Expected Utility*

Bruno Salcedo†

Winter 2020

Economic agents often lack information relevant for their choices. These notes present some of the most common models of decision making under uncertainty. The focus is on the standard model of objective expected utility, but some important departures are briefly discussed.

1. The Origin of the Expected Utility Hypothesis

The origin of the theory of probability is often attributed to a series of works from the XVIIth century. One of the central questions addressed by these works is how to assign fair prices or *values* to games of chance. The prevailing approach at the time was to use what we now call *expected value*, multiplying each possible gain or loss by its probability and adding up the products (Maĭstrov, 1974). Nicolaus Bernoulli found that such an approach could lead to paradoxical conclusions. As a thought experiment, he devised a hypothetical game with an infinite expected value, such that most people would only pay a small amount to play. The thought experiment is now called the St. Petesburg Paradox and it is described by Nicolaus Bernoulli's cousin—Daniel Bernoulli—as follows[1]

*Lecture notes for Econ 9602 at Western University.

†Department of Economics, Western University · brunosalcedo.com · bsalcedo@uwo.ca

[1] The original paradox involved a six-sided dice. It was first proposed by Nicolaus Bernoulli in private correspondence to Pierre Rémond de Montmor and published in de Montmort (1708).

> Peter tosses a coin and continues to do so until it should land "heads" when it comes to the ground. He agrees to give Paul one ducat if he gets "heads" on the very first throw, two ducats if he gets it on the second, four if on the third, eight if on the fourth, and so on, so that with each additional throw the number of ducats he must pay is doubled. Suppose we seek to determine the value of Paul's expectation…
>
> Although the standard calculation shows that the value of Paul's expectation is infinitely great, it has… to be admitted that any fairly reasonable man would sell his chance, with great pleasures, for twenty ducats. (Bernoulli (1738), p. 31)

Gabriel Cramer and Daniel Bernoulli independently came up with very similar solutions to the paradox. Instead of using the expectation of gains and losses, they proposed to use the expectation of a *utility function* defined over wealth. In a 1728 letter to Nicolaus Bernoulli, Cramer writes:

> [T]he discrepancy between the mathematical calculation and the vulgar evaluation… results from the fact that, *in their theory*, mathematicians evaluate money in proportion to its quantity while, *in practice*, people with common sense evaluate money in proportion to the utility they can obtain from it.

In turn, Bernoulli (1738) argued that:

> The concept of *value*… may be defined in a way that renders the entire procedure universally acceptable without reservation. To do this, the determination of the *value* of an item must not be based on its [*prize*],[2] but rather on the *utility* it yields. The [prize] of the item is dependent only on the thing itself and is equal for everyone; the utility, however, is dependent on the particular circumstances of the person making the estimate. There is no doubt that a gain of one thousand ducats is more significant to the pauper than to a rich man though both gain the same amount.

[2] The original work in Latin uses the word "pretium", which can mean both price or prize. The quote from a translation to English by Dr. Louise Sommer which translates "pretium" as "price". I believe that "prize" fits the context better, but I am not an expert in Latin.

> …[L]et us use this as a fundamental rule: *if the utility of each possible profit expectation is multiplied by the number of ways in which it can occur, and we then divide the sum of these products by the total number of possible cases, a mean utility will be obtained, and the profit which corresponds to this utility will equal the value of the risk in question.* (pp. 24)

Both Cramer and Bernoulli had a sense that the utility function should exhibit decreasing marginal utility. In his letter, Cramer considered two examples: a function that becomes constant for sufficiently high levels of wealth, and $\sqrt{\cdot}$. Bernoulli (1738) in part argued that

> [T]he utility resulting from a small increase in wealth will be inversely proportionate to the quantity of goods previously possessed. (p. 25)

This is the description of logarithmic utility functions. Either of the proposed utility functions would give Paul a finite expected utility. For more on the early history of the expected utility hypothesis, other resolutions, and interesting aspects of the paradox see Dutka (1988) or Seidl (2013).

2. Objective Expected Utility Theory

Consider a decision maker choosing between lotteries. Each lottery induces a probability distribution over a set of outcomes. The probabilities are objective and known to the decision maker. The expected utility hypothesis asserts that the decision maker behaves *as if* maximizing the expectation of a utility function defined over the set of outcomes.

This section presents a behavioral characterization of the expected utility hypothesis due to von Neumann and Morgenstern (1944). If you are looking for further reading material, Chapter 8 in Gilboa (2009) has an an interesting overview and philosophical discussion of the von Neumann-Morgenstern conditions. And Chapter 5 in Kreps (1988) discusses some important mathematical details and generalizations.

2.1. A Representation Theorem for Expected Utility

Fix a *finite* set of outcomes $X = \{x_1, \ldots, x_n\}$.[3] A *lottery* is a probability measure over X. Since X is finite, each lottery can be represented by a vector of probabilities $p = (p_1, \ldots, p_n) \in [0,1]^n$ with $\sum_{i=1}^n p_i = 1$, The set of lotteries is denoted by ΔX. It is a compact and convex subset of $\mathbb{R}^n$. When $n = 3$, it can be represented graphically by a triangle. Each vertex of the triangle represents a degenerate probability that assigns full probability to one of t in he three outcomes.

Suppose that we observe a preference relation $\succcurlyeq$ over simple lotteries. Let $\succ$ and $\sim$ denote strict preference and indifference, defined in the usual way. A function $U : \Delta X \to \mathbb{R}$ is said to *represent* $\succcurlyeq$ if for every pair of lotteries p and q

$$p \succcurlyeq q \quad \Leftrightarrow \quad U(p) \geq U(q). \tag{1}$$

And it is called an *expected utility function* if there exists some $u : X \to \mathbb{R}$ such that for all $p \in \Delta X$

$$U(p) = \sum_{x \in X} p(x) u(x). \tag{2}$$

In that case, the function u is called the corresponding *Bernoulli utility function*.

We are interested in whether $\succcurlyeq$ can be represented by an expected utility function. From the theory of choice under certainty, we know that a preference relation admits a continuous utility representation if and only if it is complete, transitive, and continuous. Hence, these are also necessary conditions for an expected utility representation. von Neumann and Morgenstern (1944) found a third condition called the Independence Axiom.

Axiom 2.1 (Rationality) $\succcurlyeq$ is complete and transitive.

Axiom 2.2 (Continuity) $\{p \in \Delta X \mid p \succ q\}$ and $\{p \in \Delta X \mid q \succ p\}$ are open for all $q \in \Delta X$.

[3] Assuming that X is finite greatly simplifies the exposition. The results in this section can also be generalized to cover lotteries with finite support even if X is not finite, and even to general Borel probability measures over Euclidean spaces by imposing additional continuity assumptions. See for instance Kreps (1988) pp. 57–68.

Axiom 2.3 (Independence) For all lotteries $p, q, r \in \Delta X$ and all $\mu \in (0,1)$,

$$p \succcurlyeq q \quad \Leftrightarrow \quad \mu p + (1-\mu) r \succcurlyeq \mu q + (1-\mu) r.$$

The Independence Axiom has an important and interesting normative interpretation in terms of consequentialism and dynamic consistency. See for instance Gilboa (2009) pp. 82. However, I would like to emphasize a geometric interpretation instead. Expected utility functions are linear in probabilities. Hence, we need an axiom that imposes linearity. The following proposition can be interpreted as saying that the Independence Axiom implies that indifference curves are parallel lines. See the figure. The proof is left as an exercise for the problem set.

Proposition 2.1 *If $\succcurlyeq$ satisfies the Independence Axiom, then for all $p, q \in \Delta X$ and $r \in \mathbb{R}^n$ such that $p + r, q + r \in \Delta X$,*

$$p \succcurlyeq q \quad \Leftrightarrow \quad (p + r) \succcurlyeq (q + r) \tag{3}$$

Rationality, continuity, and the independence axiom are both sufficient and necessary for the existence of an expected utility representation. The proof of the theorem is deferred to the next subsection.

Theorem 2.2 (von Neumann and Morgenstern (1944)) *A preference relation $\succcurlyeq$ over ΔX admits an expected utility representation if and only if it satisfies axioms 2.1–2.3.*

Utility functions are unique up to monotone transformations, but not all utility functions are linear. Within the class of linear utility functions there is some cardinal information.

Proposition 2.3 *Suppose that U us an expected utility representation of $\succcurlyeq$. A function $V : \Delta X \to \mathbb{R}$ is also an expected utility representation of $\succcurlyeq$ if and only if there exist $a > 0$ and $b \in \mathbb{R}$ such that $V(p) = aU(p) + b$ for all $p \in \Delta X$.*

2.2. Omitted Proofs

Establishing necessity of the axioms for the existence of an expected utility representation is straightforward. Hence, I will only prove sufficiency. I will do so by proving a series of lemmas using arguments from Herstein and Milnor (1953). In particular, note that sufficiency follows directly from lemmas 2.6 and 2.7 below.

Lemma 2.4 *If $\succcurlyeq$ is complete and satisfies the Independence Axiom, then for all lotteries $p, q \in P$ such that $p \succ q$ and all $\lambda \in [0,1]$, $\mu p + (1-\mu)q \succcurlyeq \lambda p + (1-\lambda)q$ if and only if $\mu \geq \lambda$.*

Proof. If $\mu = \lambda$, then the conclusion of the lemma follows from completeness. Now, suppose that $\mu > \lambda$. I will show that, in that case, $\mu p + (1-\mu)q \succ \lambda p + (1-\lambda)q$. There are three cases to consider. (Case 1) If $\mu = 1$, then

$$\mu p + (1-\mu)q = p = \lambda p + (1-\lambda)p \succ \lambda p + (1-\lambda)q, \tag{4}$$

where the strict preference follows from $p \succ q$ and the Independence axiom. (Case 2) If $\lambda = 0$, then

$$\mu p + (1-\mu)q \succ \mu q + (1-\mu)q = q = \lambda p + (1-\lambda)q. \tag{5}$$

(Case 3) Finally, suppose that $1 > \lambda > \mu > 0$. Let $\eta = \mu/\lambda$ and $r = \eta p + (1-\eta)q$. Since $\eta \in (0,1)$, the Independence Axiom implies that $p = \eta p + (1-\eta)p \succ r$. Using the Independence Axiom again it follows that $\lambda p + (1-\lambda)q \succ \lambda r + (1-\lambda)q$. Finally,

$$\lambda r + (1-\lambda)q = \lambda \cdot \left(\frac{\mu}{\lambda}p + \left(1 - \frac{\mu}{\lambda}\right)q\right) + (1-\lambda)q = \mu p + (1-\mu)q. \tag{6}$$

Therefore, $\lambda p + (1-\lambda)q \succ \mu p + (1-\mu)q$.

A completely analogous argument can be used to show that if $\mu \ngeq \lambda$ then $\mu p + (1-\mu)q \nsucccurlyeq \lambda p + (1-\lambda)q$, thus completing the proof of the lemma. ■

Lemma 2.5 *If $\succcurlyeq$ satisfies axioms 2.1–2.3, then for all $p, q, r \in P$ such that $p \succcurlyeq q \succcurlyeq r$ and $p \succ r$, there exists a unique $\mu^*_{pr}(q) \in [0,1]$ such that $r \sim \mu^*_{pr}(q)p + (1 - \mu^*_{pr}(q))r$.*

Proof. Consider the sets $\Lambda^+ = \{\lambda \in [0,1] \mid r \succcurlyeq \lambda p + (1-\lambda)q\}$ and $\Lambda^- = \{\lambda \in [0,1] \mid \lambda p + (1-\lambda)q \succcurlyeq r\}$. Since $\succcurlyeq$ is complete, $\Lambda^+ \cup \Lambda^- = [0,1]$. Since $\succcurlyeq$ is continuous, Λ^+ and Λ^- are closed. Since $[0,1]$ is connected, this implies that $\Lambda^+ \cap \Lambda^- \neq \emptyset$. That is, there exists some $\lambda \in (0,1)$ such that $r \sim \lambda p + (1-\lambda)q$. Uniqueness follows from Lemma 2.4. ■

As an intermediate step in order to establish the existence of an expected utility representation, I will establish the existence of a linear utility representation according to the following definition. A utility function $U : \Delta X \to \mathbb{R}$ is said to be *linear with respect to mixtures* if and only if

$$U\Big(\mu p + (1-\mu)q\Big) = \mu U(p) + (1-\mu)U(q), \tag{7}$$

for all lotteries $p, q \in \Delta X$ and all $\mu \in [0,1]$.

Lemma 2.6 *If $\succcurlyeq$ satisfies axioms 2.1–2.3, then it admits a utility representation which is linear with respect to mixtures.*

Proof. Since $\succcurlyeq$ is complete, transitive, and continuous, it admits a continuous utility representation. Since ΔX is compact, Weierstrass' Extreme-Value Theorem implies that there exist lotteries $p_0, p_1 \in \Delta X$ such that $p_1 \succcurlyeq p \succcurlyeq p_0$ for all $p \in \Delta X$.[4] If $p_0 \sim p_1$, then the agent is indifferent between all lotteries and any constant utility function works. Otherwise, we can set $U(p) = \mu^*_{p_0 p_1}(p)$, where $\mu^*_{p_0 p_1}(p)$ is the weight from Lemma 2.5. Lemma 2.4 implies that U represents $\succcurlyeq$. Hence, it only remains to verify linearity. Note that for all $pq \in \Delta X$ and $\mu \in [0,1]$,

$$\begin{aligned}
\mu p + (1-\mu)q &\sim \mu\Big[U(p)p_1 + (1-U(p))p_0\Big] + (q-\mu)q \\
&\sim \mu\Big[U(p)p_1 + (1-U(p))p_0\Big] + (1-\mu)\Big[U(q)p_1 + (1-U(q))p_0\Big] \\
&= \Big[\mu U(p) + (1-\mu)U(q)\Big]p_1 + \Big(1 - \Big[\mu U(p) + (1-\mu)U(q)\Big]\Big)\, p_0,
\end{aligned}$$

where the first two indifference comparisons follow from the Independence Axiom and the definition of U, and the equality follows from simple algebra. From the

[4]The proofs in both von Neumann and Morgenstern (1944) and Herstein and Milnor (1953) requires additional lemmas, because neither of them assumes X to be finite. Hence, they cannot guarantee the existence of a best and worst lotteries. In fact, they consider not just probability spaces, but general algebraic structures called mixture spaces.

uniqueness of $\mu^*_{p_0p_1}(\mu p + (1-\mu)q)$ it follows that

$$U\Big(\mu p + (1-\mu)q\Big) = \mu U(p) + (1-\mu)U(q). \qquad \blacksquare$$

The final step to prove Theorem 2.2 is to establish the following lemma. The proof of the lemma is left as an exercise for the problem set.

Lemma 2.7 *A function U is linear with respect to mixtures if and only if it is an expected utility function.*

Proof of Proposition 2.3. The *if* part of the proof is straightforward. Hence, I will only prove the *only if* part. Suppose that both U and V are expected utility representations of $\succcurlyeq$. Let p_0 and p_1 be the best and worst lotteries, as in the proof of Lemma 2.6. If $p_1 \sim p_0$, then both U and V are constant, and the result is trivial. Hence, suppose for the rest of the proof that $p_1 \succ p_0$.

Fix any lottery $p \in \Delta X$. Let μ be the weight from Lemma 2.5 such that $p \sim \mu p_1 + (1-\mu)p_0$. Since U is a linear representation of $\succcurlyeq$ (Lemma 2.7), it follows that

$$U(p) = U(\mu p_1 + (1-\mu)p_0) = \mu U(p_!) + (1-\mu)U(p_0). \tag{8}$$

After some simple algebra it follows that

$$\mu = \frac{U(p) - U(p_0)}{U(p_1) - U(p_0)}. \tag{9}$$

Since the same holds for V, it follows that

$$\frac{U(p) - U(p_0)}{U(p_1) - U(p_0)} = \frac{V(p) - V(p_0)}{V(p_1) - V(p_0)}, \tag{10}$$

which implies that $V(p) = aU(p) + b$ where

$$a = \frac{V(p_1) - V(p_0)}{U(p_1) - U(p_0)} \quad \text{and} \quad b = V(p_0) - \left(\frac{V(p_1) - V(p_0)}{U(p_1) - U(p_0)}\right)U(p_0). \tag{11}$$

Since $p_1 \succ p_0$, it follows that $a > 0$. Since a and b do not depend on p, the proof is complete. $\blacksquare$

3. Subjective Beliefs

The objective expected utility model assumes the existence of objective probabilities, and assumes that individuals know them. However, there are settings were the meaning of the word "probability" is not entirely clear. For instance, ask yourself what does it mean when a news outlet makes a sentence of the sort "the polls give candidate X a chance of $x\%$ of wining the election." Even if we could define probabilities, economic agents might not know how to assign probabilities to all relevant events. Even when they do, there is no guarantee that different individuals will assign the same probabilities to the same events. For example, some individuals might believe that all the outcomes of a roulette wheel are equally likely, while others might be convinced that the roulette is rigged in favor favor the house, and perhaps some believe that red spaces are "hot" at the moment and are more likely to land.

One common approach is to abandon the notion of objective probabilities. Instead, we can assume that people have an intuitive notion of 'likelihood' in their minds. Then, we can think of probabilities as a subjective measure of this likelihood. In order to distinguish objective and subjective probabilities, I will call the later ones *beliefs*. The *subjective expected utility hypothesis* is that agents behave as if maximizing the expectation of a utility function with respect to a subjective *belief*. Savage (1954) proposed a framework that allows to define beliefs and infer them from the agents' preferences over uncertain prospects

In Savage's setting, an economic agent choosing acts $a \in A$. She cares about potential outcomes $z \in Z$, and she is uncertain about the true state of the world $x_0 \in X$. The outcomes depend both on the act chosen by the agent and the true state of the world. To capture this idea, Savage defined acts as functions $a : X \to Z$ mapping states into consequences. In this way, $a(x)$ would be the consequence of act a if the true state was $x_0 = x$.

Suppose that the researcher observes the agents preferences $\succcurlyeq$ over acts. The subjective expected utility hypothesis is that the agent behaves *as if* there existed a utility function over consequences $u : Z \to \mathbb{R}$ and a probability function $p \in \Delta X$ such that

$$a \succcurlyeq b \iff \int_X u(a(x))\, dp(x) \geq \int_X u(b(x))\, dp(x),$$

for all pairs of acts $a, b, \in A$. In that case we call p the (subjective) beliefs of the agent about the state, and we say that (u, p) represent the agent's preferences.

Savage proposed with a list of axioms or postulates on preferences over acts. His postulates characterize the empirical content of the subjective expected utility hypothesis. Moreover, they allow the researcher to recover a unique belief from an idealized choice data. The introduction to Savage (1954) is an excellent read, but the axioms can be better understood from Kreps (1988) or Gilboa (2009).

The key step in Savage's work is to establish what he calls the *sure-thing principle*, which he describes as follows.

> A businessman contemplates buying a certain piece of property. He considers the outcome of the next presidential election relevant. So, to clarify the matter to himself, he asks whether he would buy if he knew that the Democratic candidate were going to win, and decides that he would. Similarly, he considers whether he would buy if he knew that the Republican candidate were going to win, and again finds that he would. Seeing that he would buy in either event, he decides that he should buy, even though he does not know which event obtains, or will obtain, as we would ordinarily say. (Savage (1954), p. 21)

More generally, suppose that an individual would prefer a to be b if they learned that event A is true. Further suppose that the individual would also prefer a to be b if they learned that event A is false. Then, the individual should also prefer a to b before learning whether A is true or false. To most people, this sounds like a natural conclusion. Consider for instance the following quote.

> Except possibly for the assumption of simple ordering, I know of no other extralogical principle governing decisions that finds such ready acceptance. (Savage (1954), p. 21)

Two of Savage's postulates are sufficient to derive the sure thing-principle. The first one is that the agent's preferences over acts should be rational.

Axiom 3.1 $\succcurlyeq$ is complete and transitive

In order to state the second postulate, we need some aditiona notation. Given acts f and g and an event $Y \subseteq X$, let $f_Y g$ denote the act given by

$$[f_Y g](x) = \begin{cases} f(x) & \text{if } x \in Y \\ g(x) & \text{if } x \notin Y \end{cases}$$

Axiom 3.2 For all acts f, g, h, k and all events $Y \subseteq X$, $f_Y h \succ g_Y h$ if and only if $f_Y k \succ g_Y k$

Postulate 3.2 allows us to define preferences conditional on an event Y. Say that $f \succ_Y g$ if there exists an act h such that $f_Y h \succ g_Y h$. Now, we are in position to formally state and prove the sure-thing principle.

Proposition 3.1 (Sure-thing principle) *Under postulates 3.1 and 3.2, given any acts f and g, and any event Y, if $f \succ_Y g$ and $f \succ_{X \setminus Y} g$, then $f \succ g$.*

Proof. Let f, g, and Y satisfy $f \succ_Y g$ and $f \succ_{X \setminus Y} g$. The definition of $\succ_A$ implies that $f_Y h \succ g_Y h$ for some act h. Postulate 3.2 thus implies that $f = f_Y f \succ g_Y f$. By an analogous argument, $f_{X \setminus Y} g \succ g_{X \setminus Y} g = g$. Note that $g_Y f = f_{X \setminus Y} g$. Hence, $f \succ g_Y f \succ g$, and the result follows from the transitivity of $\succ$ (Postulate 3.1). ■

4. Other Forms of Choice Under Uncertainty

References

Bernoulli, D. (1954/1738). Exposition of a new theory on the measurement of risk [Specimen theoriae novae de mesura sortis]. *Econometrica*, 22(1):23–36. Translated from Latin by Sommer, L.

de Montmort, P. R. (1708). *Essay d'analyse sur les jeux de hazard*. Jaques Quillau.

Dutka, J. (1988). On the St. Petersburg Paradox. *Archive for History of Exact Sciences*, 39(1):13–39.

Gilboa, I. (2009). *Theory of Decision under Uncertainty*, volume 45 of *Econometric Society Monographs*. Cambridge University Press.

Herstein, I. N. and Milnor, J. (1953). An axiomatic approach to measurable utility. *Econometrica*, 21(2):291–297.

Kreps, D. (1988). *Notes on the Theory of Choice*. Underground Classics in Economics. Westview Press.

Maĭstrov, L. E. (1974). *Probability Theory: A Historical Sketch*. Academic Press.

Savage, L. J. (1954). *The Foundations of Statistics*. John Willey & Sons.

Seidl, C. (2013). The St. Petersburg paradox at 300. *Journal of Risk and Uncertainty*, 46(3):247–264.

von Neumann, J. and Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton University Press.