# *The Foundations of Expected Utility Theory*

*October 30, 2017*

## *The St. Petersburg Paradox*

We saw that there were some problems with the idea that you rationally ought to *maximize expected monetary payouts.* Here's a fun way to bring out one of the problems:

> *The St. Petersburg Paradox.* Consider the following wager: I will flip a fair coin until it comes up heads; if the first time it comes up heads is the $n^{th}$ toss then I will pay you $\$2^n$. What's the most you'd be willing to pay for this wager? What is its expected monetary payout?

This problem was first raised by Nicholas Bernoulli. It inspired Gabriel Cramer and Daniel Bernoulli (Nicholas' brother) to solve the paradox by arguing that money has declining marginal value.

| **Toss** | **Payout** ($x_i$) | **Probability** ($p_i$) |
|---:|:---:|:---:|
| $H$ | \$2 | $\frac{1}{2}$ |
| $TH$ | \$4 | $\frac{1}{4}$ |
| $TTH$ | \$8 | $\frac{1}{8}$ |
| $\vdots$ | $\vdots$ | $\vdots$ |
| $\underbrace{T\ldots TH}_{n}$ | $\$2^n$ | $\frac{1}{2^n}$ |
| $\vdots$ | $\vdots$ | $\vdots$ |

What's its expected monetary payout?

$$\sum_i p_i \cdot x_i = \frac{1}{2} \cdot \$2 + \frac{1}{4} \cdot \$4 + \frac{1}{8} \cdot \$8 + \cdots + \frac{1}{2^n} \cdot \$2^n + \ldots$$
$$= \$1 + \$1 + \$1 + \cdots + \$1 + \cdots = \infty$$

*Cramer/Bernoulli Response:* Money has declining marginal "utility" and it is *expected utility* — not expected monetary payouts — that rationality requires us to maximize.

**Money Has Declining Marginal Utility:** If $x > y$, the difference in value between having $\$x$ and having $\$\left(\frac{x+y}{2}\right)$ is greater than the difference in value between having $\$\left(\frac{x+y}{2}\right)$ and having $\$y$.

Money has declining marginal utility, for example, if $2u(\$x) > u(\$2x)$.

If $2u(\$x) > u(\$2x)$, then $2u(\$2^n) > u(\$2^{n+1})$.

And, because $\sum_n a_n$ converges if, for all $n$, $\frac{a_{n+1}}{a_n} < 1$, the expected utility of the St. Petersburg wager ($= \sum_n \frac{1}{2^n} \cdot u(\$2^n)$) converges to a finite amount.

## *Subjective Expected Utility Theory*

Daniel Bernoulli proposed that utility is a logarithmic function of money (e.g., $u(x) = \log(x)$), but why think that utility is *objective*? Can't different people value money in different ways? And don't we value things other than money?

*Proposal:* Rationality requires you to maximize the expectation of your own *subjective* utility function.

*Worry:* What is your subjective utility function like? Can you introspect what precise utility you assign to various outcomes?

## *Von Neumann & Morgenstern*

John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior*, Princeton University Press, 1944.

Von Neumann & Morgenstern (vNM) proved that, if your preferences meet certain constraints, those preferences can be represented by a unique (up to positive linear transformation) utility-function $u$ such that you maximize expected utility relative to $u$.

A utility-function $u^*$ is a *positive linear transformation* of $u$ just in case, for all $x$, $u(x) = \alpha \cdot u^*(X) + \beta$, for any real numbers $\beta$ and $\alpha > 0$. The family of functions form an *interval scale*, which encodes information about the ratios of differences in what's being measured (e.g., Fahrenheit and Celsius measure temperature on an interval scale): $\frac{u(X)-u(Y)}{u(Z)-u(W)} = \frac{u^*(X)-u^*(Y)}{u^*(Z)-u^*(W)}$.

*vNM Utility Theory Hypothesis.* You are instrumentally rational just in case there is a utility-function $u$ such that, for any lotteries, $L = \{\langle p_1, x_1\rangle, \langle p_2, x_2\rangle, \dots\}$ and $L^* = \{\langle p_1^*, x_1^*\rangle, \langle p_2^*, x_2^*\rangle, \dots\}$, you prefer $L$ to $L^*$ if and only if the expected utility of $L$ is greater than the expected utility of $L^*$.

$$L \succ L^* \text{ if and only if } \sum_i p_i \cdot u(x_i) > \sum_i p_i^* \cdot u(x_i^*)$$

In addition, the proof requires the set of lotteries to be closed under probability mixtures (i.e., for any lotteries, $L_i$, $L_j$, and $p \in [0,1]$, there is the lottery: $\{\langle p, L_i\rangle, \langle 1-p, L_j\rangle\}$), and that compound lotteries can be reduced to simple lotteries.

What do these constraints on your preferences look like?

***vNM 1 (Transitivity)*** For any lotteries, $L_1$, $L_2$, $L_3$, if $L_1 \succ L_2$, and $L_2 \succ L_3$, then $L_1 \succ L_3$.

***vNM 2 (Completeness)*** For any lotteries, $L_1$, $L_2$, either $L_1 \succ L_2$, $L_2 \succ L_1$, or $L_1 \approx L_2$.

***vNM 3 (Continuity)*** If $L_1 \succ L_2 \succ L_3$, then there is a real number $p \in [0,1]$ such that $\{\langle p, L_1\rangle, \langle 1-p, L_3\rangle\} \approx L_2$.

***vNM 4 (Independence)*** For any lottery $L_3$ and any $p \in [0,1]$, $L_1 \succ L_2$ if and only if $\{\langle p, L_1\rangle, \langle 1-p, L_3\rangle\} \succ \{\langle p, L_2\rangle, \langle 1-p, L_3\rangle\}$.

**vNM 1** and **vNM 2** are necessary for your preferences to be represented by an *ordinal* utility-function.

**vNM 3** rules out there being infinitely good and infinitely bad outcomes and ensures that your evaluations of lotteries are appropriately sensitive their probabilities.

**vNM 4** says that your evaluation of two lotteries should only depend on the ways in which they differ. It ensures that the utility-function is *additively separable.*

## *Savage's Subjective Expected Utility Theory*

VNM was subjective in one way, but objective in another: utilities are up to the agent, but the probabilities are taken to be *objective*. Their theorem showed how to derive your subjective utility-function from your preferences over *lotteries with given probabilities*.

Leonard Savage showed how to derive both your utility-function and a *subjective probability function* — all at once — from your preferences over gambles (or "acts"). Savage's framework characterizes *decision problems* in terms of three entities:

Leonard J. Savage, *The Foundations of Statistics,* Wiley Publishing. 1954

1. ***Outcomes.*** Each act, $f$, and state, $s$, determines a unique outcome, $O[f, s]$, that describes all the desirable and undesirable things that the combination of $f$ and $s$ would bring about.
2. ***States.*** A state, $s$, is a complete description of how, for all you know, the world might be. We will call disjunctions of states *events*.
3. ***Acts.*** Savage defines acts as functions from states to outcomes. (Savage assumes that, for each outcome, there is a *constant act* that produces that outcome in every state.)

Outcomes are the objects of your non-instrumental preferences, measured by a subjective utility-function, $u$.

Your subjective probability function, $Cr$ (for "credence"), is defined over these states/events.

Acts are the objects of your instrumental preferences, measured by a subjective expectational utility-function.

DECISION MATRIX

| | $s_1$ | $s_2$ | ... | $s_n$ |
|---|---|---|---|---|
| $\mathbf{f_1}$ | $O[\mathbf{f_1}, s_1]$ | $O[\mathbf{f_1}, s_2]$ | ... | $O[\mathbf{f_1}, s_n]$ |
| $\mathbf{f_2}$ | $O[\mathbf{f_2}, s_1]$ | $O[\mathbf{f_2}, s_2]$ | ... | $O[\mathbf{f_2}, s_n]$ |
| $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ |
| $\mathbf{f_k}$ | $O[\mathbf{f_k}, s_1]$ | $O[\mathbf{f_k}, s_2]$ | ... | $O[\mathbf{f_k}, s_n]$ |

Savage's Representation Theorem shows that if your preferences obey his postulates, you will evaluate acts *as if* you are maximizing expected utility, relative to a unique $Cr$ (a probability function representing your subjective degrees of belief) and a unique (up to positive linear transformation) utility-function $u$.

***Savage's Equation*** $U(f) = \sum_s Cr(s) \cdot u\left(f(s)\right)$

Many of Savage's postulates are analogous to vNM's, so let's only look at two:

***Sure-Thing Principle*** If $f$, $g$, and $f^*$, $g^*$, are such that

(i) for all $s \in \neg E$, $f(s) = g(s)$ and $f^*(s) = g^*(s)$,

(ii) for all $s \in E$, $f(s) = f^*(s)$ and $g(s) = g^*(s)$,

Then $f \succ g$ if and only if $f^* \succ g^*$.

SURE-THING PRINCIPLE

| | $E$ | $\neg E$ |
|---|---|---|
| **f** | $X$ | $Z$ |
| **g** | $Y$ | $Z$ |
| **f**$^*$ | $X$ | $Z^*$ |
| **g**$^*$ | $Y$ | $Z^*$ |

$f \succ g$ if and only if $f^* \succ g^*$

***Stochastic Dominance*** Let $f_1$, $f_2$, $g_1$, and $g_2$ be such that

$f_1$ For all $s \in E$, $f_1(s) = X$ and, for all $s \in \neg E$, $f_1(s) = X^*$

$f_2$ For all $s \in F$, $f_2(s) = X$ and, for all $s \in \neg F$, $f_2(s) = X^*$

$g_1$ For all $s \in E$, $g_1(s) = Y$ and, for all $s \in \neg E$, $g_1(s) = Y^*$

$g_2$ For all $s \in F$, $g_2(s) = Y$ and, for all $s \in \neg F$, $g_2(s) = Y^*$

Then, if $X \succ X^*$ and $Y \succ Y^*$, $f_1 \succ f_2$ if and only if $g_1 \succ g_2$.

STOCHASTIC DOMINANCE

| | $E$ | $\neg E$ |
|---|---|---|
| $\mathbf{f_1}$ | $X$ | $X^*$ |
| $\mathbf{g_1}$ | $Y$ | $Y^*$ |

| | $F$ | $\neg F$ |
|---|---|---|
| $\mathbf{f_2}$ | $X$ | $X^*$ |
| $\mathbf{g_2}$ | $Y$ | $Y^*$ |

If $X \succ X^*$ and $Y \succ Y^*$, $f_1 \succ f_2$ if and only if $g_1 \succ g_2$

The Sure-Thing Principle is the analogue of vNM's Independence. Stochastic Dominance is what allows Savage to derive a comparative belief relation (e.g., "$E$ is *likelier than* $F$") from your preferences.

## *Problems with Savage's Framework*

1. *Constant Acts are Implausible.* Savage's framework requires there to be, for every outcome, an act that results in that outcome in each state. But in order for his postulates to be plausible, outcomes must be very specific. There is a tension here.

*Example*: Can there be a constant act that pays out "a carefree day spent basking in the sun at the park" in every state — including states in which it rains? states in which there's nuclear war? etc.

2. *Small World vs Grand World Decision Problems.* We can specify a decision-problem by partitioning the events/acts more or less finely. Savage's framework requires us to partition them as finely as possible (Grand World). But this is unrealistic. And there is no guarantee that what Savage's theory recommends in more realistic decision-problems (Small World) will align with what it would recommend in the corresponding Grand World version.
3. *Independence of Acts and States?* Consider the following decision problem:

THE BIG TEST. You have an important test tomorrow. You'd very much like to pass the test rather than fail it. Tonight, you have two options: you can *Study* or you can *Goof*. All else equal, you prefer goofing to studying. What should you do?

| | PASS | FAIL |
|---|---|---|
| *Study* | 20 | 0 |
| *Goof* | 25 | 5 |

$$\sum_E Cr(E) \cdot u\,(Study(E)) = Cr\,(\text{PASS}) \cdot 20 + Cr\,(\text{FAIL}) \cdot 0 = Cr\,(\text{PASS}) \cdot 20$$

$$\sum_E Cr(E) \cdot u\,(Goof(E)) = Cr\,(\text{PASS}) \cdot 25 + Cr\,(\text{FAIL}) \cdot 5 = Cr\,(\text{PASS}) \cdot 20 + 5$$

So, $U\,(Goof) > U\,(Study)$. But that can't be right. It isn't always irrational to study! What's gone wrong here?

## *Jeffery's Evidential Decision Theory*

Our actions sometimes affect how likely it is for the world to be in a particular state. For example, studying makes it more likely that you'll pass the test, and goofing makes it more likely that you'll fail it.

*Jeffrey's Idea:* You should evaluate your actions on the supposition that you perform them.

***Jeffrey's Equation*** $V(A) = \sum_S Cr\,(S \mid A) \cdot V\,(A \wedge S)$

Suppose you think that if you study, you're 80% likely to pass and if you don't you're 80% likely to fail. Then Jeffery's Equation says that you should prefer studying over goofing:

$$Cr\,(\text{PASS} \mid Study) = .8$$
$$Cr\,(\text{FAIL} \mid Study) = .2$$
$$Cr\,(\text{FAIL} \mid Goof) = .8$$
$$Cr\,(\text{PASS} \mid Goof) = .2$$

$$\begin{aligned} V(Study) &= \sum_S Cr\,(S \mid Study) \cdot V\,(Study \wedge S) \\ &= Cr\,(\text{PASS} \mid Study) \cdot 20 + Cr\,(\text{FAIL} \mid Study) \cdot 0 \\ &= .8 \cdot 20 + .2 \cdot 0 = 16 \\ V(Goof) &= \sum_S Cr\,(S \mid Goof) \cdot V\,(Goof \wedge S) \\ &= Cr\,(\text{PASS} \mid Goof) \cdot 25 + Cr\,(\text{FAIL} \mid Goof) \cdot 5 \\ &= .2 \cdot 25 + .8 \cdot 5 = 9 \end{aligned}$$

$V(A)$ is $A$'s "news value": it measures the extent to which you'd welcome the news that $A$ is true; i.e., $V(A)$ measures the average extent to which learning $A$ would provide you with evidence that good things are to come.

1. *Generality.* Jeffrey's proposal applies to any propositions whatsoever. Propositions describing your actions are only a special case.
2. *Simplicity.* It doesn't need to distinguish between *value* and *expected value*, and so doesn't require the problematic structural framework of Savage's (e.g., outcomes, states, acts, constant acts, etc.)
3. *Partition Invariant.* It doesn't require us to set-up decision problems so that the states are independent of the actions. It doesn't have the Small World/Grand World problem.

There is also a representation theorem for Jeffery's proposal (although it arrives at a weaker conclusion). The key constraints are:

***Averaging*** If $X$ and $Y$ are mutually incompatible, then $X \succ Y$ if and only if $X \succ (X \vee Y) \succ Y$.

***Impartiality*** Suppose that $X \succ X^*$ and that $Y$ and $Y^*$ are incompatible with each. Then you don't have the following preferences:

$$Y \succ (X^* \vee Y) \succ (X \vee Y) \succ X \succ X^*$$
$$X \succ X^* \succ (X^* \vee Y^*) \succ (X \vee Y^*) \succ Y^*$$

Averaging is analogous to Independence (and the Sure-Thing Principle). Impartiality is analogous to Stochastic Dominance.