# How big is big enough?

***By Marcel Vonk***

In poker forums, one often encounters posts of the following nature:

*I have played 50 Sit&Gos, and have an ROI of 50%. Should I move up in limits?*

The standard reply is usually brief:

*Too small a sample size! Come back when you have played N tournaments.*

In the above, N is a number that in the eyes of the poster represents a big enough sample size. Usually, the number given is in the order of 200-500. In this article, I want to take a closer look at what a "large enough sample size" really is in different cases. As we will see, a number in the above range is often correct, but there are important exceptions. If your automatic reply is like the one above, then read on – you might be surprised!

## ROI

To begin with, let me remind you what ROI, or "return on investment", means. It is the net amount of money that you have earned (or lost), expressed as a percentage of the amount of money that you have invested in buy-ins and tournament fees.

As an example, assume you have played 50 Sit&Gos with a buy-in of $10 and an entry fee of $1. This means you have invested 50 x $11 = $550. Also assume that your total prize money in these Sit&Gos has been $825. Then you have earned $825 - $550 = $275. This is exactly half of the $550 you have invested, so your return on investment, or ROI, for these tournaments is 50%. Not bad!

Had your prize money been only $495, you would have had a net loss of $550 - $495 = $55 in total. This is 10% of the $550 you paid, so your ROI would be –10%.

The return on investment is a nice number to show how well you have played, but how can we know that this number is actually a result of your playing ability, and not of some lucky run of the cards? Intuitively, it is clear to most people that if someone obtains an ROI of 50% over 50 tournaments, this is much more likely to be a result of good fortune than if that same person maintains an ROI of 50% during 1000 tournaments. Using statistics, one can make this intuition precise by calculating quantities called the variance and the standard deviation.

## A road map

Before I continue, let me give you a little road map for this article, depending on what you want to learn from it. Most readers will belong to one of the following three categories:

1. Those who already know the concepts of variance and standard deviation. If you belong to this group, feel free to skip the next section, or just read on and let it serve as a useful reminder.
2. Those who do not know these concepts, but would like to be able to calculate them for themselves. In this case, I have to warn you: I will introduce variance and standard deviations as "black boxes": I will tell you how to calculate them and how to use them, but not *why* all of this is true. If you can live with that, carry on. If you cannot, I recommend you find an introductory text about statistics (feel free to send me an e-mail for suggestions) to deepen your knowledge.
3. Those of you who are mainly interested in the answer to the question asked in the title, not in how it is obtained. If this is the case, you can safely skip to the section "The Results" below.

Having said that, here we go!

## Variance and standard deviation

The variance of a series of results measures how much these results differ from the average. To calculate it, we first calculate the average, and then list how much each of the results differ from this average. Finally, we square all these differences (that is: multiply them with themselves) and add them up.

As an example, assume a simple setup where you play six-handed Sit&Gos for $10 + $1, on a winner-takes-all base. Assume you have played four of these. Three times, you have not made the money, but the fourth time you came in first place, winning the prize of $60. On average, you have made $15 per Sit&Go. The three actual results, $0, $0, $0 and $60, differ from this average by $15, $15, $15 and $45 respectively. To find the variance, we square these three numbers and add them:

Variance = 15 x 15 + 15 x 15 + 15 x 15 + 45 x 45 = 225 + 225 + 225 + 2025 = 2700

The observant reader might note that I do not write "$2700". If you remember your high school physics classes, you might recall that when we take a length (in meters) and multiply it by another length (in meters), what we obtain is an area, expressed in *square* meters. In principle, the same holds here: I multiplied amounts (in dollars) by amounts (in dollars), so if you want to sound really impressive, you could say that the result is "2700 square dollars".

To get back to a meaningful statement in terms of an amount of money, we therefore have to take the square root of the number we have just obtained. This quantity is what is called the standard deviation. Grabbing a calculator, we find that it is

Standard deviation = sqrt(2700) = $52,

where I rounded the result to whole dollars. Before explaining what this number means, let me scale the results a bit. Suppose that now, we play 400 Sit&Gos, and we finish out of the money 300 times, and win first place 100 times. If you repeat the above calculation, you will find that the variance is now 270,000 and the standard deviation is $520. There is a general law behind this: if we scale the number of results by a factor of 100 (or more generally, N), the variance also scales by a factor of 100 (or again, N), but the standard deviation only scales with a factor of 10 (or in general, the square root of N).

So, what does this standard deviation mean? It means the following: assume that this sequence of events would be representative for how well you play. That is: whenever you play, you have a chance of one quarter that you win the tournament, and a chance of three quarters that you do not end in the money. On average, you would therefore expect to win 100 x $60 = $6000 if you play 400 times. Subtracting the 400 x $11 you have paid in entrance fees, you will have made an average of $6000 - $4400 = $1600. This corresponds to an ROI of $1600 / $4400 x 100% = 36.4%. Impressive!

However, of course you will not always exactly win the $6000 when you play 400 times. Sometimes you will win a little more; sometimes you will win a little less. This is where the standard deviation of $520 comes in: it tells you that usually, you will win between $6000 – $520 = $5480 and $6000 + $520 = $6520. Again subtracting the $4400 you have paid, this means you will usually earn between $1080 and $2120. Your ROI will then be between

$1080 / $4400 x 100% = 24.5%

and

$2120 / $4400 x 100% = 48.2%

The “usually” in the above sentences can be made mathematically precise: one can show that in 68% of the cases, your ROI over 400 tournaments will lie between the boundaries calculated here. That’s not very accurate, huh? Well, we can do better: if we subtract and add *twice* the standard deviation, we get boundaries between which our ROI over 400 Sit&Gos will be in 95% of the cases. Try the calculation for yourself; you will find that the ROI will be between 12.7% and 60%.

Now, we can clearly see what people mean when they say that a sample size is not big enough. For a very decent winning player, who should obtain an ROI of 36.4% in the long run, even a sample of 400 Sit&Gos might still give a result as low as 12.7% or as high as 60%!

Well, at least in the above example, all of these ROIs are still positive numbers: the player might sometimes win a lot over 400 tournaments, and sometimes only a little, but he is very likely to win at least something! This brings me to what I will define to be a big enough sample size:

***The sample size for a winning player is big enough if the ROI obtained is at least two standard deviations away from zero.***

In other words, for a winning player, if the ROI obtained is representative for the player's ability, he would be 97.5% sure to win again over another run of the same length. (Why not 95%? Well, of the 5% where he is *not* in the interval calculated above, he will of course actually *win even more* half of the time!)

## A sample calculation

Let us apply this knowledge to a more realistic situation. Assume we play nine-handed \$10 + \$1 Sit&Gos, where the first three prizes are \$45, \$27 and \$18. We want to consider a player whose ROI is 20% over 90 tournaments. We will make one more assumption: the distribution of the player's results is *linear*. That is: he will end in $5^{th}$ place (the average) 10 times, in $6^{th}$ place 10-n times, in $4^{th}$ place 10+n times, in $3^{rd}$ place 10+2n times, and so on. We have to figure out what n is, so that his ROI will be the 20% we are looking for. Adding up his winnings, we find that they are

45 x (10+4n) + 27 x (10+3n) + 18 x (10+2n) = 900 + 297n

dollars. Subtracting the \$990 in entry fees, he has therefore earned 297n – 90 dollars. Recall that we wanted his ROI to be 20% of \$990 = \$198, so to find n we must solve

297n – 90 = 198

meaning

297n = 288

or

n = 288/297 = 0.97.

In other words, to have an ROI of 20%, a player will have to end in $4^{th}$ place 10.97 times, in $3^{rd}$ place 11.94 times, and so on. (Of course, our player cannot exactly end in fourth place 10.97 times, but remember that we are talking about *averages* here.) With these numbers, we can now calculate the variance and the standard deviation using the recipe above. We find:

Variance = 25702
Standard deviation = \$160.32

In other words, the ROI of \$198 is a little more than one standard deviation away from zero, and the sample of 90 Sit&Gos is not big enough. What sample size would be *exactly* big enough? To answer this, we recall that if the number of tournaments scales with a factor of N, the standard deviation scales with a factor of sqrt(N). Or vice versa, if we want to scale the standard deviation by a factor of M, we have to scale the number of tournaments by M x M. This will of course also scale the expected earnings by a factor of M x M. So for these to be exactly two standard deviations above zero, we have to solve

M x M x $198 = 2 x M x $160.32

or, after removing a factor of M on each side,

M x $198 = 2 x $160.32

meaning

M = 2 x $160.32 / $198 = 1.62

Thus, we find that an exactly big enough sample size would be 1.62 x 1.62 x 90 = 236 tournaments. If a player achieves the ROI of 20% over this number of tournaments, he can be quite sure that he is a winning player.

## The results

To get a feeling for big enough sample sizes, we can do the above calculations for ROIs of –20%, -10%, -5%, 5%, 10%, 20%, and 40%. The results are given in the following table:

| ROI | 1SD | 2SD |
|---|---|---|
| -20% | 46 | 182 |
| -10% | 199 | 794 |
| -5% | 824 | 3297 |
| 5% | 878 | 3513 |
| 10% | 226 | 902 |
| 20% | 59 | 236 |
| 40% | 16 | 63 |

The first column gives the ROI of the player we are considering. The last column gives the number of tournaments this player should have played to have a big enough sample size, in the sense we explained above. That is, to be reasonably sure that the ROI he sees is a result of skill (or lack of it), not of luck. The middle column gives the sample size we would obtain if we would loosen our definition, by requiring the observed ROI to be only a single standard deviation away from zero. This reduces the sample sizes we need by a factor of four, but remember that an accuracy of a single standard deviation corresponds to an accuracy of roughly 68% - if we draw conclusions on these sample sizes, they are still quite likely to be wrong!

The numbers that are in the 5% rows might surprise the reader most. If you are only a marginally winning player, it will actually take you thousands of tournaments to be sure that you are a winning player at all! This can be intuitively explained: if you are a marginal winner, the major contribution to your results will be from luck. It takes a very long time before good luck and bad luck even out enough for your skill to show. On the other hand, if you are beating the game by a large margin, luck will only play a minor role, and you can be sure that you are a winning player after a relatively small number of tournaments.

The main conclusion of this article is therefore the following: what a “big enough sample size” is depends heavily on the ROI itself. It can vary from a couple dozen to thousands of tournaments. By studying the examples above and adapting them to your own situation, you will be able to calculate whether or not your ROI is a reliable measure of your skill.

---

This article was originally published in the Two Plus Two Online Magazine in June 2007. Here are some remarks resulting from questions I received after publication:

- The definition of standard deviation I use in this article may look strange to the reader who is familiar with the concept: usually, variance and standard deviation are introduced with extra factors of 1/N. The reason is that the standard deviation I calculate is the one for the entire series of results; not the more usual one for each individual result
- Even more precisely: in statistics, it is often more useful to define the variance containing a factor of N-1 instead of N. However, our approximation is crude enough and N is large enough for this difference not to matter
- The definition of the 95%-interval as a measure for a good sample size is a very useful one. However, one should be careful in interpreting this statement as "if I am outside this interval, there is a 97.5% chance I am a winning player". The true meaning is the other way around, as stated in the text: 97.5% of the breakeven players would *not* achieve such an ROI. To make statements of the former kind, one would need to use Bayesian statistics. The reader who is interested in this subject is encouraged to read the wonderful book "The Mathematics of Poker" by Bill Chen and Jerrod Ankenman.