Microeconomía Financiera I Andrei M. Gomberg
Fall 2007

Notes on preferences, utilities and choice.

The material in these notes is fairly standard and can be found in any textbook of advanced microeconomics (including all those mentioned in the syllabus). You should also consult your class notes for examples, explanations, etc. I may be addding to these notes as the semester progresses.

1. Preferences and choices under certainty

1.X - **consumption (state) space.** You should interpret each element x of X as a complete description of everything that happens in a state of the world (i.e., if I say that $X = \{apple, orange, banana\}$ it means that the three states are only different by consumption of a given fruit. In particular, in this example, one cannot have simultaneously consume both apple and orange - had I wanted to consider such a possibility I would say that $X = \{apple, orange, banana, appleorange\}$

In this class we have consumption spaces that are

a) finite: $X = \{apple, orange, banana\}$

b) one-dimensional continuum: $X \subset \mathbb{R}$ (usual interpretation of a consumption bundle: quantity of a single good, such as the amount of money - *e.g., \$3 pesos*).

c) multi-dimensional continuum: $X \subset \mathbb{R}^n$ (usual interpretation of a consumption bundle a vector of consumption of different goods, where n is the number of goods - *e.g.* $\begin{pmatrix} 2 \\ 3 \end{pmatrix}$ - *2 apples and 3 oranges*)

Note: if I want to talk about some particular consumption bundle (e.g. "3 apples, 4 oranges, and a cup of tea"), I would specify it exactly. When I write "for any $x, y \in X$" I mean for any combination of two bundles; I don't mean that there is a specific bundle called x and another one called y (I know, it may seem ridiculous to review this here, but a lot of you seem to need it).

2. $\succsim$ - "at least as good as" (as in $apple \succsim banana$) - subjective individual **preference relation**.

$\succ$ - "strictly better than" ($apple \succ banana$ if $apple \succsim banana$ but not $banana \succsim apple$)

$\sim$ - "indifferent to" (if both $apple \succsim banana$ and $banana \succsim apple$

3. We say that the preference relation is **complete** if for any two elements x, y of X either $x \succsim y$ or $y \succsim x$ (or both)

We say that the preference relation is **transitive** if for any three elements x, y and z such that $x \succsim y$ and $y \succsim z$ it follows that $x \succsim z$ (note, that if X has only two elements transitivity is satisfied automatically).

If the preference relation is both complete and transitive we say that it is **rational.**

NOTE 1: if completeness of preferences fails, one can't make a choice between some pair of alternatives. If transitivity fails, three or more alternatives form a cycle, making it impossible to choose among them: think "money pumps".

NOTE 2: if $X = \mathbb{R}^2_+$ (the case for which we draw most of the diagrams), completeness means that once you draw an indifference curve, everything must be either on it, better than it, or worse than it. Transitivity means that indifference curves don't intersect.

NOTE 3: in this class rationality **never** means anything else: monotonicity, continuity, convexity, etc. are not part of rationality.

4. If the consumption space is $X \subset \mathbb{R}^n$, $n \geq 1$, then we studied other properties that preferences may (but don't have to) satisfy. These are: **strict monotonicity**, **convexity (non-increasing *marginal rate of substitution*)**, **strict convexity (decreasing *marginal rate of substitution*)** and **continuity** (the latter is a technical property you may simply assume holds in every example in class, unless stated otherwise). For definitions you should consult any textbook (e.g., Varian or Jehle and Reny).

5. Preferences are just **ordinal** rankings of objects. Of course, we would like to have **cardinal** measures of intensity of preferences, but this is not possible, since we do not know how to observe them.

6. In contrast, to observe preferences we only have to observe individual **choices:** we believe that individuals choose the "best" state (consumption bundle) according to their preferences. Note: if an individual with rational preferences has to choose one state out of a finite collection of states he can always find the best one according to his/her preferences (might not be true if preferences are not rational).

2. Utility function representation of preferences

1. It is hard to work with preferences directly, so we want to **represent** them with **utility functions.** Given a preference relation $\succsim$, a *real-valued* utility function u represents it if the following two statements are equivalent:We say that a **utility function** $u : X \to \mathbb{R}$ **represents** a preference relation $\succsim$ on X if for all x, y in X

$$\begin{aligned} x &\succsim y \\ &\text{is equivalent to} \\ u(x) &\geq u(y) \end{aligned}$$

$$\begin{aligned} x &\succsim y \\ &\text{and} \\ u(x) &\geq u(y) \end{aligned}$$

In other words, a utility function assigns more "utils" to better alternatives, and fewer utils to worse alternatives.

Example1: the preference relation $a \succ b \succ o$ over the set $X = \{a, o, b\}$ can be represented by the utility function $u(a) = 5, u(b) = 1, u(o) = -3$

Example 2: the preference relation $x \succsim y$ if and only if $x \geq y$ over the real line $\mathbb{R}$ can be represented by the utility function $u(x) = x$ or by the utility function $u(x) = x^3 - 5$.

2. Since "bigger than or equal" $\geq$ is a complete and transitive relation on the real numbers $\mathbb{R}$ it should be clear that it can only be used to represent complete and transitive (*rational*) preferences. Thus, rationality is a **necessary** condition for **representability** (existence of utility functions representing a preference). So, whenever you can write a utility function, you have rationality (rationality is implicitly in the utility function language).

3. If X has finitely many elements, all rational preferences over it can be represented by utility functions. If X is infinite (i.e., the usual $\mathbb{R}^n$ commodity space you worked with in your Intermediate Micro or the Propedeutico) there are some rational preferences that can't be represented (but this is just a technical quirk which we will ignore in this class).

4. If $w(x) = g(u(x))$ for some strictly increasing function $g : \mathbb{R} \to \mathbb{R}$ we say that u and w are monotonic transformations of one another.

5. Two utility functions represent the same preference if and only if they are monotonic transformations of one another. You should always apply monotonic transformations that simplify your work. But be (a bit) careful: a monotonic transformations of parts of a utility function do not have to result in a monotonic transformation of the entire utility function (the sum of the squares is not equal to the square of the sum - this is particularly relevant when we talk about intertemporal choice below).

6. Only **ordinal** properties (those unaffected by monotonic transformations) of utility functions matter. I.e., we can speak about increasing, or decreasing, or constant utility functions, but we can't speak about rates of increase *(there is no such thing as "twice as good", there is only "at least as good"; there is no such thing as "decreasing marginal utility" - marginal utility can only be positive, negative or zero)*. Concavity and convexity are meaningless in this context - though quasi-concavity is not. Never give any intensity interpretation to utility functions - "utils" don't exist.

7. 6. The following ordinal properties of utility functions represent properties of preferences:

A. strictly increasing functions represent strictly monotone preferences
B. quasi-concave functions represent convex preferences
C. strictly quasi-concave functions represent strictly convex preferences

8. The problem of finding "the best" state (consumption bundle) according to one's preferences now reduces to the problem of maximizing a utility function.

3. Intertemporal Preferences.

1. Until now all individuals "lived" for one period: born, act, realize uncertainty, enjoy consumption, die. What if individuals have multiple periods to act?

2. We can extend the model to incorporate time by assuming that each good is a distinct commodity at each moment. Suppose each individual consumes only

one good c at each moment and lives for T periods. The consumption space in this case is $\mathbb{R}_+^T$, and individual preferences are defined over **streams of consumption** $c = (c_1, c_2, ..., c_t, ...c_T)$. We assume the same rationality hypothesis as before (and if there is uncertainty, we also assume independence).

3. Under the same conditions as before (rationality and some technical assumptions) we can represent the preferences over the intertemporal consumption with a utility function $U(c_1, c_2...c_T)$.

4. For simplicity, we often assume *intertemporal additive separability* of preferences:

$$U(c_1, c_2...c_T) = \sum_{t=1}^{T} u_t(c_t)$$

This assumption implies that the **intertemporal marginal rate of substitution** between consumption in periods t and t'

$$\frac{\partial U(c_1, c_2...c_T) / \partial c_t}{\partial U(c_1, c_2...c_T) / \partial c_{t'}}$$

is independent of consumption in other periods.

5. We normally assume that the functions $u_i(c_i)$ are strictly concave. This implies the strict quasi-concavity of the function U (i.e., the convexity of preferences), which, in turn, implies preference for **consumption-smoothing** (i.e., it is better to consume something each period, than to consume a lot in one period and nothing in another period).

6. We normally assume that people are *impatient*: $u_t(c_t) > u_{t'}(c_{t'})$ whenever $t' > t$.

7. One way of modelling impatience is with **exponential discounting of utility:**

$$U(c_1, c_2...c_T) = \sum_{t=1}^{T} \beta^t u(c_t)$$

for some β between zero and one.

8. If people discount exponentially, their decisions are **time-consistent:** they may adopt a course of action at period 1 and even given the opportunity to re-decide later they won't change their original decision. Otherwise, their preferences

are not time consistent: if given a chance to re-decide in the future, they will want change what they do.

4. Probability:

0. Set notation:

$A \cup B-$union of A and B (A or B)

$A \cap B-$intersection of A and B (A and B)

$A \subset B - A$ is a subset of B

$A \backslash B-$ difference of A and B (in A, but not in B)

1. X - **state space** (what may occur)

In this class we, primarily, deal with money lotteries (i.e. $X \subset R$). In what follows, unless otherwise defined, every state is a real number $x \in X$, which denotes how much money a person has.

2. $E \subset X$ - **event (**e.g., the person has between \$2 and \$3)

3. **Probability** is a function P that assigns to each possible event a number such that

i) $0 \leq P(E) \leq 1$

ii) $P(X) = 1$

iii) For any two disjoint events ($A \cap B = \emptyset$) $P(A \cup B) = P(A) \cup P(B)$

(Note: strictly speaking, property (iii) has to be strengthened slightly to hold for countably infinite unions, but for the purposes of this class we will never need this).

4. An immediate consequence:

$P(X \backslash A) = 1 - P(A)$ (probability of A not happening is 1 minus probability of A happening). In particular, $P(\emptyset) = 0$.

Note: we say that event E occurs with **positive** probability if $P(E)$ is **strictly bigger** than zero.

Comment: in economics we often think of probabilities as **subjective**: probability of an event is *what an individual believes it to be.*

5. Lotteries

1. Let $X = \{x_1, x_2 ... x_N\}$ be the **state space** and $P = (p_1, ... p_N)$ a **probability distribution**. We call the pair $L = \{X; P\} = \{x_1, x_2, ... x_N; p_1, p_2, ... p_N\}$ a **lottery** (some books use the term **gamble** instead).

Examples:
A coin toss: $L = \left\{Head, Tail; \frac{1}{2}, \frac{1}{2}\right\}$
Throwing dice: $L = \left\{1, 2, 3, 4, 5, 6; \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}, \frac{1}{6}\right\}$
Lottery from the St. Petersburg paradox (due to Bernoulli): $L = \left\{2, 4, 8, 16, ... 2^n ...; \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, ..\right.$

2. If states are real numbers ($X \subset \mathbb{R}$) we can define *expected value of a lottery* (for X consisting N elements) as

$$E(X|L) = \sum_{i=1}^{N} p_i x_i$$

Note: expected value may not always exist. *Example:* try computing the expected value for the St. Petersburg paradox lottery.

3. From now on we shall assume that X is fixed and known. Given a fixed state space X we shall be interested in the set $\mathcal{L}(X)$ of all possible lotteries over X. When X is known, we may suppress it in writing and just identify the lottery with a probability distribution $L = \{p_1, p_2, ... p_N\}$ (notice, this is just an n-dimensional vector)

4. Besides the **simple** lotteries defined above, we may want to consider **compound** lotteries, *i.e.* lotteries over lotteries. For instance, what if individual is not sure, which of the lotteries $L = \{p_1, p_2, ... p_N\}$ or $L' = \{p'_1, p'_2, ... p'_N\}$ he/she is facing, but assigns a probability α to it being L (and, correspondingly, $(1 - \alpha)$ to it being L'). Consider the resulting *compound* lottery

$$L'' = \{L, L'; \alpha, 1 - \alpha\}$$

5. It is not hard to see (why?), that L'' can actually be viewed as a *simple* lottery

$$L'' = \{\alpha p_1 + (1 - \alpha) p'_1, \alpha p_2 + (1 - \alpha) p'_2, ..., \alpha p_N + (1 - \alpha) p'_N\} = \alpha L + (1 - \alpha) L'$$

A compound lottery is, thus, simply an average of simple lotteries (note, that an average of the lotteries is still a lottery, while the sum of lotteries is not - think why). Nonetheless, it is sometimes convenient to use the compound lottery interpretation, as we shall see shortly.

6. Preferences over lotteries.

1. Consider space X and the corresponding lottery space $\mathcal{L}(X)$. We are interested in the way individuals choose actions, when these choices result in uncertain consequences - i.e., in the way individuals choose among lotteries.. The standard hypothesis in economics is to assume that individuals have preferences $\succsim$ over the lottery space $\mathcal{L}(X)$ (we say $L \succsim L'$ - lottery L is at least as good as lottery L'; think of these in the same way as you thought about preferences over states).

2. As in the case of preferences over certainty, define the properties of preferences reflexivity, transitivity and completeness (**rationality**). Recall, that completeness means that the individual can compare any two lotteries, while transitivity means that if, in his view, one lottery is at least as good as another, which, in turn, is at least as good as the third, then this individual will rank the first lottery as being at least as good as the third. Since the formal definitions are, essentially, identical to the certainty case we shall skip them here.

7. Utility function representation and Expected Utility.

1. As in the certainty case, given rationality of preferences $\succsim$, together with an additional regularity assumption ("continuity"), we can conclude that there exists a utility function

$$U : \mathcal{L}(X) \to \mathbb{R}$$

which represents these preferences, i.e.

$$L \succsim L' \Leftrightarrow U(L) \geq U(L')$$

2. However, if the theory is to be workable, we want to be able to take expectations of utilities. However, consider the "no uncertainty" lotteries $L_1 = \{1, 0, 0, ..., 0\}$, $L_2 = \{0, 1, 0, 0...0\}$, ..., $L_N = \{0, 0, ...1\}$. It seems reasonable to

view utility of such a lottery as simply utility of having the corresponding state for sure. Let us denote

$$u(x_k) = U(L_k)$$

It turns out, however, that IT IS NOT generally true that utility of a lottery is equal to the expected utility of the states, given the probability distribution of the lottery.

3. The expected utility theory which is sketched below was developed by John von Neumann and Oscar Morgenstern in their book *Game Theory and Economic Behaviour.*

4. **Assumption (independence axiom)**: for any lotteries L, L', L'' and for any α between zero and one

$$L \succsim L' \Leftrightarrow \alpha L + (1-\alpha) L'' \succsim \alpha L' + (1-\alpha) L''$$

5. Interpretation of independence: suppose, when comparing L and L' I choose L; now suppose that, whether I choose L or L', with some probability $(1-\alpha)$ I will be facing a different lottery L''. Independence says, that this additional uncertainty *which is the same whatever my choice is*, should not matter for my choice (if the chance of being robbed en route to cashing your lottery ticket is the same, whatever the ticket you bought, possibility of being robbed should not affect which ticket you buy).

Note: if independence fails, individuals's choices would not be time-consistent: their rankings would change after the realization of uncertainty. Think "Dutch books".

6. Von Neumann and Morgenstern showed that under assumptions of rationality, continuity AND independence it is possible to choose a particular - utility function U, which represents the preferences AND has an expected utility form:

$$U(L) = \sum_{i=1}^{N} p_i u(x_i) = E(u(X)|L)$$

The "small" u (defined over the STATES) is known as **Bernoulli utility function**, while the "big" U (defined over the lotteries) is known as **von Neumann - Morgenstern (vNM) utility function**. (Note, that the nomenclature is not yet settled - many books use a different terminology, actually calling u **vNM utility function**, while U - **expected utility function**).

7. Knowledge of an individual's Bernoulli utility function u over the states, uniquely defines his/her preferences over the larger domain of the lotteries. Since it carries a lot more information than an arbitrary utility function over X, a monotonic transformation may, actually, distort some of this information. It is only when the two utility functions are *linear* monotonic transformations of one another:

$$w(x) = au(x) + b$$

where $a > 0$, b - any real number, that they give rise to representations of the same preferences over uncertainty.

Comment: since we are primarily dealing with lotteries over amounts of money, and since it is not unreasonable to assume that without uncertainty everyone would prefer more money to less money, we shall always assume in what follows, that all Bernoulli utility functions are strictly increasing.

8. Attitude towards risk. Risk Aversion.

1. In fact, the concavity or convexity of a Bernoulli utility function (properties that are not generally preserved by arbitrary monotonic transformations but are preserved by the linear transformations) actually carries a lot of information about individual attitude to risk. In fact, the mathematical property known as **Jensen's inequality** tells us that for any concave u

$$E(u(X)|L) = \sum_{i=1}^{N} p_i u(x_i) \leq u(\sum_{i=1}^{N} p_i x_i) = u(E(X|L))$$

In words: an individual, whose preferences are represented by a concave Bernoulli utility function prefers to have the amount of money equal to the expected value of the lottery to having the lottery itself. An individual like this is known as **risk-averse**. The inequality is reversed for convex u - such individuals are called **risk-loving**. Finally, the inequality becomes *equality* when u is linear (e.g., $u(x) = x$). Such individual cares only about the expected value of the lottery, and is indifferent to risk; he/she is known as **risk-neutral**. We shall primarily consider risk-averse individuals in what follows.

2. A risk-averse individual would, generally, exchange a lottery for a fixed sum of money somewhat smaller than the lottery's expected value. This amount of money, denoted $c(L;u)$, given by the equality

$$u\left(c\left(L;u\right)\right)=E\left(u\left(X\right)|L\right)$$

the individual's **certainty equivalent** of the lottery. The smaller is the certainty equivalent, the more the individual is willing to pay to get rid of uncertainty.

2. Consider two individuals, with preferences represented using Bernoulli utility functions u and w, respectively. If, for every lottery L it turns out that $c(L;u) \leq c(L;w)$ it is reasonable to say that the individual with Bernoulli utility function u is more risk-averse than the one with Bernoulli utility function w.

Alternatively, we may want to say that an individual with a "more concave" utility function is the more risk-averse one. One way to formally define this is to say that u is more concave than w if it is a concave transformation of w, *i.e.*

$$u\left(x\right)=\varphi\left(w\left(x\right)\right)$$

for some increasing and concave function φ.

A third way to define "more concave" (and, therefore, more risk-averse) is to define a curvature coefficient of the Bernoulli utility function, which is known as **Arrow-Pratt coefficient of absolute risk aversion***:*

$$r_A(x;u)=-\frac{u''(x)}{u'(x)}$$

3. Fortunately, it turns out that the three definitions of "more risk-averse" coincide. In other words, the following three statements are equivalent:

(i) $c(L;u) \leq c(L;w)$ for every lottery L
(ii) $u(x)=\varphi(w(x))$ for some increasing concave φ
(iii) $r_A(x;u) \geq r_A(x;w)$ for every state x in X

If you find that any of the above holds you may conclude that individual with B.u.f. u is more risk averse than the one with B.u.f w.

4. You should keep in mind, that this is only an incomplete ranking: it is possible that you can't compare two risk-averse individuals in this way. However, within an important subclass of utility functions, known as **constant absolute**

risk-aversion (CARA) utility functions, the ranking is complete. These are the utility functions of the form

$$u(x) = -ae^{-rx} + b$$

For these, the coefficient of absolute risk-aversion

$$r_A(x; u) = r$$

5. The problem with CARA utility functions is, that such individuals are concerned only with the *absolute* size of the risk they are exposed to. Whether a possible loss of \$1000 represents .5%, 5%. or 50% of his net worth, an individual with CARA preferences will view it in the same manner. If you believe, that the possibility of a loss of \$1000 for Bill Gates now that he has his billions, is less unpleasant, than when he was a poor Harvard drop-out, you believe that his preferences exhibit decreasing, not constant absolute risk-aversion.

6. Therefore, it may make sense to consider the Arrow-Pratt **coefficient of relative risk-aversion**

$$r_R(x; u) = -x\frac{u''(x)}{u'(x)}$$

which focuses on the size of the gamble as a proportion of one's wealth. Constant relative risk-aversion implies decreasing absolute risk-aversion (think why?)

7. You will commonly encounter the *constant relative risk-aversion (CRRA)* utility functions

$$\begin{aligned} u(x) &= \frac{x^{1-\sigma}-1}{1-\sigma}, 0<\sigma<1 \\ u(x) &= \ln(x), \sigma = 1 \end{aligned}$$

where

$$r_R(x; u) = \sigma$$

9. Ranking of lotteries by payoff (FOSD) and by risk (SOSD).

1. When can we unambiguously say that one lottery gives a higher payoff than another? Clearly, simply considering expected value of the lottery is insufficient (why?), and expected utilities reflect necessarily subjective individual preferences.

2. We say that lottery L has higher payoff than lottery L' (L **First Order Stochastically Dominates (FOSD)** L') if every individual who prefers more money to less money agrees that he/she prefers L to L'. In other words

$$E\left(u\left(X\right)|L\right) \geq E\left(u\left(X\right)|L'\right)$$

for every increasing u.

3. It turns out, that this is equivalent to saying that lottery L can be obtained from lottery L' by an upward shift of the probability distribution (recall discussion in class). Recalling a definition from your probability theory class, an alternative way of saying the same thing is to say that the distribution function F associated with L is always smaller than or equal to the distribution function G associated with L'

4. If lotteries have the same expected value we, obviously, cant' compare them with respect to payoff. However, we can compare them with respect to risk. We say that lottery L is *less risky* than lottery L' (L **Second Order Stochastically Dominates (SOSD)** L') if their expected values are equal

$$E(X|L) = E(X|L')$$

and every *risk-averse* individual who prefers more money to less money agrees that he/she prefers L to L'. In other words

$$E\left(u\left(X\right)|L\right) \geq E\left(u\left(X\right)|L'\right)$$

for every increasing and *concave* u.

5. It turns out, that this is equivalent to saying that lottery L' can be obtained from lottery L by a *mean-preserving spread* of the probability distribution (once again, recall discussion in class).