Risk and Uncertainty: A Fallacy of Large Numbers (Reprint)

by Paul A. Samuelson

RISK AND UNCERTAINTY: A FALLACY OF LARGE NUMBERS[1]

Experience shows that while a single event may have a probability spread, a large repetition of independent single events gives a greater approach toward certainty. This corresponds to the mathematically provable Law of Large Numbers of James Bernoulli. This valid property of large numbers is often given an invalid interpretation. Thus people say an insurance company reduces its risk by increasing the number of ships it insures. Or they refuse to accept a mathematically favorable bet, but agree to a large enough repetition of such bets: e. g., believing it is almost a sure thing that there will be a million heads when two million symmetric coins are tossed even though it is highly uncertain there will be one head out of two coins tossed. The correct relationship (that an insurer reduces total risk by *subdividing*) is pointed out and a strong theorem is proved: that a person whose utility schedule prevents him from ever taking a specific favorable bet when offered only once can never rationally take a large sequence of such fair bets, if expected utility is maximized. The intransitivity of alternative decision criteria-such as selecting out of any two situations that one which will more probably leave you better off-is also demonstrated.

1. INTRODUCTION. - « There is safety in numbers. » So people tell one. But is there? And in what possible sense?

The issue is of some importance for economic behavior. Is it true that an insurance company *reduces* its risk by *doubling* the number of ships it insures? Can one distinguish between risk and uncertainty by supposing that the former can count on some remorseless cancelling out of actuarial risks?

To throw light on a facet of this problem, I shall formulate and prove a theorem that should dispell one fallacy of wide currency.

2. A TEST OF VALOR. - S. Ulam, already a distinguished mathematician when we were Junior Fellows together at Harvard a quarter century ago, once said: « I define a coward as someone who will not bet when you offer him two-to-one odds and let him choose *his* side.»

With the centuries-old St. Petersburg Paradox in my mind, I pedantically corrected him: « You mean will not make a *sufficiently small* bet (so that the change in the marginal utility of money[2] will not contaminate his choice). »

3. A GUINEA PIG SPEAKS. - Recalling this conversation, a few years ago I offered some lunch colleagues to bet each \$200 to \$100 that the side of a coin *they* specified would not appear at the first toss. One distingui-

[1] See the article of M. B. De Finetti, "La decisione nell' incertezza," in *Scientia*, April-May 1963, p. 61.

[2] I might have quibbled that the chap could have a corner in his Bernoulli-Ramsey-Neumann utility function at his initial point, and thus escape the charge of cowardice or (even worse) irrationality. This, however, would have been a quibble since Ulam could move him from the corner by giving him a dollar and then test his "courage." As for the "St. Petersburg Paradox," see footnote 2, Section 5.

shed scholar - who lays no claim to advanced mathematical skills - gave the following answer:

« I won't bet because I would feel the $100 loss more than the $200 gain. But I'll take you on if you promise to let me make 100 such bets».

What was behind this interesting answer? He, and many others, have given something like the following explanation. « One toss is not enough to make it reasonably sure that the law of averages will turn out in my favor. But in a hundred tosses of a coin, the law of large numbers will make it a darn good bet. I am, so to speak, virtually sure to come out ahead in such a sequence, and that is why I accept the sequence while rejecting the single toss. »

4. *Maximum loss and probable loss.* - What are we to think about this answer? Here are a few observations.

a) If it hurts much to lose $100, it must certainly hurt to lose 100 x $100 = $10,000. Yet there is a distinct *possibility* of so extreme a loss. Granted that the probability of so long a run of repetitions is, by most numerical calculations, extremely low: less than 1 in a million (or $1/2^{100}$), still, if a person is already at the very minimum of subsistence, with a marginal utility of income that becomes practically infinite for any loss, he might act like a minimaxer[1] and eschew options that could involve any losses at all. [Note: increasing the sequence from n = 100 to n = 1,000 or $n \to \infty$, will obviously not tempt such a minimaxer - even though the probability of any loss becomes gigantically tiny].

b) Shifting your focus from the maximum possible loss (which grows in full proportion to the length of the sequence), you may calculate the probability of making no loss at all. For the single toss, it is of course one-half. For 100 tosses, it is the probability of getting 34 or more correct heads (or, alternatively, tails) in 100 tosses. By the usual binomial calculation and normal approximation,[2] this probability of making a gain is found to be very large, $P_{100} = .99+$. If this has not reduced the probability of a loss by enough, it is evident that by increasing n from 100 to some larger number will succeed in reducing the probability of a loss to as low as you want to prescribe in advance.

c) Indeed, James Bernoulli's so-called Law of Large Numbers guarantees you this: « Suppose I offer you favorable odds at each toss so that your mathematical expectation of gain is k per cent in terms of the money you put at risk in each toss. Then you can choose a long-enough sequence of tosses to make the probability as near as you like to one that your earnings will be indefinitely near k per cent return on the total money you put at risk ».

[1] In the literature of statistical decision making, a minimaxer is defined as one who acts so as to insure that his maximum possible loss is at a minimum.

[2] I assume the coin is a reasonably new one. If it has developed some bias toward landing on one side, and if prior experimentation leads you to prefer one side to bet on, you can hope to do even better than as given above. Note: for definiteness I assume that when you decide to bet on a sequence of tosses, you are held to the full contract and cannot opt out in midstream; nor can you learn the coin's bias in the early tosses, since you are told immediately the result of your 100-toss play.

5. Irrationality of compounding a mistake. - The « virtual certainty » of making a large gain must at first glance seem a powerful argument in favor of the decision to contract for a long sequence of favorable bets. But should it be, when we recall that virtual certainty cannot be complete certainty and realize that the improbable loss will be very great indeed if it does occur?

If a person is concerned with maximizing the expected or average value of the utility of all possible outcomes[1] and my colleague assures me that he wants to stand with Daniel Bernoulli, Bentham, Ramsey, v. Neumann, Marschak, and Savage on this basic issue - it is simply not sufficient to look at the probability of a gain alone. *Each outcome must have its utility reckoned at the appropriate probability; and when this is done it will be found that no sequence is acceptable if each of its single plays is not acceptable.* This is a basic theorem.

One dramatic way of seeing this is to go back to the St. Petersburg Paradox itself. No matter how high a price my colleague agreed to pay to engage in this classic game, the probability will approach one that he will come out as much ahead as he cares to specify in advance.[2]

6. An alternative axiom system of maximizing probabilities. No slave can serve two independent masters. If one is an expected-utility-maximizer he cannot generally be a maximizer of the probability of some gain. However, economists ought to give serious attention to the merits of various alternative axiom systems. Here is one that, at first glance, has superficial attractiveness.

Axiom: In choosing between two decisions, *A* and *B*, select that one which will more probably leave you better off. *I.e.*, select *A* over *B* if it is more probable that the gain given by *A* is larger than that on *B*, or, in formulae:

$$\text{Prob}\left\{ \text{A's gain} > \text{B's gain} \right\} > \frac{1}{2}$$

[abbreviate the above to A > B].

Similarly with respect to any pair of (A, B, C, D, ...).

In terms of the above system, call *A* agreeing to bet on one toss; B deciding not to toss at all; and C agreeing to a long sequence of tosses. Then clearly,

$$A = B, \; C > B, \; C > A \, .$$

So my friend's decision to accept the long sequence turns out to agree with this axiom system. However, if D is the decision to accept a sequence of two tosses, my friend said he would not undertake it; and yet, in this

[1] *I. e.*, he acts to maximize $U = p_1 U_1 + p_2 U_2 + \ldots + p_n U_n$, where U_i represents the utility of each possible outcome and p_i represents its respective probability.

[2] The « Paradox » (Daniel Bernoulli, St. Petersburg, 1738) says, that turnig a coin until head appears for the first time, and to get $1, or 2, 4, ..., 2^{n-1}, ... according to the number of turns required, is a favorable bet no matter how large the amount to be paid for it. To avoid such a paradox, D. Bernoulli suggested dealing with the utilities rather than with money values (that is, with a concave scale with diminishing increments). To get rid of any initial infinity in the problem, see the modified sequence of finite tosses for the Petersburg situation in P. A. Samuelson, *The St. Petersburg Paradox as a Divergent Double Limit* « International Economic Review » Vol. 1, N. 1, January, 1960), pp. 31-37.

system, D > B. Moreover, call E the decision to accept the following bet: you win a million dollars with probability .51 but lose a million with probability .49. Few could accept such a bet; and of those who could, few would. Yet in this axiom system E > B.

There is a further fatal objection to this axiom system. It need not satisfy transitivity relations among 3 or more choices. Thus, it is quite possible to have X > Y, Y > Z and Z > X.

One example is enough to show this pathological possibility. Let X be a situation that is a shade more likely to give you a small gain rather than a large loss. By this axiom system you will prefer it to the Situation Y, which gives you no chance of a gain or loss. And you will prefer Y to Situation Z, which makes it a shade more likely that you will receive a small loss rather than a large gain. But now let us compare Z and X. Instead of acting transitively, you will prefer Z to X for the simple reason that Z will give you the better outcome in every situation except the one in which simultaneously the respective outcomes would be the small gain and the small loss, a compound event whose probability is not much more than about one-quarter (equal to the product of two independent probabilities that are respectively just above one-half).

7. PROOF THAT UNFAIRNESS CAN ONLY BREED UNFAIRNESS. - After the above digression, there remains the task to prove the basic theorem already enunciated.

Theorem. If at each income or wealth level within a range, the expected utility of a certain investment or bet is worse than abstention, then no sequence of such independent ventures (that leaves one within the specified range of income) can have a favorable expected utility.

Thus, if you would always refuse to take favorable odds on a single toss, you must rationally refuse to participate in any (finite) sequence of such tosses.

The logic of the proof can be briefly indicated. If you will not accept one toss, you cannot accept two - since the latter could be thought of as consisting of the (unwise) decision to accept one plus the open decision to accept a second. Even if you were stuck with the first outcome, you would cut your further (utility) losses and refuse the terminal throw. By extending the reasoning from 2 to 3 = 2 + 1, ..., and from n-1 to n, we rule out any sequence at all.[1]

[1] Mathematically, if you start at a known utility U_t, the probability of ending after one venture with at least U_{t+1} can be written as $F(U_{t+1}, U_t)$. By hypothesis, in the utility metric each toss is an unfair game (even though it may be more than fair game in the money metric). Or

$$E(U_{t+1}/U_t) = \int_{-\infty}^{\infty} U_{t+1}\, dF(U_{t+1}, U_t)\, U < t.$$

It is an easy theorem that repeated (identical and independent) fair games yield a fair game; and repeated unfair games yield an unfair game. Specifically, the probability of getting at least $U_{t+k} = X$, after starting out with $U_t = Y$ and playing a sequence of k games, is given by
$F_k(X, Y) = F(X, Y)^* F_{k-1}(X, Y) = \ldots = F(X, Y)^* F(X, Y)^* \ldots {}^* F(X, Y)$, where $F(X, Y)^* G(X, Y)$ is the integral $\int_{-\infty}^{\infty} F(X, S)\, dG(S, Y)$. And, if $\int_{-\infty}^{\infty} X dF(X, Y) < Y$ then necessarily $\int_{-\infty}^{\infty} X\, dF_1(X, Y) < Y$ and $\ldots \int_{\infty}^{\infty} X\, dF_k(X, Y) < Y$.

8. Conclusions. - Now that I have demonstrated the fallacy that there is safety in numbers - that actuarial risks must allegedly cancel out in the sense relevant for investment decisions - a few general remarks may be in order.

Firstly, when an insurance company doubles the number of ships it insures, it does also double the range of its possible losses or gains. (This does not deny that it reduces the probability of its losses.) If at the same time that it doubles the pool of its risks, it doubles the number of its owners, it has indeed left the maximum possible loss per owner unchanged; but - and this is the germ of truth in the expression « there is safety in numbers » - the insurance company has now succeeded in reducing the probability of each loss; the gain to each owner now becomes a more certain one.

In short, it is not so much by *adding* new risks as by *subdividing* risks among more people that insurance companies reduce the risk of each. To see this, do not double or change at all the original number of ships insured by the company: but let each owner sell half his shares to each new owner. Then the risk of loss to each owner per dollar now in the company will have indeed been reduced.

Undoubtedly this is what my colleague really had in mind. In refusing a bet of $100 against $200, he should not then have specified a sequence of 100 such bets. That is adding risks. He should have asked to subdivide the risk and asked for a sequence of 100 bets, each of which was 100th as big (or $1 against $2). If the *money* odds are favorable and if we can subdivide the bets enough, any expected-utility-maximizer can be coaxed into a favorable-odds bet - for the obvious reason that the utility function's curvature becomes more and more negligible in a sufficiently limited range around any initial position. For sufficiently small bets we get more-than-a-fair game in the utility space, and my basic theorem goes nicely into reverse.[1]

Secondly, and finally, some economists have tried to distinguish between risk and uncertainty in the belief that actuarial probabilities can reduce risk to « virtual » certainty. The limit laws of probability grind fine but they do not grind that exceeding fine. I suspect there is often confusion between two similar-sounding situations. One is the case where the owner of a lottery has sold out *all* the tickets; the buyers of the tickets then face some kind of risky uncertainty, but the owner has completely cancelled out his risks whatever the draw may show - which is not a case of risk as against uncertainty, but really reflects a case of certainty without any risks at all. Another case is that in which the management of Monte Carlo or of the « numbers game » do business with their customers. The management makes sure that the odds are in their favor; but they can never make *sure* that a run of luck will not go against them and break the house (even though they can reduce this probability of ruin to a *positive* fraction).

In every actuarial situation of mathematical probability, no matter

[1] Cf. my cited 1960 paper. I should warn against undue extrapolation of my theorem. It does not say one must always refuse a sequence if one refuses a single venture: if, at higher income levels the single tosses become acceptable, and at lower levels the penalty of losses does not become infinite, there might well be a long sequence that is optional.

how large the numbers in the sample, we are left with a finite sample: in the appropriate limit law of probability there will necessarily be left an epsilon of uncertainty even in so-called risk situations. As Gertrude Stein never said: Epsilon ain't zero. This virtual remark has great importance for the attempt to create a difference of kind between risk and uncertainty in the economics of investment and decision-making.

P. A. SAMUELSON

Cambridge, Massachusetts Institute of Technology.