Samuelson's Fallacy of Large Numbers With Decreasing Absolute Risk Aversion

Whelan, Karl

University College Dublin

4 July 2024

Samuelson's Fallacy of Large Numbers With Decreasing Absolute Risk Aversion

Karl Whelan*

University College Dublin

July 2024

Abstract

Samuelson (1963) conjectured that accepting multiple independent gambles you would reject on a stand-alone basis violated expected utility theory. Ross (1999) and others presented examples where expected utility maximizers would accept multiple gambles that would be rejected on a stand-alone basis once the number of gambles gets large enough. We show that a stronger result than Samuelson's conjecture applies for DARA preferences over wealth. Expected utility maximizers with DARA preferences have threshold levels of wealth such that those above the threshold will accept N positive expected value gambles while those below will not and these thresholds are increasing with N.

Keywords: Risk aversion; Paul Samuelson; Law of large numbers

JEL Classification: D81

*karl.whelan@ucd.ie. Thanks to Paul Devereux and Margaret Samahita for assistance.

1. Introduction

Life regularly offers individuals and businesses opportunities that are on average profitable but that risk losses. In a famous paper, Paul Samuelson (1963) discussed the question of how people should behave when offered multiple simultaneous such opportunities. Samuelson described a conversation in which he "*offered some lunch colleagues to gamble each $200 to $100 that the side of a coin they specified would not appear at the first toss. One distinguished scholar—who lays no claim to advanced mathematical skills—gave the following answer: I won't gamble 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 gambles.*" Their explanation was "*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.*"

Samuelson believed his colleague was engaging in a "fallacy of large numbers." He argued that an expected utility maximizer who turns down a gamble as a stand-alone proposition should also reject an offer of multiple such gambles. Samuelson used an inductive argument: "*If you will not accept one toss, you cannot accept two – since the latter would 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 my sequence at all.*"

This paper describes the optimal decision rules for accepting or rejecting multiple independent identical gambles of an expected utility maximizer with decreasing absolute risk aversion (DARA) preferences over wealth. DARA is a standard assumption that includes the constant relative risk aversion (CRRA) utility functions typically used in modern economic theory. It has empirical support from experimental evidence (Levy, 1994, Huber et al, 2023) and from studies based on financial asset holdings (Chiappori and Paiella, 2011, Calvet and Sodini, 2014, Meeuwis, 2022).

Samuelson included a qualification in a footnote—"*I should warn against undue extrapolation of my theorem. It does not say that one must always refuse a sequence if one refuses a single venture: if, at higher income levels the single losses become acceptable, and at lower levels the penalty of losses does not become infinite, there might well be a long sequence that it is optimal*." Here, Samuelson was suggesting that declining risk aversion might over-turn his result.

We show that DARA preferences do not over-turn Samuelson's "theorem" and are actually consistent with a stronger version of it. Specifically, we show that for once-off gambles that might be accepted at some wealth level, DARA preferences imply a threshold level of wealth W_N^* such that only those with $W \geq W_N^*$ will accept N independent such gambles and this threshold is increasing with N. Increasing the number of gambles makes it less likely that people with DARA preferences will accept an offer, so some people will accept one gamble but not two, while others will accept two gambles but not three and so on. We prove this analytically for local gambles and illustrate it numerically by showing how various multiple gambles will be considered by people with CRRA preferences.

Our results run counter to those presented by Nielsen (1985), Hellwig (1995) and Ross (1999). These papers provided conditions (with Ross's being the most general) under which expected utility maximizers who turned down a once-off gamble with a positive expected value would accept N of them once N was sufficiently large, which Ross termed the "eventual acceptance property". Our paper differs in assuming DARA preferences, which Ross noted did not satisfy the conditions he outlined for an eventual acceptance property to hold. Our numerical calculations also differ in ruling out gambles that could lead to negative values of wealth. We do not consider this possibility because of the difficulty in interpreting negative wealth and also because the CRRA utility functions that we use for illustration are not defined for negative wealth.

The paper is structured as follows. Section 2 presents the theoretical results on local gambles. Section 3 contains our numerical calculations. Section 4 discusses our findings.

2. Theory

Pratt (1964) defined utility functions as exhibiting DARA if $r'(W) < 0$ where

$$r(W) = -\frac{U''(W)}{U'(W)} \tag{1}$$

Dybvig and Lippman (1983) showed DARA preferences implied that for any risky gamble with a positive expected value, there is a threshold level of wealth such that those with wealth above this level will accept the gamble and those below it will not. Here, we show how the thresholds for accepting multiple gambles can depend positively on the number of gambles.

Ross (1999) described how the conditions for eventually accepting an offer of multiple positive expected value gambles as the number of gambles gets large required a lower bound on utility losses from losing gambles and that utility gained from winning gambles was unbounded. The following simple result, however, shows that increasing the number of positive expected value gambles does not generally make offers more attractive for people with concave utility.

Proposition 1: Let $U(W)$ be a twice-differentiable concave utility function of wealth. For positive expected value gambles small enough that second-order Taylor series approximations to the utility function around initial wealth accurately describe utility outcomes from accepting multiple gambles, an expected utility maximizer with utility function $U(W)$ that is indifferent between accepting or rejecting an offer of N independent positive expected value gambles will reject offers of $N+k$ gambles where $k > 0$.

Proof. Consider a gamble defined by a profit X that has a mean of $\mu > 0$ and a variance of σ^2.

Expected utility for someone with initial wealth W from taking one of these gambles can be written as

$$E\left[U\left(W+X\right)\right]=U\left(W\right)+U'\left(W\right)E\left(X\right)+\frac{U''\left(W\right)}{2}E\left(X^{2}\right) \tag{2}$$

Using the relationship between the second moment and the variance

$$E\left(X^{2}\right)=\sigma^{2}+\mu^{2} \tag{3}$$

this can be written as

$$E\left[U\left(W+X\right)\right]=U\left(W\right)+U'\left(W\right)\mu+\frac{U''\left(W\right)}{2}\left(\sigma^{2}+\mu^{2}\right) \tag{4}$$

Denote the outcome from accepting N independent gambles of this type as X_N. This has mean $N\mu$ and variance $N\sigma$. The expected utility from accepting this offer is

$$E\left[U\left(W+X_N\right)\right]=U\left(W\right)+U'\left(W\right)N\mu+\frac{U''\left(W\right)}{2}\left(N\sigma^{2}+N^{2}\mu^{2}\right) \tag{5}$$

Now consider someone with wealth W_N^* such that they are indifferent between accepting or rejecting an offer of N independent gambles, so $E\left[U\left(W_N^*+X_N\right)\right]=U\left(W_N^*\right)$. This means

$$U'\left(W_N^*\right)N\mu+\frac{U''\left(W_N^*\right)}{2}\left(N\sigma^{2}+N^{2}\mu^{2}\right)=0 \tag{6}$$

So this person's coefficient of absolute risk aversion is

$$r\left(W_N^*\right)=-\frac{U''\left(W_N^*\right)}{U'\left(W_N^*\right)}=\frac{2\mu}{\sigma^{2}+N\mu^{2}} \tag{7}$$

Their expected utility from accepting $N+k$ gambles would be

$$\begin{aligned} E\left[U\left(W_N^*+X_{N+k}\right)\right] &= U\left(W_N^*\right)+U'\left(W_N^*\right)\left(N+k\right)\mu+\frac{U''\left(W_N^*\right)}{2}\left(\left(N+k\right)\sigma^{2}+\left(N+k\right)^{2}\mu^{2}\right) \\ &= U\left(W_N^*\right)+U'\left(W_N^*\right)\left(N+k\right)\left[\mu-\frac{r\left(W_N^*\right)}{2}\left(\sigma^{2}+\left(N+k\right)\mu^{2}\right)\right] \end{aligned}$$

Substituting in the formula for the coefficient of absolute risk aversion from equation 7, this becomes

$$E\left[U\left(W_N^*+X_{N+k}\right)\right]=U\left(W_N^*\right)+U'\left(W_N^*\right)\left(N+k\right)\mu\left[1-\frac{\sigma^{2}+\left(N+k\right)\mu^{2}}{\sigma^{2}+N\mu^{2}}\right]<U\left(W_N^*\right) \tag{8}$$

which is the required result. □

This result holds because, relative to accepting N gambles, the positive first-moment term due to accepting $N+k$ gambles increases by a factor of $\frac{N+k}{N}$ while the negative second-moment term increases by more. The result does not imply that all expected utility maximizers with concave utility who would reject a single gamble should reject two such gambles. Tversky and Bar-Hillel (1983) and Nielsen (1985) presented counter-examples involving concave utility functions with kinks at a specific value of wealth such that at that value, people would turn down some single gambles while accepting two of them. These examples worked because neither the first or second derivative of the utility function were defined at the kink point, so the Taylor series approximation we have used could not apply in these cases.

However, for the traditional assumption of DARA preferences, such that $r'(W)$ is always defined and negative, the result has strong implications. Ross (1999) noted that the conditions for multiple gambles to be eventually accepted once N became large enough were not satisfied by people with DARA preferences. In fact, for local gambles the opposite result applies. Stated formally, the result is as follows.

Proposition 2: Under the conditions for utility functions and gamble size described in Proposition 1, DARA preferences imply that the threshold level of wealth for being willing to accept N independent positive value gambles is increasing in N.

Proof. Monotonicity of $r(W)$ ensures there is a unique mapping from $r(W)$ to W. The level of wealth at which people are indifferent between accepting or rejecting N independent gambles is defined implicitly by equation 7. The indifference level of $r(W)$ falls as N increases. DARA means the indifference level of wealth rises as N increases and that everyone below this level of wealth will decline the offer of N gambles and everyone above it will accept. □

This result means that for a specific positive expected value gamble and a given distribution of wealth among a population of expected utility maximizers with DARA preferences, the fraction of people that will accept N gambles will decline with N. This does not mean that increasing the number of favorable gambles can't at some point be considered a good thing. The following result shows this property applies for DARA preferences once wealth reaches a certain level.

Proposition 3: Under the conditions for utility functions and gamble size described in Proposition 1, DARA preferences imply that for each value of N, there is a threshold level of wealth such that above this threshold, larger multiples of positive expected value gambles are preferred to multiples of size N.

Proof. Define

$$F(W,N) = U'(W)\left(N\mu - \frac{r(W)}{2}\left(N\sigma^2 + N^2\mu^2\right)\right) \tag{9}$$

Larger multiples will be preferred if $\frac{\partial F}{\partial N} > 0$. This can be calculated as

$$\frac{\partial F}{\partial N} = U'(W)\left(\mu - \frac{r(W)}{2}\left(\sigma^2 + 2N\mu^2\right)\right) \tag{10}$$

This will be positive if

$$r(W) < \frac{2\mu}{\sigma^2 + 2N\mu^2} \tag{11}$$

DARA preferences imply that there will eventually be a level of wealth such that $r(W)$ is below this value, so above that threshold people will prefer larger to smaller multiples. Also, the level of $r(W)$ defined in equation 11 is lower than the one defined in 7. This means there will be an interval of wealth in which people will be willing to accept multiple gambles of a certain size but would prefer smaller ones. □

These results support Samuelson's conjecture that an expected utility maximizer who turns down one gamble should turn down multiple such gambles. But they do so without his assumption that the once-off gamble would always be turned down across the full range of wealth values being considered. Instead, DARA preferences mean someone could turn down a gamble at one value of wealth and accept it if wealth was higher. The results also tell us more about how people would behave than Samuelson's conjecture. Some people who would accept one of these gambles would turn down two of them and some who would accept two would reject three. For these people, increasing the number of gambles makes the offer less attractive, not more. And some people who would accept four gambles would prefer to accept three. But then, for high enough wealth, people will prefer larger multiples to a smaller one.

Our proof method differs from Samuelson's and raises questions about his inductive approach to assessing offers of multiple independent gambles. Samuelson's argument implied that when thinking about whether to accept N gambles, someone may start from the point of view of having already decided to accept at least $N-1$ gambles and then deciding whether to raise this to N.

Samuelson was correct that if someone would reject a stand-alone gamble, then the outcome of this thought experiment would be to reject the Nth additional gamble, and then to reflect that they should also reject the $(N-1)$th gamble and so on. But we also have a converse result. Willingness to accept a once-off gamble does not mean accepting an Nth gamble conditional on having already accepted $N-1$ of them. In fact, we have shown there are people who would accept the once-off gamble but decide to limit the number of gambles they will accept to $N-1$ rather than accept an Nth gamble. This shows that, as a general matter, decision-making about aggregated risks cannot be modeled as a sequence of separate once-off decisions and solved via backward induction. A commitment to accept at least $N-1$ gambles affects the tolerance for risk when considering adding another gamble.

But what about larger gambles? One could flag two different concerns about the second-order approximation used here. The first is that we are evaluating attitudes to potentially large risks using relative risk aversion at initial wealth and this may not be a good way to consider whether or not one would accept $N+1$ gambles given you are willing to accept N. If you already know you would accept N positive expected value gambles, you may factor in that you are probably going to be sufficiently wealthy for your attitudes to taking an additional risk to be different than at your current wealth level. This concern is easily addressed. In an appendix, we derive the same findings as here by comparing the decisions to accept either N or $N+1$ gambles using a Taylor expansion around $W+\mu N$, the expected wealth from accepting N gambles.

Still, if the amounts involved are large enough, it may be that no second-order approximation is accurate. From Pratt (1964), we know that DARA requires the utility function to have a positive third derivative so could adding third-order terms deliver different results than the second-order approximation? Our numerical calculations below show that at least for CRRA utility, the answer is no.

3. Numerical Calculations

Here we report some numerical illustrations illustrating that the results for local gambles just presented generalize to gambles that can be large relative to initial wealth.

3.1. Calibration

We didn't needed stronger assumptions than DARA to prove the results about local gambles but to assess attitudes towards larger gambles, we use numerical calculations and thus specific utility functions. As discussed by Pratt (1964), one can construct non-CRRA examples of DARA utility functions but they tend to be relatively unintuitive, so we will restrict ourselves to the CRRA class

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma} \tag{12}$$

which have the property that $r(W) = \frac{\gamma}{W}$ and thus satisfy DARA.

In assessing the appropriate levels of γ for modeling someone who would turn down the kind of offer made by Samuelson, it should be noted that the amounts in the offered gamble were not as trivial as they might seem from today's perspective. The CPI has increased by a factor of 10 since 1960, so in today's money the offer would be an equal chance of winning \$2,000 and losing \$1,000.

However, even assuming modest levels of wealth, turning down this offer would require a relatively high level of relative risk aversion. Given this, our reported calculations will have $\gamma > 1$, which implies more risk aversion than the log utility function that various empirical studies have found to be a good approximation to real-world risk aversion.[1]

CRRA utility functions with $\gamma > 1$ are bounded from above by zero and tend to minus infinity as wealth goes to zero. These are the opposite of the features Ross (1999) showed were required for people who would not accept once-off gambles to accept multiple independent gambles once N became large enough. As such, it should not be surprising that our numerical calculations do not display the eventual acceptance property. Instead, they confirm that the local results just described also apply to larger gambles without relying on Taylor approximations.

3.2. Restriction to Positive Wealth

As noted above, our calculations do not consider offers that have the potential for people to lose all (or more) of their wealth. There are two reasons for this restriction.

First, from a practical perspective, for the CRRA utility functions that we use for our calculations, the possibility of even one negative outcome for wealth produces complex-valued expected utilities which cannot be interpreted. In any case, for the $\gamma > 1$ examples we will use for illustration, people would never consider taking a gamble that could potentially (even with a very small probability) reduce their wealth to (close to) zero. The almost-infinite potential losses in utility in close-to-worst outcomes would always outweigh any potential gain. In this case, gambles involving losing more than one's wealth are irrelevant.

Second, there are various problems with interpreting negative wealth. The literatures on both multiple gambles and St. Petersburg paradoxes have debated how to interpret the possibility of negative wealth, considering issues such as whether the possibility of bankruptcy places a lower bound on utility for very bad outcomes.[2] For CRRA utility functions with $0 < \gamma < 1$, one could construct a continuous real-valued "extension" to the utility function that generates negative utility values, for example by assuming $U(-x) = -U(x)$ where $x > 0$ but this is ultimately an ad hoc approach and does not work for $\gamma > 1$ because it implies non-monotonic utility functions.

There is also a question of realism and enforceability of gambles that result in negative wealth. While people can certainly have negative net wealth in the real world, these cases generally occur because they are expected to pay off this debt using future labor income. This suggests it is best to interpret wealth here in the broadest possible sense—as both the sum of current financial assets and the present discounted value of future labor income. Under this interpretation, our restriction

[1]See Chetty (2006), Layard, Mayraz and Nickell (2008) and Gandelman and Hernandez-Murillo (2013) for various estimates of γ using different methodologies.

[2]See Samuelson (1977) and Ross (1999) for discussions of this issue

could be interpreted as allowing the consideration of gambles that could wipe out someone's current financial wealth but not ones that had the potential to also take away the lifetime value of their labor income.

3.3. Results

Our numerical calculations describe how expected utility maximizers behave when offered multiple independent gambles with equal chances of winning G or losing L where $G > L$. We calculate exact expected utilities from accepting multiple independent gambles using the true binomial distributions of all possible outcomes. In other words, we assume people will accept an offer of N gambles if

$$U(W) < E\left[U(W + X_N)\right] = \left(\frac{1}{2}\right)^N \sum_{i=0}^{N} \binom{N}{i} U\left(W + iG - (N - i)L\right) \tag{13}$$

and report the smallest value of W for which people will accept the offer.

Table 1 reports calculations for $\gamma = 1.5$ and $\gamma = 2$ and gambles with equal chance of losing 100 and winning one of three different values of G (Samuelson's offer of 200 and also 150 and 110). We only consider gambles for which the worst possible outcome is to be left with W = 1 and model decisions at integered values of wealth.[3] The table shows the lowest integer-valued levels of wealth required to accept N independent gambles. It can be seen from each of the columns that thresholds rise with N and the magnitude of this effect is often big. For γ = 2 and G = 200, the threshold rises from 401 for a once-off gamble to 1103 for 11 independent gambles.

The limitation on gambles to those that do not potentially generate negative wealth has an impact on the calculations as N rises. For example, for $\gamma = 1.5$ and $G = 200$, the table ends with the value of 801 for $N = 8$. This means that everyone that can be offered 8 independent gambles will take the offer. Calculations show that once we have reached a value of N for which this is true, it is also true for higher values of N. For gambles with smaller expected gains, the percentage rise in the thresholds for accepting larger multiples are smaller but it takes higher values of N to get to the point where the largest allowable multiple will be accepted.

The figures in Table 1 are thresholds such that everyone above this level will accept N independent gambles. But at the thresholds listed here, the agents will generally prefer a smaller multiple than one of size N. Table 2 shows the preferred multiple sizes at each of these thresholds. For example, with $\gamma = 1.5$, G = 110 and W = 1719, the table shows that N = 12 will be accepted, but this person's favorite offer would be N = 6. Perhaps in this case, rather than accepting the largest feasible multiple, this person may attempt to negotiate a smaller one. As wealth grows, for each gamble we eventually reach a point where the preferred offer is the largest one allowable, as suggested for local gambles by Proposition 3.

[3] The results are essentially the same if we allow smaller worst possible levels of wealth.

These results show that a property that looks a little bit like the "eventual acceptance property" applies for DARA utility. Once people are wealthy enough to be able to consider large gambles, then compounding makes the offers more attractive. But this does not refute Samuelson's argument because it does not show that people poor enough to turn down one gamble would find multiple independent gambles more acceptable. Indeed, if Samuelson's colleague had DARA preferences, then not only should he have been unwilling to accept multiples of the rejected gamble at his current wealth level but we could construct an increasing set of payments to him as N rose such that, even after receiving these payments and thus having lower absolute risk aversion, he should still have turned down the offers of multiple independent gambles.

We also did these calculations for values of $\gamma < 1$. The results are consistent with our other findings but less interesting. For example, for $\gamma = 0.9$, the threshold for accepting one, two and three gambles are W = 181, W = 234 and $W = 306$ but for larger multiples, everyone eligible to accept does so. These utility functions satisfy Ross's conditions when applied to positive wealth values but they still do not imply that increasing the number of gambles makes the offers more attractive for those with low wealth. Ross's results allowed those with low wealth to consider a large number of independent gambles that could lead to large negative wealth outcomes. We restrict the set of options to offers that rule out negative wealth and, within the restricted choice available, those who would turn down a one-off gamble will also turn down the larger feasible multiple gambles available to them.

One thing to note about our tables is the number of gambles considered falls well short of the colleague's counter-offer to take 100 gambles. This is because, in our calculations, people who are wealthy enough to be feasibly offered large multiples are all willing to accept them at levels of N well below 100. As Rabin (2000) has pointed out, the N = 100 offer with G = 200 and L = 100 seems incredibly attractive. It has an expected value of $5000 and a probability of only 0.00004 of losing money. Rabin argued that nobody would turn down this gamble. In our calculations, since we are not considering negative wealth outcomes, the N = 100 option can only be offered to people with $W >$ 10,000. It takes extremely high levels of risk aversion for anyone with these wealth levels to turn down this offer. For $W = 10,001$, it takes $\gamma >$ 8.5 to turn it down, well above the levels estimated from empirical studies.

This suggests the greater culprit in the violation of expected utility theory in the story is probably not the willingness to accept the large multiple but rather the rejection of the once-off gamble. A CRRA expected utility maximizer with W = 10,001 should only turn down the G = 200, L = 100 once-off gamble if they had $\gamma >$ 48. The colleague cannot simultaneously have had both $\gamma <$ 8.5 and $\gamma >$ 48.

Table 1: Threshold wealth levels for accepting N independent gambles with equal chances of winning G and losing 100 for different values of G and γ

	$\gamma = 1.5$			$\gamma = 2$		
N	$G = 200$	$G = 150$	$G = 110$	$G = 200$	$G = 150$	$G = 110$
1	299	450	1650	401	600	2201
2	338	475	1656	435	622	2205
3	386	503	1662	475	646	2210
4	446	538	1668	523	673	2215
5	520	581	1674	580	704	2221
6	608	638	1680	648	741	2226
7	703	714	1686	728	788	2231
8	801	804	1693	815	849	2236
9		901	1699	908	924	2241
10			1706	1005	1012	2247
11			1712	1103	1106	2252
12			1719	1202	1203	2258
13			1726	1301	1302	2263
14			1734		1401	2269
15			1741			2275
16			1749			2280
17			1757			2286
18			1801			2292
19						2298
20						2304
21						2310
22						2317
23						2324
24						2401

Table 2: Preferred multiple size for threshold wealth levels for accepting N independent gambles with equal chances of winning G and losing 100 for different values of G and $\gamma = 1.5$

Will Accept	G = 200	Preferred	G = 150	Preferred	G = 110	Preferred
1	299	1	450	1	1650	1
2	338	1	475	1	1656	1
3	386	2	503	2	1662	2
4	446	3	538	2	1668	2
5	520	4	581	3	1674	3
6	608	5	638	4	1680	3
7	703	6	714	6	1686	4
8	801	7	804	7	1693	4
9	901	8	901	8	1699	5
10	1001	9	1001	9	1706	5
11	1101	10	1101	10	1712	6
12	1201	12	1201	11	1719	6
13	1301	13	1301	13	1726	7
14	1401	14	1401	14	1734	8
15	1501	15	1501	15	1741	8
16	1601	16	1601	16	1749	9
17	1701	17	1701	17	1757	9
18	1801	18	1801	18	1801	13
19	1901	19	1901	19	1901	18
20	2001	20	2001	20	2001	20

3.4. Comparison with Rabin (2000)

These calculations bear some resemblance to those in Rabin's (2000) famous paper on risk aversion. In an approach similar to Samuelson's theorem, Rabin assumes that if people turn down a specific gamble, then they will turn it down at all possible values of wealth. Based on this assumption, he illustrates, based only on concavity of the utility function, how turning down small positive expected value gambles would imply turning down larger gambles with much higher expected values that it might be imagined most people would take.

The structure of our problem is different. Rabin considers gambles with two possible outcomes while accepting N independent gambles implies $N + 1$ possible outcomes. But our calculations also show how turning down one kind of gamble can imply turning down larger gambles with greater expected value and the economics underlying why people turn down these larger gambles is similar.

For example, consider someone with $W = 299$ and $\gamma = 1.5$. Table 1 shows they will just about accept a gamble with equal chances of gaining \$200 and losing \$100 but will not accept two of these gambles. A quick calculation shows that this person would not accept an equal chance of winning \$4050 and losing \$200. So we should not be surprised that people who are essentially indifferent between winning \$200 and losing \$100 will turn down two of these gambles, which implies a one-quarter chance of winning \$400, a half chance of winning \$100 and a one-quarter chance of losing \$200. Similar calculations of how turning down equal chances of winning \$200 or losing \$100 implies an extreme aversion to losing \$200 can be obtained from Rabin's non-parametric method without assuming a particular form for the utility function. A key difference, however, is that our numerical results do not rely on any global assumption about rejecting certain gambles.

4. Discussion

Samuelson's paper has been discussed often over the years. A common reaction has been to consider his colleague's response to be perfectly reasonable. After all, people who accept multiple positive expected value gambles will generally be better off than those who reject them. Psychologist Lola Lopes (1996) described Samuelson's colleague's counter-offer as "eminently sensible" and viewed the story as an illustration of the flaws inherent in expected utility theory. In other cases, such as the previous literature on multiple gambles discussed here, it has been argued that perhaps the colleague's preferences were consistent with expected utility theory. This appears to have been the colleague's preferred explanation. Samuelson noted that his colleague assured him that he "*wants to stand with Daniel Bernoulli, Bentham, Ramsey, v. Neumann, Marschak and Savage on this basic issue*" of being an expected utility maximizer.

What we can say is that if the colleague was an expected utility maximizer, then it seems his preferences did not satisfy DARA and as such were unusual. We have also shown that you do not need to invoke Samuelson's assumption that someone would turn down a gamble at all wealth levels to obtain the result that turning down one gamble should imply turning down multiple such gambles. In addition, we have described how the inductive method he used to arrive at his conclusion is not generally the right way to evaluate whether to accept multiple independent gambles. And contrary to Samuelson's footnoted qualifier, DARA preferences do not over-turn his theorem but actually strengthen it.

Two final points are worth making about this question. First, there is a possible alternative interpretation of Samuelson's colleague's counter-offer. If the colleague had in mind that they had the option to stop taking gambles at some point then Gollier (1996) has shown that an expected utility maximizer with DARA preferences will accept gambles as part of a sequence that they would not accept on a standalone basis. The possibility of multiple wins combined with being able to stop after losses makes the offers more attractive as part of a sequence than when presented alone or as part of a bundle of independent gambles. Gollier showed that the wealth threshold for accepting gambles gets lower as there are more rounds remaining. This provides one way to possibly view the colleague as still being an expected utility maximizer with DARA preferences. This interpretation however, runs counter to the colleague's request to "let me make 100 such gambles" and the appeal to the law of large numbers, since ending a sequence early would rule out the sample getting large enough to invoke this law.

Second, as an MIT professor in the early 1960s, the colleague almost certainly earned at least $10,000 a year and likely had a comfortable lifestyle. Our calculations suggest he was either joking or was not an expected utility maximizer. Most likely, Benartzi and Thaler (1996) and others who have put down his reluctance to take the once-off gamble as being due to loss aversion are right. But I guess we will never know.

References

[1] Benartzi, Shlomo and Richard Thaler (1996). "Risk Aversion or Myopia? Choice in Repeated Gambles and Retirement Incentives," *Management Science*, Volume 45, pages 364-81.

[2] Calvet, Laurent and Paolo Sodini (2014). "Twin Picks: Disentangling the Determinants of Risk-Taking in Household Portfolios," *Journal of Finance*, Volume 69, pages 867-906.

[3] Chetty, Raj (2006). "A New Method of Estimating Risk Aversion," *American Economic Review*, Volume 96, pages 1821-1834.

[4] Chiappori, Pierre-André and Monica Paiella (2011). "Relative Risk Aversion is Constant: Evidence from Panel Data," *Journal of the European Economics Association*, Volume 9, pages 1021-1052.

[5] Dybvig, Philip and Steven Lippman (1983). "An Alternative Characterization of Decreasing Absolute Risk Aversion," *Econometrica*, Volume 51, pages 223-234.

[6] Gandelman, Nestor and Ruen Hernandez-Murillo (2013). "What Do Happiness and Health Satisfaction Data Tell Us About Relative Risk Aversion?" *Journal of Economic Psychology*, Volume 39, pages 301-312.

[7] Gollier, Christian (1996). "Repeated Optional Gambles and Risk Aversion," *Management Science*, Volume 42, pages 1524-1530.

[8] Harper, James and Kenneth Ross (2005). "Stopping Strategies and the Gambler's Ruin," *Mathematics Magazine*, Volume 78, pages 255-268.

[9] Hellwig, Martin (1995). "The Assessment of Large Compounds of Independent Gambles," *Journal of Economic Theory*, Volume 67, pages 299-326.

[10] Huber, Tobias, Johannes Jaspersen, Andreas Richter and Dennis Strumpl (2023). "On the Change of Risk Aversion in Wealth: A Field Experiment in a Closed Economic System," *Experimental Economics*, Volume 26, pages 1-26.

[11] Layard, Richard, Guy Mayraz and Steve Nickell (2008). "The Marginal Utility of Income," *Journal of Public Economics*, Volume 92, pages 1846-1857.

[12] Levy, Haim (1994). "Absolute and Relative Risk Aversion: An Experimental Study," *Journal of Risk and Uncertainty*, Volume 8, pages 289-307.

[13] Lopes, Lola (1996). "When Time Is Of the Essence: Averaging, Aspiration and the Short Run," *Organizational Behavior and Human Decision Processes*, Volume 65, pages 179-189.

[14] Meeuwis, Maarten (2022). "Wealth Fluctuations and Risk Preferences: Evidence from US Investor Portfolios," Washington University in St. Louis.

[15] Nielsen, Lars Tyge (1985). "Attractive Compounds of Unattractive Investments and Gambles," *Scandinavian Journal of Economics*, Volume 87, pages 463-473.

[16] Pratt, John (1964). "Risk Aversion in the Small and in the Large," *Econometrica*, Vol. 32, pages 122-136.

[17] Rabin, Matthew (2000). "Risk-Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica*, Volume 68, pages 1281-1292.

[18] Ross, Stephen (1999). "Adding Risks: Samuelson's Fallacy of Large Numbers Revisited," *Journal of Financial and Quantitative Analysis*, Volume 34, pages 323-339.

[19] Samuelson, Paul (1963). "Risk and Uncertainty: A Fallacy of Large Numbers," *Scientia*, Volume 98, pages 108-113.

[20] Samuelson, Paul (1977). "St. Petersburg Paradoxes: Defanged, Dissected, and Historically Described," *Journal of Economic Literature*, Volume 15, pages 24-55.

[21] Tversky, Amos and Maya Bar-Hillel (1983). "Risk: The Long and the Short," *Journal of Experimental Psychology: Learning, Memory and Cognition*, Volume 9. pages 713-717.

A Alternative Taylor Expansion

Here is an alternative way to prove the theoretical results in Section 2 that does not rely on using the level of wealth when rejecting the gamble to calibrate risk aversion. Rather than assess the expected utility from multiple independent gambles by approximating around W, we can compare the the option to take N or $N+1$ gambles by approximating around $W + N\mu$, the expected value of wealth when accepting N gambles.

Writing X_N as the outcome from accepting N independent gambles, expected utility can be written as

$$E\left[U\left(W + X_N\right)\right] = U\left(W + N\mu\right) + \frac{U''\left(W + N\mu\right)}{2} N\sigma^2 \tag{A.1}$$

Writing X_{N+1} as the outcome from accepting $N+1$ gambles, expected utility is

$$E\left[U\left(W + X_{N+1}\right)\right] = U\left(W + N\mu\right) + \mu U'\left(W + N\mu\right) + \frac{U''\left(W + N\mu\right)}{2} E\left(X_{N+1} - \mu N\right)^2 \tag{A.2}$$

This last term can be written as

$$\begin{aligned} E\left(X_{N+1} - \mu N\right)^2 &= E\left(X_{N+1}\right)^2 - 2\mu N E\left(X_{N+1}\right) + N^2\mu^2 \\ &= \left(N+1\right)\sigma^2 + \left[\left(N+1\right)^2 - 2N\left(N+1\right) + N^2\right]\mu^2 \\ &= \left(N+1\right)\sigma^2 + \left[2N^2 + 2N + 1 - 2N^2 - 2N\right]\mu^2 \\ &= \left(N+1\right)\sigma^2 + \mu^2 \end{aligned} \tag{A.3}$$

Someone will prefer accepting N gambles to $N+1$ if

$$\mu U'\left(W + N\mu\right) + \frac{U''\left(W + N\mu\right)\left[\left(N+1\right)\sigma^2 + \mu^2\right]}{2} \leq \frac{U''\left(W + N\mu\right) N\sigma^2}{2} \tag{A.4}$$

which simplifies to

$$\mu U'\left(W + N\mu\right) + \frac{U''\left(W + N\mu\right)\left(\sigma^2 + \mu^2\right)}{2} \leq 0 \tag{A.5}$$

So the smaller multiple will be preferred if

$$r\left(W + N\mu\right) \geq \frac{2\mu}{\mu^2 + \sigma^2} \tag{A.6}$$

For a fixed value of N, DARA means this condition is more likely to be satisfied for smaller values of W, meaning people with wealth below a specific threshold will prefer the smaller multiple. As wealth rises, this inequality will be reversed so eventually people prefer larger multiples. One can use these results to calculate a set of threshold levels of wealth rising in N for being willing to accept multiple gambles of various sizes.