) \geq U(y) \iff x \succsim y$. Given the final two assumptions described above (more-is-better-than-less and convexity), we can show that there exists a particularly well-behaved utility function that represents these preferences. (More formally: the utility function is increasing and quasi-concave.) Having utility functions—especially well-behaved utility functions—is good because utility functions are easier to understand and study than preference relations. (One reason for this may be that the mathematical operation $\geq$ is more familiar than the preference relation $\succsim$.)

Having said that, it is important to remember that the utility function $U(\cdot)$ is simply a handy representation of consumer preferences. Unlike economists around the turn of the 20th century, who discussed using "hedonmeters" or "psychogalvanometers" to measure utility, modern economic theory has a rather subtle view of utility, one based on the underlying supremacy of preferences.

Notably, it turns out that any single set of preferences can be described equally well by any of a number of utility functions: if $U(\cdot)$ represents the consumer's preferences, then so will $V(\cdot) = 2U(\cdot)$ or $W(\cdot) = e^{U(\cdot)}$ or (in general) any $f(U(\cdot))$ when $f(\cdot)$ is a strictly increasing function, i.e., where $f(a) > f(b) \iff a > b$. Because utility functions are defined only up to a strictly increasing transformation, many results about utility functions are fundamentally meaningless: when we say that $U(x) = 2U(y)$, we do not mean that the consumer is twice as happy with x than with y. (After all, we could easily describe this consumer's preferences using a utility function under which the utility from x is 100 times that of y.) The properties of utility functions that are invariant under strictly increasing transformations—i.e., that hold *no matter which utility function we use to represent preferences*—are called **ordinal properties**.

3 An easy way to incorporate uncertainty

We can incorporate uncertainty into our discussion of preferences simply by expanding our idea of a commodity, e.g., including such things as "gaining \$100 if a flipped coin comes up heads" or "having a root beer float if the temperature is more than 100 degrees". (These types of commodities are called **state-dependent commodities** because their existence depends on the state of the world.) What we are basically doing here is imagining parallel universes and having the consumer express preferences over the various consumption bundles in all the different universes.

This approach is useful in many contexts, and in fact it will serve as the benchmark for our study of insurance. In particular, we will focus on an uncertain situation with two possible states, "accident" and "no accident". If w_1 and w_2 represent the individual's wealth level in the accident and no-accident states, respectively, we can draw indifference curves showing the individual's willingness to make trade-offs between w_1 and w_2, write a utility function $U(w_1, w_2)$ that represents this individual's preferences, &etc.

4 A hard way to incorporate uncertainty

Forget about state-dependent commodities and return to the world of full certainty: there is a set S of consumption options, and the consumer has preferences over those consumption options that satisfy the properties described above. Now consider the set of *lotteries* over those consumption options, e.g., a 40% chance of getting consumption bundle x and a 60% chance of getting consumption bundle y. Using the preference relation $\succsim$ to describe consumer preferences *over lotteries*, we now make the following assumptions:

Completeness (Same as above) For any lotteries a and b, either $a \succsim b$ or $b \succsim a$ (or both).

Transitivity (Same as above) For any lotteries a, b, and c, if $a \succsim b$ and $b \succsim c$ then $a \succsim c$.

Continuity (Similar to above, but phrased differently) Given lotteries a, b, and c with $a \succsim b \succsim c$, there is a **compound lottery** involving a and c such that the consumer is indifferent between that lottery and b.

Independence (This one is new!) Given lotteries a, b, and c with $a \succsim b$, the consumer always prefers some combination of a and c to an otherwise-identical combination of b and c.

Given these assumptions, von Neumann and Morgenstern showed that the consumer's utility function over *lotteries* can be expressed using a utility function over *consumption options* that has an attractive form: it is linear in probabilities, meaning that the utility from a given lottery can be expressed nicely in terms of the utility of the various consumption options that arise from that lottery.

This could not be any more opaque, so let's do an example. Return to the accident example above, where the consumer has wealth level w_1 if there's an accident and wealth level w_2 if there's no accident. (Let the probability of an accident be 0.4, meaning that the probability of no accident is 0.6.) We know from above that there's a utility function $U(w_1, w_2)$ that reflects this consumer's preferences. Given the assumptions of von Neumann and Morgerstern, the **Expected Utility Theorem** that they proved shows that there is *another* utility function, $u(w)$, called a **von Neumann–Morgenstern utility function (VNM)**, such that $U(w_1, w_2) = 0.4 \cdot u(w_1) + 0.6 \cdot u(w_2)$. The consumer, in other words, is maximizing expected utility. The importance of the expected utility theorem can be seen in at least two ways: first, this result basically established the foundation for game theory and other theories of uncertainty; second, economists continue to make extensive use of the expected utility theory despite empirical evidence showing that it is (in the words of Matthew Rabin) an “ex-hypothesis”.

One final issue: recall that our original utility functions were unique only up to strictly increasing transformations, i.e., that the same set of preferences could be expressed by many different utility functions. It turns out that preferences *over lotteries* can also be expressed by many different utility functions, but here there are some additional limitations. Instead of allowing any strictly increasing transformations, VNM utility functions are unique up to *affine* transformations, meaning that if $u(\cdot)$ and $v(\cdot)$ are VNM utility functions that reflect the same consumer preferences, then there are some constants α and $\beta > 0$ such that $v(\cdot) = \alpha + \beta\, u(\cdot)$. The additional restrictions on VNM utility functions means that they have some **cardinal properties**. For example, it makes sense to talk about second derivatives of VNM utility functions because these are invariant under affine transformations: $u''(\cdot) = v''(\cdot)$ for any u and v that represent the same consumer preferences.