Examining Expected Utility Theory from Descriptive and Prescriptive Perspectives

A draft by Silviu Pitis, dated January 2, 2010

1. Abstract

At the very foundation of financial theory lies the theory of decision-making under uncertainty. Understanding how and why people make decisions, as well as how they *should* be making decisions, is crucial. The quantitative treatment of preferences under uncertainty constitutes some key assumptions in Markowitz's Modern Portfolio Theory and the widely used Capital Asset Pricing Model, among others. Several decades ago, this quantitative treatment could be equated with Expected Utility Theory, and according to Machina (1987), "was considered one of the success stories of economic analysis". Since then however, Expected Utility theory has been challenged on several fronts, and its descriptive validity today is, no pun intended, quite uncertain.

This paper is written as a thorough examination of Expected Utility theory, from its roots, to its validity as a model, both descriptive and prescriptive. The failures of Expected Utility as a descriptive model are outlined and examined in detail. Observed dynamic inconsistency among rational agents calls into question the prescriptive validity of Expected Utility. The issue is resolved with a discussion of outcome spaces, from which several complications arise. Financial applications of Expected Utility are discussed, and an alarming observation is made: many of the descriptive failures of Expected Utility serve as the root cause for its incorrect usage as a prescriptive model. These incorrect applications are identified, and the paper concludes with a few brief examples of how one might properly construct a Utility function.

2. Expected Utility: It's history and usage

The roots of Expected Utility theory trace back to the early 18th century, when Bernoulli proposed that evaluating uncertain monetary outcomes using expected value was neither an adequate method of describing human behavior, nor did it always produce rational choices. He proposed that individuals attribute utility values to each final state of wealth, and attempt to maximize this utility when making decisions. The theory resolved several paradoxes in decision theory, including the notable St. Petersburg Paradox.

A breakthrough in Expected Utility theory came in 1944, when John von Neumann and Oskar Morgenstern formalized the theory with an axiomatic approach. They showed that if rational agents follow four basic axioms of choice[1], then they possess what is termed a *von Neumann-Morgenstern (VNM) utility function*. The theory states that rational agents always attempt to maximize their expected utility, defined as the probability-weighted average of the VNM utility of each final state of wealth.

As an example, an individual might be presented with the choice of receiving \$10 or having a 50-50 chance of receiving either \$15 or \$5. We call each option a lottery, and express it in shorthand as $\{x_1, p_1; x_2, p_2; \ldots; x_n, p_n\}$[2], where outcomes $x_1, x_2, \ldots, x_n$ are mutually exclusive, unambiguous, and collectively exhaustive, with $p_1, p_2, \ldots, p_n$ as their respective probabilities. Obviously, $\Sigma p = 1$. The individual in the above example would be choosing between $\{10, 1\}$ and $\{15, .5; 5, .5\}$. According to Expected Utility theory, the individual with wealth *w* would possess a VNM utility function, $U(\cdot)$, and would value the utility of each lottery as $\Sigma U(w+x_i)p_i$. There are two important things to note here. First, utility is calculated based on final wealth states and not on absolute changes in wealth. Second, utility is linear in terms of probabilities, so that a 100% chance of receiving any amount is exactly twice as valuable as a 50% chance of receiving that same amount (and a 50% chance of receiving nothing), regardless of what VNM utility function is used.

A common assumption is that individuals are *risk-averse*, which implies that their VNM utility functions are concave. That is, the marginal utility of an additional dollar of wealth falls as wealth increases. The individual above, would therefore choose to receive \$10, since the concavity of $U(\cdot)$ implies $(1)U(w+\$10) > (0.5)U(w+\$15) + (0.5)U(w+\$5)$. Even though both lotteries have the same expected value (w+\$10), the risk-averse individual prefers the one that has less uncertainty.

As a result of the logical validity of its axioms, the Expected Utility hypothesis quickly garnered wide acceptance in the field of decision theory. It has since been applied extensively in Economics, Finance, and Game Theory. It has been used, among other things, for the purposes of personal financing planning, formulating asset allocation strategies, pricing insurance, and determining optimal strategies in games of chance.

In contrast to this widespread use however, is a substantial body of evidence suggesting that the theory fails as a descriptive model. Empirical results have shown that the vast majority of people make choices that systematically violate the predictions of Expected Utility theory. More

[1] The four axioms are *completeness, transitivity, continuity,* and *independence*. A full statement of the axioms and the resulting proof of the Expected Utility hypothesis can be found at http://homepage.newschool.edu/het//essays/uncert/vnmaxioms.htm

[2] Curly brackets { } have been used as opposed to the usual ( ) for the purpose of clarity

specifically, it has been shown that the *independence axiom*, which implies linearity in probabilities, does not hold descriptively.

To demonstrate, suppose an individual has a strict preference between two lotteries. Say they prefer receiving $10000 with certainty, {10000, 1}, to receiving $20000 75% of the time, {20000, .75}[3]. The independence axiom states that for any probability *p*, they will always prefer the "mixture lottery" consisting of a *p* chance of the preferred option and a (1-*p*) chance of any lottery X to the one consisting of a *p* chance of the less preferred option and a (1-*p*) chance of lottery X. Consider the two lotteries above, and suppose *p* = 0.1 and lottery X is the zero lottery (0, 1). Then the independence axiom implies that the lottery {10000, .1} is preferred to {20000, .075}. In reality, the majority of individuals may prefer $10000 with certainty to a 75% chance of $20000, but contradict the independence axiom by choosing a 7.5% chance of $20000 over a 10% chance of $10000.[4] This contradiction of the independence axiom is just one of the many empirical failures of the independence axiom and Expected Utility.

3. Failures of Expected Utility as a Descriptive Model

There are several well-documented systematic failures of Expected Utility theory as a descriptive model. These can be broadly classified into *perceptual failures* and *non-linearity failures*.[5] The first, which results entirely from biases in human perception, includes the *isolation effect, framing effects,* and *preference reversal phenomenon*. The second implies that humans have non-linear probability preferences, and is entirely independent of perceptional biases. It includes the *Allais Paradox*, *common ratio effect,* and *calibration failure,* among many others. Many empirical examples have a hint of both perceptual failures and non-linear probability failures, but the difference, as will be demonstrated, is very clear. Before examining the data, it is important to note that these empirical results have absolutely no bearing on the normative validity of Expected Utility theory, for such an evaluation could be done only logically and philosophically.

Borrowed Examples of Perceptual Failures

The three selected examples below illustrate empirical failures of Expected Utility that occur because of biases in perception. The biases at fault are examined afterwards.

[3] Note that {$2000,.75} is unambiguous shorthand for {$2000,.75; 0,.25}
[4] This statement not tested, but rather inferred from several very similar and statistically significant results found in Kahneman & Tversky (1979) among others
[5] This is an original classification created for the purposes of this paper

Example 1*: The Isolation Effect*[6]

> *"*Consider the following two-stage game. In the first stage, there is a probability of .75 to end the game without winning anything, and a probability of .25 to move into the second stage. If you reach the second stage you have a choice between {4000, 0.8} and {3000, 1}. Your choice must be made before the game starts, i.e., before the outcome of the first stage is known."

If the first stage is ignored and only the second stage is presented, the majority of people choose {3000, 1}. Accordingly, when presented with both stages isolated as above, the majority of people choose {3000, 1}. However, when the two stages of the game are not isolated, i.e., respondents are asked to choose between {4000, .2} and {3000, .25}, the majority choose {4000, .2}. Thus, two different representations, or framings, of the same problem generate contradictory choices. This effect is also an example of a *framing effect*.

Example 2: *Framing Effects*[7]

> *"*Q1*:* In addition to whatever you own, you have been given $10000. You are now asked to choose between {5000, 1} and {10000, 0.5}.
>
> Q2: In addition to whatever you own, you have been given $20000. You are now asked to choose between {-5000, 1} and {-10000, 0.5}."

The two choice problems above are identical, in that they produce the exact same distributions of final wealth. Both are equivalent to a choice between {15000, 1} and {10000, 0.5; 20000, 0.5}. Risk aversion predicts that {15000, 1} is chosen in both cases, but the majority of people select {10000, 0.5; 20000, 0.5} when the choice is framed as in Q2. This is contradictory to the risk aversion assumption, but more importantly, it is self-contradictory since the majority of people presented with both Q1 and Q2 will choose {15000, 1} in Q1 and {10000, 0.5; 20000, 0.5} in Q2.

Example 3*: Preference Reversal*[8]

> *"*In this study [Lichtenstein & Slovic (1971)], subjects were first presented with a number of pairs of bets and asked to choose one bet out of each pair. Each of these pairs took the following form:
>
> P-bet {X, *p*: x, 1-*p*} versus $-bet {Y, *q*; y, 1-*q*}

[6] Borrowed from Kahneman & Tversky (1979)

[7] Adapted from Kahneman & Tversky (1979)

[8] Borrowed from Machina (1987) who borrowed it from Lichtenstein & Slovic (1971)

Where X > x, Y > y, p > q, and Y > X (the names "P-bet" and "$-bet" come from the greater probability of winning in the first and the greater possible gain in the second)."

After this first choice problem, subjects were asked to value each bet individually (state the maximum price they would pay for the bet). Of the 173 subjects who were each presented with 6 such choices between P-bets and $-bets, 73% made a preference reversal every time they chose the P-bet over the $-bet. That is, when asked to value the P-bet and $-bet individually, they gave a higher value to the $-bet despite having chosen the P-bet in an outright choice.

Where Perception Fails

In his Nobel lecture paper "Maps of Bounded Rationality: Psychology for Behavioral Economics", Kahneman explores several of the biases at fault for the apparently contradictory behaviors above. A distinction between two different 'modes' of decision-making, intuition and reasoning, is crucial. "While intuition is fast, automatic, effortless, and emotional, reasoning is slow, controlled, effortful, and neutral."[9]

Note that in all three examples above, the choices made are largely intuitive. In order to make these choices, most individuals do not sit down, consider their VNM utility function (or any other decision making model for that matter), and make careful calculations to determine their choice. While there may be some reasoning involved in making the choices, the final decision relies largely on intuition. This tendency to rely on intuition for decisions involving uncertainty is not isolated to the three choice examples above, but can be generalized to most choice problems. Furthermore, the rules governing intuition are generally similar to those governing perception, and thus our decision making is subject to several biases in perception.[10]

The most notable biases arise from differences in accessibility. Accessibility is defined as the ease with which we can interpret certain information about things we perceive (i.e. people, objects, or for our purposes choice problems).

Consider three different representations of the same number: its canonical form $2^{12}3^{3}17^{1}$, its decimal form 1880064, and its hexadecimal form 1cb000. All three forms represent the same number, and therefore contain the exact same information. However, some pieces of information are more accessible than others. The magnitude of the number is easily accessible from its decimal form, but not so accessible from its canonical form. It is much easier to determine the greater of 1880064 and 4826809 than it is to determine the greater of $2^{12}3^{3}17^{1}$ and 13^{6}. On the

[9] Kahneman (2003) pg. 1451
[10] Kahneman (2003) pg. 1450

other hand, the prime factorization of the number is immediately accessible from its canonical form, but not from its decimal or hexadecimal forms. Determining that 2^{12} is a factor of 1880064 requires some calculation, but is a trivial task given the canonical representation. To one familiar with hexadecimal notation, it is trivial that 2^{12} (or 16^3) divides 1cb000. This last point demonstrates that sufficient training or familiarity with specific types of representations might make certain information more accessible, and therefore help in making intuitive judgments.

In a similar fashion, the method in which a choice problem is represented or framed alters the accessibility of its various attributes. Because our intuitive judgment passively accepts the information given and is not concerned with reframing problems, we are left only with the narrow frame provided and have only the information that was readily accessible. Furthermore, it should be noted that our perception and thus intuition is concerned primarily with changes in states, as opposed to the states themselves. Both the *isolation effect* and *framing effects* can be explained as a direct consequence of these factors.

Examine the *framing effects* problem. In both Q1 and Q2, the first sentence about a certain fixed gain is almost ignored completely by our intuition. This changes the state of our wealth, but the change is prior to the choice problem, which is what our intuition is most concerned with. In Q1, the frame provided makes accessible to us information about a potential gain, whereas in Q2, the frame makes accessible to us information about a potential loss. As discussed later in this report, it is common for people to exhibit risk-averse behavior when faced with potential gains and risk-seeking behavior when faced with potential losses[11], thus explaining their contradictory choices in Q1 and Q2. See you if you can use the perceptual biases discussed to explain the *isolation effect* failure.

Providing an explanation for the *preference reversal* phenomenon is not as easy, and requires some consideration. It is thus surprising that *preference reversal* was in fact predicted by Slovic and Lichtenstein before it was empirically observed. Their prediction arose from an observation that individuals value probabilities of positive outcomes more in direct choice problems, but value the $-amount of positive outcomes more when determining buying and selling prices of lotteries.[12] In psychological terms, this is called a *response mode effect*, and is a result of people using a different cognitive method for responding to choice problems than the one they use in responding to valuation problems.[13]

[11] This behavior is by itself contradictory to Expected Utility theory, but it is unclear whether it is a perceptual failure or a non-linearity failure. This will be discussed later.

[12] Slovic and Lichtenstein (1983), as referenced by Machina (1987)

[13] Machina (1987)

Non-linearity Failures

A second class of systematic failures arises from preferences that are non-linear with respect to probabilities. The perceptual failures that were discussed are clearly caused by irrational behavior, but it is not so evident whether this is true of non-linearity failures. While you read the following examples of non-linearity failures, ask yourself if the observed behavior is rational or not.

The Allais Paradox and Common Ratio Effect

In 1953, Allais proposed one of the first challenges to Expected Utility theory, the *Allais Paradox*. Consider the following two choice problems:

i. Choose between *a1*: {1 million, 1} and *a2*: {5 million, .10; 1 million, .89; 0, .01}
ii. Choose between *a3*: {5 million, .10} and *a4*: {1 million, 0.11}

If the independence axiom holds, then a choice of *a1* in the first pair implies a choice of *a4* in the second pair and a choice of *a2* in the first implies a choice of *a3* in the second.[14] Contrary to this prediction, several researchers have shown that the majority of people presented with the choice problems above pick *a1* in the first and *a3* in the second. While the *Allais Paradox* is an isolated example of non-linear preferences, there are several classes of systematic violations of the independence axiom.[15]

The *common ratio effect* describes the class of violations that take the form of the two choice problems below, where $p > q$, $Y > X > 0$, $0 \leq r \leq 1$:

i. Choose between *b1*: $\{X, p; 0, 1\text{-}p\}$ and *b2*: $\{Y, q; 0, 1\text{-}q\}$
ii. Choose between *b3*: $\{X, rp; 0, 1\text{-}rp\}$ and *b4*: $\{Y, rq; 0, 1\text{-}rq\}$

Recall the violation of independence given in the last paragraph on page 2. This was an example of the *common ratio effect* using X = \$10000, Y = \$20000, $p = 1$, $q = .75$, and $r = 0.1$. While the independence axiom would mean a choice of *b1* in the first pair implies a choice of *b3* in the second, empirical evidence shows that for at least some choices of *X, Y, p, q,* and *r;* the majority of people exhibit preferences of *b1* in the first and *b4* in the second.

[14] Consider lotteries *x:* {5M, 10/11; 0, 1/11} and *y*: {1M, 1}; the two choice problems in the Allais Paradox are both mixture lotteries involving *x* and *y*, and so according to the independence axiom, the choices should be consistent with the preference ordering of *x* and *y.*

[15] Actually, the Allais Paradox is an example of the *common consequence effect*, a general class of violations that is not dealt with in this paper. For more information on this see Machina (1987) pg. 129.

A Failure in Calibration

Another troublesome failure of Expected Utility as a descriptive model is that it cannot explain the significant risk aversion that individuals exhibit over small or even modest stakes. As discussed in Rabin (2000), the Expected Utility curve cannot be calibrated to match this 'extreme' risk aversion. Empirical results, such as in Kahneman & Tversky (1991) among others, have demonstrated that a majority of people generally reject 50-50 gain/loss lotteries of the form $\{G, .5; -L, .5\}$ unless the potential gain is in the order of two times the loss. That is, regardless of stakes, unless G is about twice as large as L, people tend to reject the lottery even though its expected value might be positive. Rabin demonstrates the troublesome implications of assuming individuals possess a VNM utility curve given their observed risk aversion over small and modest stakes.

Take for example an Expected Utility maximizing individual who turns down the lottery {11, .5; -10, .5} at all wealth levels. Since the individual would turn down {11, .5; -10, .5} if he had 10 more dollars, this means the first additional 10 dollars have more marginal utility than the next 11 dollars. Let X be the average utility of the first 10 dollars and Y be the average utility of the next 11. Then 10X > 11Y, so that $\frac{10}{11}$ X > Y. Due to risk aversion, the 1st dollar is the most valuable of the first 10, and the 21st dollar is the least valuable of the next 11. It follows that the 21st dollar is worth less than $\frac{10}{11}$ of the 1st. By an analogous argument, the 42nd dollar is worth less than $\frac{10}{11}$ of the 21st, the 63rd less than $\frac{10}{11}$ of the 42nd, and so on. Suppose the 1st dollar is worth 1, then the summed utility of every 21st dollar is less than $1 + \frac{10}{11} + \left(\frac{10}{11}\right)^2 + \left(\frac{10}{11}\right)^3 \ldots = 11$ and the summed utility of all dollars is less than 21*11 = 231. Starting at any wealth level and using a similar argument, it can be shown that the utility of every 21st dollar lost is at least $\frac{11}{10}$ of the last dollar. Then the summed utility of the first 21*8 = 168 dollars lost is at least $21*(1 + \frac{11}{10} + \left(\frac{11}{10}\right)^2 + \ldots + \left(\frac{11}{10}\right)^8) = 240.15$. It follows that if Expected Utility maximizing individuals turn down the lottery {11, .5; -10, .5} at all wealth levels, then they will turn down *any* lottery that involves a 50% chance of losing $168, regardless of how large the potential gain is (even if that gain is infinite!).

Using a generalized algebraic argument of the same flavor with slightly stronger inequalities, Rabin reduces the above conclusion to turning down *any* lottery that involves a 50% chance of losing $100. As another example, he shows that an individual who turns down {105, .5; -100, .5} at all wealth levels will turn down *any* lottery involving a 50% chance of losing $2000. This means that the latter individual prefers to have $200,000 to a 50-50 chance of having either $198,000 or $50,000,000,000 (or any other arbitrarily high number).

While these rather ridiculous conclusions hinge on the unrealistic assumption of small stakes risk aversion for *all* wealth levels, Rabin further generalizes his theorem to include specified ranges of wealth. Take for example an individual known to turn down the lottery {105, .5; -100, .5} for all levels of wealth less than $300,000. Note that such an individual should be relatively easy to find given the findings of Kahneman & Tversky (1991). Suppose this person is at a current wealth level of $290,000 and behaves according to Expected Utility theory. Then Rabin's theory concludes that they will turn down {71799110, .5; -20000, .5} – still a rather ridiculous conclusion.

The ultimate implication of Rabin's work on calibration is that Expected Utility maximizers should be very close to risk neutral (use expected value to make decisions) over modest stakes. This is clearly not the case and the empirical evidence shows that the majority of individuals display risk-averse behavior even over very small stakes.[16]

Distinguishing between Perceptual and Non-linearity Failures

I propose that the ultimate rule for distinguishing between perceptual and non-linearity failures is as follows:

> If, after being made aware of his perceptual biases and all relevant information regarding the problem, a subject changes his response to a choice problem from a non-expected utility response to a response that is consistent with expected utility theory, then the failure of his initial response was a *perceptual failure*.
>
> If a subject maintains a response inconsistent with expected utility theory, even after being made aware of his perceptual biases and all relevant information regarding the problem, then his response is a *non-linearity failure.*

We can conclude from the findings of Macrimmon (1968), Moskowitz (1974), and Slovic and Tversky (1974) as referenced in Machina (1987), that individuals who display non-expected utility preferences in the *Allais Paradox* generally do not change their opinion, even after being

[16] For example, see Kahneman & Tversky (1979)

presented with several arguments against their choice and careful consideration. Thus, the *Allais Paradox* can generally be considered a non-linearity failure.

Certain distinctions are not so easy to make. The reader may have noticed that the first example of a perceptual failure given, the *isolation effect*, is almost analogous to the *common ratio effect*. Determining whether an individual making a non-expected utility choice in such a situation is exhibiting a perceptual failure or a non-linearity failure of Expected Utility therefore comes down to whether or not they revise their initial choice upon further consideration.

As an example, consider an individual who is presented with the choice problem outlined in the *isolation effect* on page 4[17]. Suppose that they display non-expected utility preferences by initially choosing {3000, 1} over {4000, .8} and, as shown by the arrow in **FIG 1**, choosing {4000, .20} over {3000, .25}. Then, they are presented with the two-stage framing of the same problem **(FIG 2)[18]**, and allowed time to critically analyze their choice.

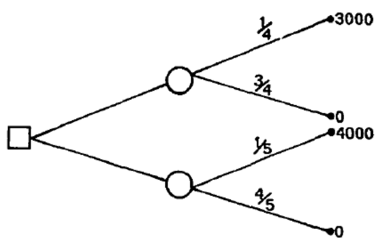

Fig 1:[19]

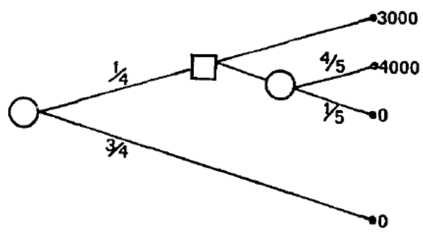

Fig 2:

Observing that given the choice, they prefer {4000, .20} over {3000, .25}, the subject would choose the down path in figure 1. In order to stay consistent with this choice over final wealth distributions, the subject's planned choice in Fig 2, made before the first chance node is resolved, would be to go down, since this would produce the 'preferred' distribution of final outcomes, {4000, .20}.

However, upon reaching the decision node in figure 2, a large part of the uncertainty has been resolved, and the subject is confronted with a decision between {3000, 1} and {4000, .8}. Because they prefer the former, they choose to go up, as indicated by the arrow. Thus, the subject's actual choice is inconsistent with their planned choice. This inconsistency is termed *dynamic inconsistency.* Any direct violation of the independence axiom can be framed in this way, so that anyone violating the axiom is dynamically inconsistent in this specific context.[20]

[17] This problem is borrowed from Kahneman & Tversky (1979)
[18] Figure 1 and 2 have identical distributions over final wealth, and therefore represent the same choice
[19] Fig. 1 and 2 borrowed directly from Kahneman & Tversky (1979). The squares in each tree represent a decision node, and the circles represent random chance events.
[20] Note that *dynamic consistency* is a more general term that refers to any scenario in which current preferences contradict past preferences

The subject now has two choices. The subject can choose to resolve this issue of dynamic inconsistency, in which case he must change one of the two preference orderings involved in this choice so that his choices are consistent with Expected Utility axioms. If this is the case, then the failure of Expected Utility in this problem is rooted in perceptual biases.

Alternatively, he can accept that he is dynamically inconsistent, in which case he retains his non-expected utility preference ordering. If this is the case, then the failure of Expected Utility to describe the subject's behavior is rooted in the non-linearity of probabilities.

Conscious Acceptance of Dynamic Inconsistency[21]

Consciously accepting that one is dynamically inconsistent has some particularly concerning implications. Dynamic inconsistency implies *information aversion*. The subject in the above example preferred {4000, .20} over {3000, .25} and planned to go down at the choice node in figure 2. However, the free information that the subject received from the first chance node in figure 2 caused the subject to deviate from their planned choice and settle for the less preferred lottery {3000, .25}. Therefore, if one truly preferred the final wealth distribution resulting from choosing {4000, .20} over {3000, .25}, it would be in one's best interest to reject free information!

That one might be consciously information averse is alarming, if not nonsensical. An assumption of *information seeking*, the idea that reliable information should never hurt a decision maker, seems very attractive and highly "rational". In fact, the information seeking assumption could effectively replace the independence axiom in the development of Expected Utility theory.[22]

Thus, the conscious acceptance of dynamic inconsistency and therefore information aversion calls into question not only the descriptive validity of Expected Utility, but also its normative validity. Indeed, it would be very difficult to argue that the conscious choice leading to a non-linearity failure, made after all comprehending all available information and arguments, is irrational.

---

[21] The word conscious is used to refer to the comprehension of all available and relevant information

[22] Logically: 'not independence axiom' implies 'dynamically inconsistent' implies 'information averse', therefore 'information seeking' implies 'dynamically consistent' implies 'independence axiom'

4. The Normative Debate

The various empirical failures of Expected Utility theory suggest that it may not be suitable for descriptive purposes. Indeed, these failures have led to the formulation of several Non-Expected Utility models. These models do not assume linearity in probabilities (the independence axiom), and more accurately describe human behavior in certain scenarios. Non-Expected Utility models include Prospect Theory, Subjective Expected Utility, and Subjectively Weighted Utility, among others.[23]

Although these models were made specifically for descriptive purposes, the question of whether or not these models might be valid for prescriptive purposes is still open to debate. In fact, there is still debate with regards to the normative validity of Expected Utility. For example, Machina (1989) makes the argument that *consequentialism* is an implicit assumption of Expected Utility. Consequentialism means that unrealized past possibilities have no influence on the value of current states, or as Machina puts it, "whenever a choice in a decision tree is reached the part of the tree before the current choice node is 'snipped' off". Some thought reveals that consequentialism is essentially equivalent to the independence axiom[24], and that people who are not consequentialist over the stated outcomes are actually considering outcomes beyond the current *outcome space*. In other words, the x_i defined in the lottery (see page 2) are ambiguous and do not describe all dimensions of the relevant outcomes. Interestingly enough, rational people who behave as *non-Expected Utility maximizers* over one outcome space might behave as *Expected Utility maximizers* over another outcome space, and vice versa.

The Outcome Space Argument

An example given by Machina (1989) against consequentialism involves "Mom" and her two children, Abigail and Benjamin. She has an indivisible treat, which she can choose to give to either Abigail (lottery A) or Benjamin (lottery B). Mom strictly prefers A or B to not giving either child the treat. However, she doesn't want to play favorites, and strictly prefers to flip a coin for a 50:50 chance of either A or B. She prefers this 50-50 flip to any other mixture of probabilities since it is the fairest.

Benjamin, having been exposed to the dynamic consistency argument, decides to impose consequentialism on Mom. He gets Mom to write down that she prefers a flip to either A or B, and if he loses the flip (the coin lands on A), he shows her the paper so that instead of giving the

[23] See Kahneman & Tversky (1979), Savage (1954), and Karmarkar (1978, 1979) respectively
[24] And therefore also equivalent to *dynamic consistency* and *information seeking*

treat to Abigail, Mom decides to ignore past uncertainty and do another flip. By repeating this, Benjamin ensures that lottery B is chosen, and he gets the treat.

Clearly, it would not make sense for Mom to flip the coin twice, as this would produce a different distribution over final outcomes and be dynamically inconsistent with her initial choice to flip once. Mom rejects consequentialism, and does not flip the coin twice. To any onlooker, Mom's behavior is perfectly rational. Her preferences {A, .5; B, .5} > {A, 1}, {B, 1} make sense, and so does her rejection of Benjamin's proposal for another flip. However, even her initial preference ranking, as stated, outright contradicts the independence axiom. Has Expected Utility failed normatively, and is Mom a non-Expected Utility user?

The point to be made here is that the possible outcomes (A: Abigail receives treat; B: Benjamin receives treat) are ambiguous; they do not fully describe the relevant outcomes. Mom prefers {A, .5; B, .5} to A or B because she doesn't want to play favorites. In other words, she is considering "playing favorites" as a relevant dimension of the potential outcomes. The relevant outcome space is therefore (A_1: Abigail receives treat, Mom plays favorites; A_2: Abigail receives treat, Mom is fair; B_1: Benjamin receives treat, Mom plays favorites; B_2: Benjamin receives treat, Mom is fair). By expanding the outcome space, we eliminate all ambiguity in outcomes, and the independence axiom holds. Her preference ranking becomes {A_2, .5; B_2, .5} > {A_1, 1}, {B_1, 1}, and Benjamin's strategy can no longer be used, since the second flip with be a choice between {A_2, 1} and {A_1, .5; B_1, .5}, not {A_1, 1} and {A_2, .5; B_2, .5}. Over the relevant outcome space, Mom is behaving consistently with consequentialism and the independence axiom.

In cases of dynamic inconsistency and non-linearity failures of Expected Utility, an analogous outcome space argument can be made. By broadening the outcome space and considering all relevant outcomes, violations of the independence axiom can be resolved. Recall the *isolation effect* example:

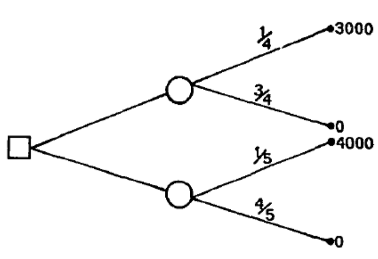

Fig 1:[25]

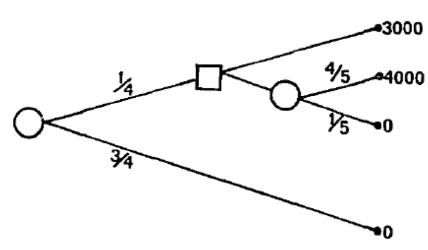

Fig 2:

Here the possible outcomes are all $-values, and the stated outcome space is dollars. If an individual behaves in a dynamically inconsistent manner (by choosing the red arrows), they are

[25] Fig. 1 and 2 borrowed directly from Kahneman & Tversky (1979). The squares in each tree represent a decision node, and the circles represent random chance events.

violating the independence axiom. However, this discrepancy can be easily resolved by expanding the outcome space to include the dimension of resolved risk, which the individual considers relevant to the decision. In figure 1, the individual is choosing between wealth distributions that are risky, whereas in figure 2, much of the risk has already been resolved. While it is true that the risk of the wealth distribution is the same before the first node of either tree, the resolved risk at the decision point might change the outcomes from the perspective of the decision maker. In figure 1, the possible outcomes are (0: it probably would've happened anyways; 3000: if I got this lucky I'd probably have gotten 4000; 4000: lucky!). In figure 2, the possible outcomes of the decision are (0: could have been 3000 richer!; 3000: good thing I didn't risk it; 4000: lucky!). The 'free' information discussed previously is no longer free – it makes the final outcomes less attractive.

Drawbacks of the Outcome Space Argument

The outcome space argument just discussed seems very attractive at first glance; it provides a simple, yet powerful, defense for Expected Utility on both normative and descriptive fronts. However, there are two critical drawbacks to the argument that cannot be ignored.

First, expanding the outcome space to include all relevant outcomes is difficult and can quickly become too cumbersome for any application. For Mom, the second dimension of the outcome space (Mom's fairness) was binary. Either she was fair, or she wasn't. In this particular example, measuring other dimensions was easy, but this is not always the case. In the *isolation effect* example, the second dimension of resolved risk was highly subjective. While it is easy to measure $-values, it is very difficult to quantify resolved risk. "It probably would've happened anyways" is not specific, and it is unclear how such a statement can be valued. Creating a VNM Utility function that included this dimension would be nigh-impossible. In some cases the relevant outcome space might include several dimensions, making the application of Expected Utility far too cumbersome.

Second, while expanding the outcome space to include dimensions of feeling or emotion might allow Expected Utility to serve in descriptive contexts, it is unclear as to whether or not the added dimensions are *rationally* relevant. Should outcomes such as "lucky!" and "could have been 3000 richer!" be considered in normative contexts? This is a worthwhile topic for philosophical debate, and depending on the context either side might be correct.

Relevant Outcomes for Financial Applications

For strictly financial applications of Expected Utility, one might say that *money* is the only rationally relevant outcome. Non-human rational agents, in particular corporations, should not be subject to emotions. Although a recent trend towards sustainability and the triple bottom line has expanded the relevant outcome space for some corporations, generating cash has always been the most important objective in business. [26]

Individual investors would be well-advised to leave their emotions and any other non-monetary aspects out of their financial decision-making. By considering more than just financial objectives, investors might find their wealth growing at a less than optimal rate. Of course, it is not easy to ignore all non-monetary factors (exemplified by the descriptive failures of Expected Utility), and so many investors choose to trust their wealth with an investment advisor. In this case, the investor is not the decision-maker, and only financial outcomes are relevant.

It should also be interesting to note that if investments are carried out in such a way that decisions are not in accordance with Expected Utility over the strictly monetary outcome space, opportunities for arbitrage may surface. This is beyond the scope of this paper, and a few arguments showing how arbitrage might occur can be found in Machina (1989).

If we consider money to be the only relevant outcome for financial applications, then it follows that Expected Utility theory should be the only valid prescriptive model. If it were not, this would mean that investors and corporations *should* be information averse and dynamically inconsistent over the monetary outcome space. However, as demonstrated above, this means that they are considering at least some non-monetary dimension of the outcomes– a contradiction. [27]

[26] The "triple bottom line" expands on the traditional "bottom line", profit, by adding in the dimensions of ecological and social performance

[27] It is important to note that although it can generally be ignored by discounting future cash flows to their present value, the *timing* of cash flows might also be considered a relevant dimension of the outcome space. In most cases, a two dimensional function of money (x dollars) and time (t years) can easily be defined as $V(x, t) = U(xe^{-rt})$ where U is the VNM utility function over current cash, and r is the continuous compounded annual rate of return.

5. Proper Prescriptive Usage of Expected Utility

The acceptance and use of Expected Utility as a prescriptive model is widespread in a variety of financial applications, including both Investment Theory, and Corporate Finance. In order for optimal financial decision making, it is in the best interests of the decision maker to use Expected Utility theory *properly*. However, the failures of Expected Utility as a descriptive model cross over and create significant challenges to proper usage in prescriptive applications.

Mean-variance Preferences as they relate to Expected Utility

The decision-making model generally used to evaluate portfolios is not actually Expected Utility, but rather based on Mean-Variance (MV) preferences. MV theory was developed in the 50s and 60s by Markowitz (1952), Tobin (1958), and Sharpe (1964), among others.[28] The use of MV preferences is fundamental in Markowitz's Modern Portfolio Theory, which is commonly applied in investment scenarios.

MV theory is accepted for the purposes of Modern Portfolio Theory as a general approximation to Expected Utility theory, and used in its stead for its simplicity. In the words of Morone (2008), MV theory "is the simplest model of investment that is sufficiently rich to be directly useful in applied problems … it is, also, rather obvious that Expected Utility should perform better than Mean Variance."[29]

Levy and Markowitz (1979) and Kroll, Levy, and Markowitz (1984) demonstrated empirically that MV preferences and Expected Utility preferences produce nearly identical choices among different sets of portfolios. The VNM utility functions examined in these studies include the logarithmic utility function, $U(x) = \ln(x)$, among others. It was found that unless investors have, as Levy and Markowitz (1979) write, "some very strange preferences among probabilities of return", the correlation between their Expected Utility preferences and MV preferences is nearly perfect. For example, the MV approximation of logarithmic utility produces a .995 correlation.[30]

Note that in these studies, MV preferences were found by approximating various VNM utility functions. Therefore, to properly apply MV preferences, either one must know what subject's VNM utility function is, or one must obtain their MV preferences directly. The latter option is subject to the same challenges as the former, which will now be discussed.

[28] As referenced in Morone (2008)
[29] Morone (2008)
[30] Levy and Markowitz (1979)

Challenges in Determining U(·)

Consider how one would go about constructing a VNM utility function for themselves, for a client, or for a subject. A common treatment is an elicitation or recovery, which would generally involve obtaining some information about the subject's preferences and matching this information to a utility function.[31]

Take for example the fractile method, which can be used to construct a subject's utility curve for a given wealth range, say 0 to W. The method first assigns U(0) = 0, and U(W) = 1. It then proceeds by arbitrarily selecting some mixture probability $\underline{p}$ and using it to determine wealth values, w_i, so that $U(w_1) = \underline{p}$, $U(w_2) = \underline{p}^2$, and so on. For example, say W = 100 and we use $\underline{p}$ = 0.5. Then, to find w_1 so that $U(w_1) = 0.5$, we elicit the subject's *certainty equivalent* of {100, 0.5; 0, 0.5}.[32] This can be done in a number of ways, including direct elicitation or questioning. Since U({100, 0.5; 0, 0.5}) = 0.5U(0) + 0.5U(100) = 0.5, $U(w_1) = 0.5$. This process is repeated for other mixtures lotteries involving $\underline{p}$, and the subject's 'recovered' VNM utility function from 0 to W is found.[33]

What is important to note here is that determining the wealth values w_i involves eliciting the certainty equivalents of various mixture lotteries from subjects, which a process that is heavily influenced by the descriptive failures of the independence axiom. Thus the validity of the elicited VNM utility function will be highly questionable. Indeed, various researches cited by Machina (1987) including Karmarkar (1974, 1978) and McCord have found that using a different value for $\underline{p}$ will elicit a different VNM utility function! As seen below, higher values of $\underline{p}$ generally produce functions with greater concavity (risk aversion).[34]

Fig 3:[35]

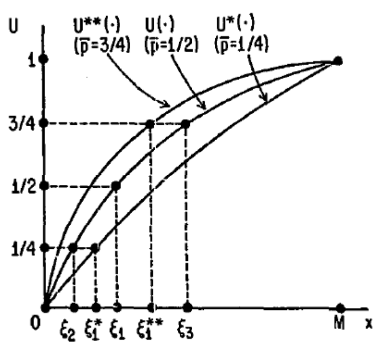

---

[31] Machina (1987), pg 124

[32] The certainty equivalent of a lottery is the largest certain amount a subject would pay to 'buy' the lottery, or alternatively, the smallest certain amount for which they would 'sell' the lottery

[33] Machina (1987, pg 124

[34] Machina (1987), pg 131 - 132

[35] Fig. 3 borrowed directly from Machina (1987) pg. 131

This is not an isolated failure of the fractile method, but will extend to any elicitation method that involves obtaining non-expected utility information from subjects and using it to create an expected utility function.

Without knowing the subject's VNM utility function one cannot prescriptively apply Expected Utility theory, and evidently, elicitation methods can't help. So is proper prescriptive usage of Expected Utility a hopeless case? A rather interesting detour into the world of professional gambling says otherwise.

A Trip to the Poker Table

The game of poker and its countless variants is a fascinating arena for the analysis and application of Expected Utility theory. The variety of poker players provides an excellent mix of non-expected utility users. Some have overly risk-averse preferences, while others have overly risk-seeking preferences. Some repeatedly make the same stochastically dominated[36] choices, while others play almost entirely randomly. Some play the game professionally, and they come extremely close to being Expected Utility maximizers. Because of the psychological barriers of intuition and emotion, as well as the limits on our cognitive ability to reason, nobody plays the game optimally. That is – nobody is a perfect Expected Utility maximizer.

However, the fact that optimal play is equivalent to maximizing one's Expected Utility[37], combined with the fact that optimal play can be analyzed through other means, provides a context in which Expected Utility theory can be analyzed in and of itself. The really interesting thing about this is that poker is not played in the monetary outcome space, but rather monetary outcomes result from game factors, most notably the exchange of table chips. Whenever outcomes in the relevant space (real dollars) are derived from an *acting* space (table chips), we have an *outcome space transformation*. If the VNM utility function over the relevant space is known, a VNM utility over the acting space can be derived.[38]

The general treatment of optimal poker play is to maximize *expected value* (EV) of real dollars. Note that this is consistent with the idea that over modest stake sizes, Expected Utility maximizers will maximize their EV (as concluded in the section on *calibration failure* on pg. 8). Thus, given the relevant function, U($) = $, and the rules of the game, one can find the acting function, U'(chips).

---

[36] A lottery is said to stochastically dominate another if it can be obtained from it by shifting probabilities from lower outcomes to higher ones. For example, {10, .5} dominates {10, .4} since .1 of the probability of receiving 0 in the second, can be shifted to a .1 probability of receiving 10 to obtain {10, .5}
[37] This is assumed, not proven; consider the examples given as evidence
[38] Again this is not proven; in fact, I am not sure it is always true, but it can be done in the examples given

In most cash games U'(chips) = chips, and optimal play consists of maximizing the expected value of one's chips, or table wealth. [39] However, unique aspects of the game might cause U'(chips) to deviate considerably from expected value.

Here is one example of such a scenario:

> I was sitting at a table with one of the largest *fish*[40] I had ever played against. He was making huge –EV decisions in nearly half of every hand dealt to him. To this point he had gotten insanely lucky; he had won several 20% lotteries, and was sitting with more than 5 times his initial buy-in. At this particular table, my strategy deviated considerably from maximizing EV. Here are 2 examples:
>
> When I was sitting with my initial buy-in, I entered into a –EV lottery with the fish. My expert assessment at the time was that my EV was about –10% of a buy-in.
>
> Later on, after having won some money, I was sitting with 4 times my initial buy-in. I gave up a lottery against the fish, which in my expert assessment at the time had about a 65% chance of netting me 3 buy-ins and about a 35% chance of losing me 4 buy-ins. Thus, I knowingly gave up a lottery with an EV of about 0.55 buy-ins – a very substantial sum, and perhaps the largest EV I have ever knowingly given up.

My seemingly irrational behavior can be explained by examining the characteristics of the game and constructing U'(chips) given U($) = $. It was obvious that the fish was either going to leave the table or lose all his money, and so I was competing not only against the other players, but also against time. The opportunities for winning a large hand against the fish were running out fast, and so it would be in my best interest to have as many chips at the table as the fish – this way, I could take all of his money in a single hand, rather than having to beat him several times. However, the maximum one can buy-in to the table for limited, and the fish had more than 5 times this amount. It would be in my best interest to try and increase my table wealth as fast as possible, and since I could always re-buy if I ran out of chips, the optimal play was to be risk-seeking for small values of table wealth.

At higher values of table wealth, the opposite is true. Having the ability to win all or most of the fish's money in a single hand, it was in my best interest to wait until my odds were good enough

[39] Cash games refer to games where the buy-in is cash, players bet and therefore win/lose cash in every hand, and players can join or leave the game at any time. This contrasts with tournament play, where a certain number of players pay cash for chips, and play until they are eliminated or win the tournament. The payoff structure in tournaments is determined before play starts.

[40] A fish is a common term for a poor player from who other players profit; a large fish is a fish who is so bad that he might be termed 'free money'

to risk losing all of my chips. If I had lost all my chips, it would have been difficult to rebuild my table wealth and it is likely that the fish would leave or go broke by the time I did. It was therefore optimal to be risk-averse at larger values of table wealth.

Consider a possible graphing of U'(chips) below. The difference between the red line (Absolute Utility) and the blue line (Table Wealth) is equal to the amount of money I would pay to be sitting at the table with that much wealth. For the graphing I assumed this was equal to 40% of my effective wealth versus the fish.[41] Note that the effective wealth used when my table wealth is less than 1 is equal to 1, because if my table wealth drops below 1, I can (and it is assumed that I do) always add additional money up to a table wealth of 1.

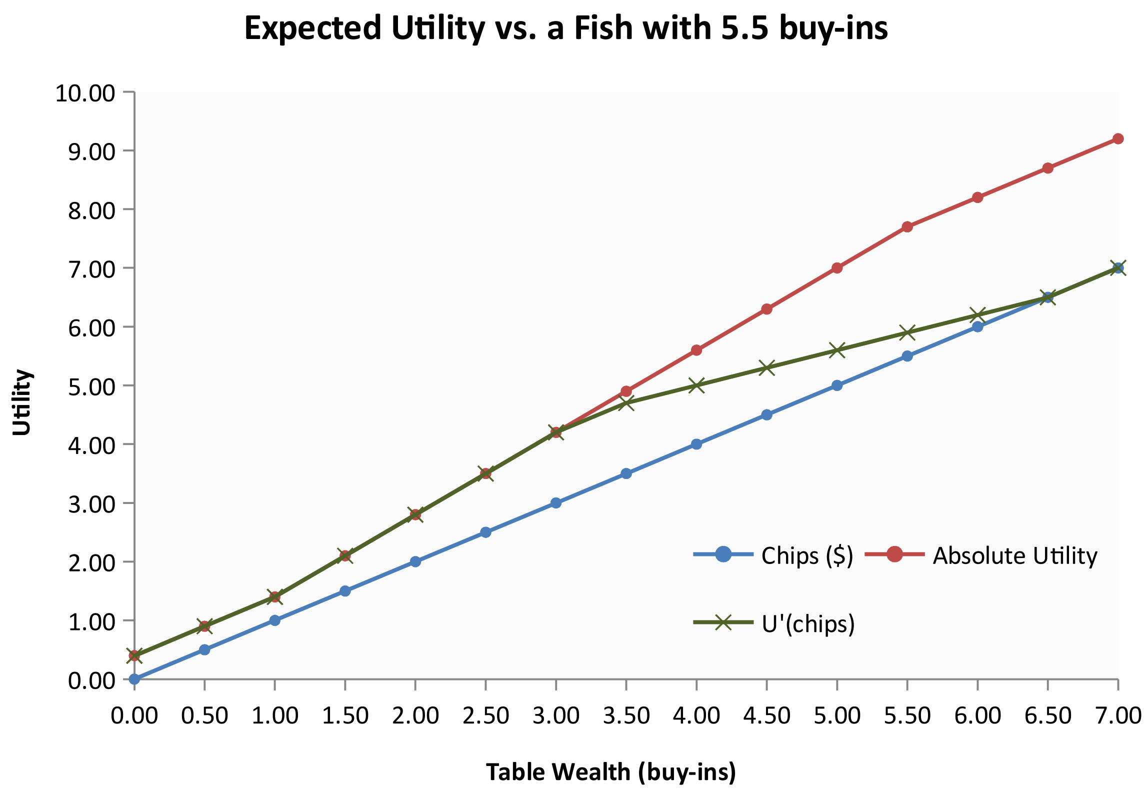

The green line represents my adjusted utility under the assumption that if my table wealth rises above 1, the fish's table wealth drops by an equal amount. In other words, I assume all money won is won from the fish, and that if my table wealth rises, the fish's table wealth falls. [42]

[41] Effective wealth refers to the shared wealth between me and the fish, or equivalently the minimum of my table wealth and fish's table wealth (assumed to be 5.5 buy-ins). The assumption that added utility = 40% of effective wealth is arbitrary. The added utility could be precisely identified through more complex mathematics, but this is beyond the scope of this discussion.

[42] This assumption is required in order to make gains in table wealth unambiguous and 1-dimensional. The real outcome space at the table has 6 dimensions, one for each player's table wealth. This therefore also assumes that

Therefore the effective wealth and added utility of sitting with the fish begins to fall if my table wealth rises above 2.75 (since this implies the fish's table wealth is now less than 2.75). This adjusted utility curve is my U'(chips).

Notice two things about U'(chips). It is convex or risk-seeking for small wealth amounts, and concave or risk-averse for larger wealth amounts. Consider my decision to turn down the lottery {3, .65; -4, .35}. This lottery, would have an expected utility of 0.65*U(0) + 0.35*U(7) = 4.69, which was less than my expected utility of not taking the lottery, or U(4) = 5. This shows that my seemingly irrational behavior was actually consistent with Expected Utility preferences.

In this scenario we were able to construct an objective VNM utility curve in the acting outcome space based solely on understand how acting outcomes (table wealth) materialized into relevant outcomes (real dollars), and knowledge of the relevant VNM utility function. Because the relevant VNM utility function was assumed, quite fairly, to be U($) = $, no elicitation of (biased) preferences was required in this construction. This suggests that objective VNM utility curves might be constructed for financial applications by assuming a VNM utility function over a relevant outcome space.

Objective VNM Utility Curves for Financial Applications

In the poker example, the relevant U(x) was known because outcomes were sufficiently small in comparison to total wealth. However, this is not always the case, and so one might consider other possible relevant U(x) functions. The chosen function should closely match one's strategic objectives over the relevant outcome space.

Consider a company operating as a going concern[43] whose only strategic objective was to maximize the long-run compounded annual growth rate (CAGR) of their assets or wealth. It has been shown that they should make decisions that maximize the geometric mean of outcomes. This is equivalent to them having logarithmic VNM utility preferences, U(x) = ln(x).[44] Similarly, any subject whose single strategic objective was long-run wealth maximization should have logarithmic utility, and this can explain why U(x) = ln(x) is "the most common and oldest utility function in the economist's toolkit".[45]

---

the table wealth of each of the other four players remains constant. This is a *fair* assumption given the circumstances, since I was virtually ignoring other players, and they were virtually ignoring me (we only cared about playing hands vs. the fish).

[43] A company that uses an infinite timeframe

[44] I need to find a proper reference for this.

[45] Moshe A. Milevsky (2009)

A very important observation in the poker example is that the optimal strategy (rational behavior) was risk-seeking for at least some decisions. This outright contradicts the assumption of risk-aversion, which is commonly perceived by economists to require no explanation[46], and is critical in certain applications of Expected Utility theory. A well-known agency cost of debt demonstrates another scenario where risk-seeking is optimal.

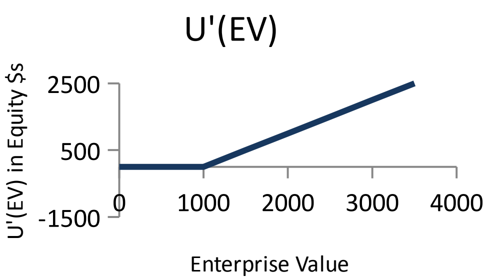

The goal of a publicly-traded company is to maximize shareholder returns. This is equivalent to maximizing the value of the equity, and because shareholders can diversify away individual company risk, U(equity) is the risk-neutral expected value function. The acting outcome space of the company is not equity however, but its enterprise value (EV). The company has debt obligations which it must pay before equity, which when combined with the limited liability of equity, produce U'(EV) = MAX(EV – debt, 0). If debt = 1000, this produces the function to the right. The extreme risk-seeking behavior over relatively low enterprise values is an agency cost of debt known as the over-investment problem.[47]

For example, if the company's management was to choose between {1000, 1} and {0, .9; 2000, .1} it would choose the latter since it maximizes U'(EV). The debt-holders bear the cost of this risk-seeking behavior, and the shareholders would benefit.

Relating Individual Risk Preferences to Consumption

Up to this point, the relevant outcome space has been monetary; however, it might be useful to think of money as the acting outcome space. If one equates the utility of money with the value it provides, one might want to consider the relevant outcome space as their consumption opportunities. That is, an individual's VNM utility function over money might be derived from their consumption goals. The outcome space consisting of all potential consumption opportunities is large and very difficult to measure, and this makes it almost impossible to apply Expected Utility directly to consumption. However, by making a set of simplifying assumptions about an individual's consumption patterns and goals, it is not only possible to construct a VNM utility curve over their finances, but it may be far more 'correct' than an elicitation, or any other method that does not consider consumption needs. Indeed, the process of eliciting preferences

[46] Footnote 3 of Rabin (2000)

[47] It is an 'agency' cost, because the company, the agent, in trying to maximize shareholder value is assuming risk-seeking behavior that is beneficial for shareholders, but a cost for debtholders, the principals.

over consumption would not be subject to the same perceptual biases as that of eliciting preferences over lotteries.

As an example, pretend that Bob was given a brief questionnaire about his consumption preferences. The utility function of his savings was assumed, and the utility of each consumption item was found through elicitation with respect to his savings (at what point would he buy each item, and in which order would he buy the items?). The following information was elicited:

- He wants to rent an apartment (as opposed to living with his parents), \$1200, U = 4
- He would like to buy a car, \$500, U = 2
- He wants to save some money, U(\$x saved) = ln(5 + x) - 2

Now, assuming that Bob consumes in such a way that maximizes his utility, we can construct Bob's VNM utility curve in the acting outcome space. With \$0, he will not be able to consume anything. If he has \$575, buying a car will outweigh the extra \$500 in savings. If Bob has \$1325, paying rent is more valuable that any combination of savings and buying a car. Once he has more than \$1775, he will be able to comfortable pay for rent and buy a car. Graphing Bob's utility at each \$25 interval of wealth produces the following curve:

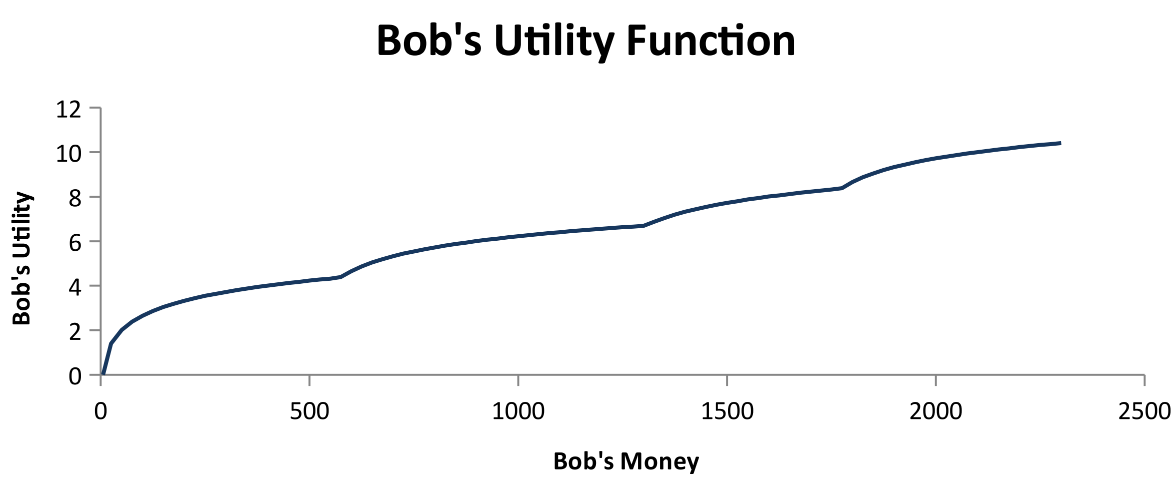

By making assumptions about how the acting space (money) impacts the relevant space (consumption), we were able to construct a VNM utility function over the acting space for Bob. Notice that Bob is risk-seeking over certain ranges of money – he is willing to gamble so that he may have a car or apartment immediately.

Although this approach seems like it might work if used properly, the drawbacks are plentiful. First, certain consumption opportunities and preferences change very quickly. In fact, the act of consumption itself will cause preferences over consumption to change. Changing income and debt capacity would also cause preferences to change. Even the oversimplified curve constructed for Bob would likely become obsolete very quickly. Second, eliciting consumption preferences has its own challenges. Although not subject to same perceptual biases as the elicitation of preferences over lotteries, there is no reason to believe that elicited consumption preferences will be equivalent to actual consumption preferences. Finding consumption preferences for wealth levels outside of a person's current wealth bracket would be highly subjective and very likely erroneous.

There is one saving grace to this idea however: retirement planning. Eliciting preferences over consumption in retirement might be less prone to error than elicitation of current preferences. Furthermore, future wealth is well defined, in that the retirement fund will be sole provider of income and debt capacity can be reasonably estimated (as opposed to Bob's case, in which it is unclear what his future income and debt capacity is). Finally, future consumption preferences are bound to be relatively stable, even over long periods of time. The VNM curve constructed for future wealth at retirement would not need to be constantly revised, as would likely be the case for Bob's utility function.

By eliciting a client's consumption preferences in retirement, a personal financial planner would be able to construct the client's VNM utility function, and manage the client's finances in such a way that maximizes their future utility. Using such a method for determining the asset allocation of a retirement fund would be far more accurate than trying to elicit a client's preferences without any reference to their consumption preferences.

6. Concluding Remarks & Further Investigation

The theory of Expected Utility is one of the oldest and most important discoveries in the field of decision-making under uncertainty. It lies at the core of several widely used financial models, and is used for a variety other prescriptive purposes. Despite the theory's numerous descriptive failures, its normative validity given a well defined and relevant outcome space is undeniable. Nevertheless, proper prescriptive application of Expected Utility is fraught with obstacles; the relevant outcome space may be difficult to quantify and may have many dimensions, and the descriptive failures of Expected Utility make direct and accurate elicitation of VNM utility curves from subjects a hopeless endeavor.

There are several scenarios in which VNM utility curves might be objectively constructed rather than elicited, and if one hopes to apply Expected Utility properly in a normative context every effort must be made to do so. The construction of VNM utility curves over acting outcome spaces by considering relevant outcomes and strategic objectives will not always be easy, but it is far preferable to relying on the interpretation of potentially irrational elicited input. Mean-variance preferences, which have been shown to be highly correlated with Expected Utility preferences, must be treated in a similar fashion.

The treatment of deriving VNM curves over acting outcome spaces was only briefly touched on, and it is likely that much progress could be made in this area. Some ideas worth further investigation:

- Are there scenarios, other than the ones mentioned, in which one knows precisely what U(x) is over the relevant outcome space?
- In the discussion the overinvestment problem, U(x) was assumed to be risk-neutral. U(x) was really just the aggregate of the utility functions of each shareholder. It would be interesting to observe how an aggregate of two or more utility functions is formed, and what the implications of this are. In Markowitz's MPT, each investor uses his preferences to pick an asset allocation, so that the aggregate asset allocation is determined by the weighted average of individual asset allocations. This has some deep implications of the overall asset allocation of the market, and the implied aggregate U(x). The market should, in theory, behave as the corporation trying to maximize its long-run CAGR. U(x) should therefore be ln(x), but it is unlikely that the implied aggregate U(x) is ln(x). I hypothesize that arbitrage opportunities do exist in this area, but are well hidden.
- Deriving U(x) from consumption preferences seems like a very attractive idea, yet it is burdened by some nasty drawbacks. It would be interesting to see what kind of work can be done on this, and how this kind of construction performs empirically.
- Is there a mathematically precise way to do the kind of outcome space and VNM curve 'transformation' that was done when going from the relevant space to the acting space? Can such a transformation always be made, and will Expected Utility always be valid in the acting space?

The topic of decision-making under uncertainty is certainly deep, and considering how critical its role in our global economy and everyday lives is, it is a wonder that it remains concealed within the depths of academia. Financial education for practitioners skips over the basic assumptions of the models, and it is unlikely that many practitioners, including the reader's own personal financial planner, truly understand the limitations and potential pitfalls involved.

Bibliography

Fonseca, Gonçalo L.. "The von Neumann-Morgenstern Expected Utility Theory." *The History of Economic Thought Website*. N.p., n.d. Web. 2 Jan. 2010. <homepage.newschool.edu/het//essays/uncert/vnmaxioms.htm>.

Kahneman, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47.2 (1979): 263-291`. *ABI/INFORM Global*. Web. 20 Nov. 2009.

Kahneman, Daniel. "Maps of Bounded Rationality: Psychology for Behavioral Economics." *The American Economic Review* 93.5 (2003): 1449-1475. *nobelprize.org*. Web. 22 Nov. 2009.

Kroll, Yoram, Haim Levy, and Harry M. Markowitz. "Mean-Variance Versus Direct Utility Maximization." *The Journal of Finance* 39.1 (1984): 47-61. *JSTOR*. Web. 20 Nov. 2009.

Levy, H, and H. M. Markowitz. "Approximating Expected Utility by a Function of Mean and Variance." *The American Economic Review* 69.3 (1979): 308-317. *JSTOR*. Web. 20 Nov. 2009.

Lichtenstein, Sarah, and Paul Slovic. "Reversals of Preference Between Bids and Choices in Gambling Decisions." *Journal of Experimental Psychology* 89.1 (1971): 46-55. *PsycARTICLES*. Web. 23 Nov. 2009.

Machina, Mark J.. "Choice Under Uncertainty: Problems Solved and Unsolved." *The Journal of Economic Perspectives* 1.1 (1987): 121-154. *ABI/INFORM Global*. Web. 27 Sept. 2009.

Machina, Mark J.. "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty." *Journal of Economic Literature* 27.4 (1989): 1622-1668. *ABI/INFORM Global*. Web. 27 Sept. 2009.

Milevsky, Moshe A.. "Utility Functions and Risk Theory." *FINE 4050, SSB* Fall (2009): 1. *SSB*. Web. 20 Oct. 2009.

Morone, Andrea. "Comparison of Mean-Variance Theory and Expected-Utility Theory through a Laboratory Experiment." *Economics Bulletin* 3.40 (2008): 1-7. *economicsbulletin.vanderbilt.edu*. Web. 30 Nov. 2009.

Rabin, Matthew. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem." *Econometrica* 68.5 (2000): 1281-1292. *ABI/INFORM Global*. Web. 22 Nov. 2009.