BLACK SWANS ARE NOT THAT BLACK

Valentyn Khokhlov, MBA
Corporate Finance and Investments Consultant
Tel: +380-50-4779316
E-mail: val.khokhlov@gmail.com

Abstract:

Some researchers have recently criticized using the normal distribution for modeling stock returns. While it's true that the normal distribution is inappropriate and leads to the extreme outliers, known as the Black Swans problem, other elliptical distributions allow addressing this issue. The Student's t-distribution with 3 to 4 degrees of freedom and the Laplace distribution both can be used to largely eliminate Black Swans in daily returns. Both distributions are compatible with the modern portfolio theory. We also show that no single distribution is clearly preferred when describing periodic returns, but the Black Swans problem is not so acute when considering returns over holding periods longer than one month.

Keywords:

Security returns distribution, Black Swans, fat tails problem, lognormal distribution, Student's t-distribution, Laplace distribution.

The modern portfolio theory has recently been severely criticized for its inability to account for extreme losses. One of the most prominent critics, Nassim Taleb, argues in Taleb (2007) that the underlying assumption of the securities returns being normally distributed is wrong, and that the extreme outliers ("Black Swans") can't be effective modeled by any suitable random distribution to make them explainable and predictable. While Mr. Taleb gives very strong arguments, his books lack more solid theoretical basis in prove his hypothesis, such as parametric or non-parametric distributional tests on historical data samples. So, speaking of Black Swans, we first need to check the hypothesis that the normal distribution actually fails to account for them.

Second, the modern portfolio theory, despite common misperception, doesn't actually rely on returns being normally distributed. The theory only assumes that all portfolio decisions can be made using the expectation and volatility of returns. Mathematically speaking, it supposes that the portfolio optimization can effectively be performed in the mean-variable space. Thus, any distribution that is fully defined by these two parameters is suitable for the modern portfolio theory. As it was stated in Chamberlain (1983) and Owen and Rabinovitch (1983), at least one broad class of distributions, elliptical distributions, fits this requirement. And the normal (Gaussian) distribution is just one particular instance of this class. Even if it fails to describe Black Swans, we still need to check other distributions to see it there's a better substitute.

Black Swans in daily returns

Let's start looking for Black Swans in daily returns, the most exhaustive source of data for distributional analysis. The longest sample we consider in this research is the Dow Jones Industrial Average (DJIA) index sample that contains about 95 years of daily quotes. While DJIA has long history, in practice the Standard & Poor's 500 index is mostly used as a proxy for large-cap market portfolio. We'll consider a sample of daily quotes for an ETF (NYSE: SPY) that closely tracking this index, the sample contains 17 years of daily quotes. Since both DJIA and SPY are indices, i.e. represent diversified portfolios, they may not be representative for judging about single securities.

To make this research more comprehensive, we will additionally consider samples for two different liquid stocks, a high-beta stock of General Electric (NYSE: GE) and a low-beta stock of Exxon Mobil (NYSE: XOM). Both samples contain 20 years of market data.

The data samples analyzed in this research are:

- DJIA sample of 23,981 observations from January 5, 1915 to July 27, 2010;
- SPY sample of 4,510 observations from February 1, 1993 to December 23, 2010;
- GE sample of 5,042 observations from January 2, 1990 to December 31, 2010;
- XOM sample of 5,042 observations from January 2, 1990 to December 31, 2010.

All samples contain logarithmic returns and are built using Yahoo Finance market data (closing prices adjusted for corporate actions).

The DJIA sample has a large excess kurtosis (22.04) and is somewhat asymmetrical (skewness -0.56), which constitute a significant deviation from normality. Will it lead to Black Swans? Let's consider an extreme outcome of catastrophic daily loss of -6.9%, which is about 6 standard deviations below the mean daily return. If DJIA returns were normally distributed, such an outcome would occur once in 3.5 million years, so we are highly unlikely to meet even a single such a loss in our 95 years sample. In fact, there're 28 days with returns below -6.9%. It seems we've encountered a real Black Swan. The actual probability of such an outcome is 1 million times higher than the normal distribution predicts. If we took less extreme losses, like -5.8% (about 5 standard deviations below the mean) and -4.6% (about 4 standard deviations below the mean), their actual probability would be 7,000 times and 130 times higher than the normal distribution predicts.

Maybe the DJIA is a rare exception in the stock market? Let's check if SPY has any Black Swans. The SPY sample has lower but still high excess kurtosis (10.11), but it almost symmetric (skewness -0.05). For SPY, the actual probability of extreme loss of -7.4% (about 6 standard deviations below the mean) is 900 thousand times higher than the normal distribution predicts. This number is very close to that of DJIA, and this resemblance holds for other outliers considered. For extreme losses of -6.2% (about 5 standard deviations below the mean) and -5.0% (about 4 standard deviations below the mean), their actual probability would be 6,200 times and 132 times higher than the normal distribution predicts.

The Black Swans we met in S&P 500 are very similar to those in DJIA. Is the same true for single stocks? Despite GE and XOM are very different securities with quite different volatility, they both have a high excess kurtosis (8.07 for GE, 9.01 for XOM) and insignificant skewness (0.02 for GE, 0.06 for XOM). The normal distribution does a terrible job predicting extreme outliers for both securities. If we had assumed their returns are normally distributed, we would have underestimated true probability of catastrophic losses of about 6 standard deviations below the mean by 400,000 times for GE and by 700,000 times for XOM. While the numbers are slightly less than those for DJIA and SPY, they're still huge.

The Black Swans are there, and Mr. Taleb seems to be absolutely right regarding the normal distribution. But is he right regarding failure of the modern portfolio theory? Luckily, no. Some other elliptical distributions can be used to model daily returns that eliminate the Black Swans problem. As you can see from Table 1, if the Student's t distribution with 3 degrees of freedom does much better job in estimating probability of extreme outliers. For DJIA, SPY and GE it slightly (1.1 to 1.5 times) underestimates the actual probability, and for XOM it even overestimates it. The Laplace distribution, while being substantially worse the t distribution, still is much better is predicting catastrophic events, its worst mismatch is only about 20 times instead of 1 million times for the normal distribution.

Table 1. The number of years it take to encounter an extreme daily loss

Extreme Daily Loss	Number of Years Assuming Distribution of Daily Returns Normal	Student's t (df=3)	Laplace	Actual Number of Years
		Panel a. DJIA		
-6.9%	3,466,480	4.14	73.75	3.41
-5.8%	15,878	2.50	17.28	2.17
-4.6%	120	1.28	3.55	0.93
		Panel b. SPY		
-7.4%	3,249,506	4.12	57.72	3.59
-6.2%	13,994	2.46	13.74	2.25
-5.0%	148	1.33	3.27	1.12
		Panel c. GE		
-11.3%	4,141,433	4.20	55.20	10.04
-9.4%	13,745	2.46	12.54	2.01
-7.5%	121	1.28	2.85	0.84
		Panel d. XOM		
-9.3%	4,646,854	4.24	37.36	6.70
-7.7%	13,312	2.45	8.79	3.35
-6.2%	140	1.32	2.26	1.83

Goodness of fit tests

While checking the difference between actual and theoretic probability at several arbitrary selected points allowed us to make a conclusion that the normal distribution can't effectively model the log-returns of securities, that test was somewhat artificial. Now we can use a more theoretically sound way of testing distributional assumptions. The typical question asked by a statistician is whether the hypothesis that the observed data fit the theoretical distribution holds with a given significance level. We'll use a slightly different question — which of the elliptical distributions would fit the actual data best.

The first test we use in this research is the well-known Pearson's chi-squared test. Test formula for the test statistics is

$$\chi^2 = \sum_{i=1}^{n+1} \frac{\left[O_i - N\left(F(x_i) - F(x_{i-1})\right)\right]^2}{N\left(F(x_i) - F(x_{i-1})\right)}$$

where O_i denotes the actual number of observations between x_i and x_{i-1}, x_i are boundaries of bins (x_0 is $-\infty$ and x_{n+1} is $+\infty$), F is the CDF of the theoretical distribution, N is the same size. The test was performed using $n = 201$, $x_1 = -10\%$, $x_2 = -9.9\%$, …, $x_n = 10\%$, so the χ^2 is expected to adhere to the chi-squared distribution with 200 degrees of freedom (we have 203 intervals and estimate 2 parameters from the sample).

Results of the chi-squared test are given in Table 2. It's quite clear than the normal distribution is inferior to both the t-distribution and the Laplace distribution. The t-distribution is preferred for DJIA and SPY, somewhat better for GE and approximately as good as the Laplace distribution when modeling XOM log-returns. One can perceive the difference better when looking at the graphic of PDF for the actual and theoretic distribution, which is shown on Figure 1. The normal PDF clearly misfits the actual PDF, and the difference is really huge in the tails.

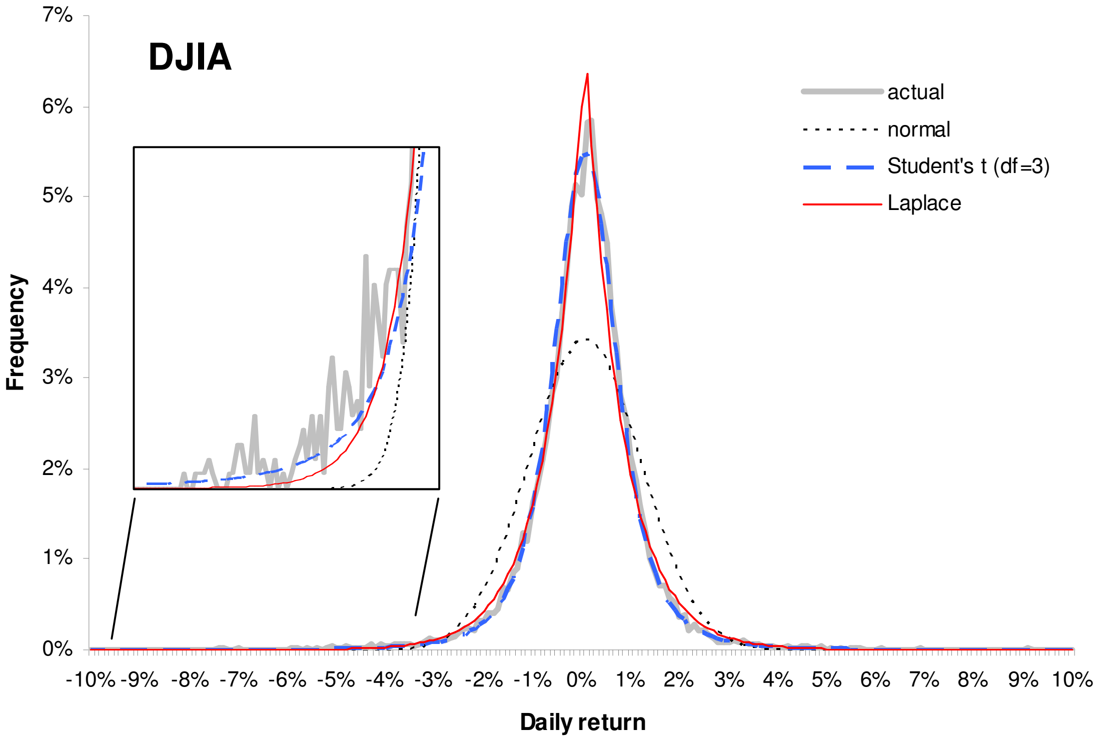

Figure 1. The probability density functions for the DJIA log-returns distributions

Note: the box, which contains a magnified version of the left tail, clearly shows that the normal distribution largely underestimates probability of the extreme outliers.

Table 2. The chi-squared test results and critical values

Security	χ^2 Values for Distributions			χ^2 Critical Values		
	Normal	Student's t	Laplace	5%	1%	0.1%
DJIA	1,427,363,104,075	368.98	3,106.25	233.99	249.45	267.54
SPY	7,697,207,461	260.49	628.59			
GE	91,023	439.25	518.93			
XOM	4,056,522	424.59	421.57			

The seconds test used is the Korlogorov-Smirnov (K-S) goodness of fit test, a non-parametric test that makes no assumption of the underlying distribution. While more powerful tests exist for normality, such as Shapiro-Wilk and Adnerson-Darling tests, we're using the K-S test since we need to check hypothesis for distributions other than the normal. The K-S test statistics is

$$D = \max\left(D^+, D^-\right), D^+ = \max_{1\le i\le N}\left(\frac{i}{N} - F\left(y_i\right)\right), D^- = \max_{1\le i\le N}\left(F\left(y_i\right) - \frac{i-1}{N}\right)$$

where y_i denotes the i-th element of the sorted sample for daily log-returns. $K = D\sqrt{N}$ is expected to follow the Kolmogorov distribution.

Results for the K-S test are given in Table 3. Please note that we use K values, for D values in this table. The normal distribution is inferior according to this test as well. Unlike the chi-squared test, the t-distribution doesn't seem preferred to the Laplace distribution any more — for latter fits somewhat better for SPY and much better for XOM, while the former is significantly better only for DJIA.

Table 3. The Kolmogorov-Smirnov test results and critical values

Security	Sample Size	*K* Values for Distributions			*K* Critical Values		
		Normal	Student's t	Laplace	5%	1%	0.1%
DJIA	23,981	13.3216	2.9779	3.4750	1.22	1.36	1.63
SPY	4,510	5.6397	1.8815	1.8229			
GE	5,042	5.5924	2.0564	2.1811			
XOM	5,042	3.7918	3.2614	1.9703			

The actual and theoretic cumulative probability functions for the SPY log-returns are shown on Figure 2. While both the t-distribution and the Laplace distribution CDFs are quite close to the actual function, the normal CDF clearly misfits it up to -2.5% and lies well below. In terms of probability, this means that the probability of losses calculated assuming the log-returns are normally distributed (i.e. the returns are log-normally distributed) would underestimate the actual losses a great deal, which is corresponds to the results displayed in Table 1.

Using the Student's t distribution for modeling daily returns

The Student's t-distribution with 3 degrees of freedom proved to be the preferred distribution for modeling log-returns according to the chi-squared test, and performed quite well according to the K-S test. But why have we chosen 3 degrees of freedom and how do other t-distributions behave? The t-distribution, unlike the normal and the Laplace ones, doesn't have a constant kurtosis — its excess kurtosis is equal to $6/(\nu-4)$, where $\nu > 4$ denotes the number of degrees of freedom. Excess kurtosis for the normal distribution is 0, and for the Laplace distribution is 3.

The sample excess kurtosis for all the securities reviewed is above 8. Therefore, the t-distributions with $\nu > 4$ fail to model the sample kurtosis. So our primary candidates for testing are the t-distributions with $\nu = 3$ and $\nu = 4$ (since the variance is infinite for $\nu \leq 2$, but all the samples clearly have finite variances). Both distributions have interesting statistical properties, and the distribution with $\nu = 4$ is practically appealing for modeling since its inverse CDF has a simple analytical expression. The Student's t-distribution is not limited to integer number of degrees of freedom. Non-integer values of ν were also tested in this research. We used different values around $\nu = 3$ and $\nu = 4$ with step of 0.10 in an attempt to find local minimums for K.

The classic t-distribution has only one parameter, the number of degrees of freedom. So we use a generalized t-distribution that allows choosing the location and the scale parameters. A random variable that adheres to the generalized t-distribution is expressed as follows:

$$\tilde{t}_\nu = \frac{t_\nu - \mu}{\sigma_t}, \quad \sigma_t^2 = \frac{\nu-2}{\nu}\sigma^2$$

where $\tilde{t}_\nu$ is a random variable following the generalized t-distribution, t_ν is a random variable following the t-distribution, μ is the expected value, σ^2 is the variance.

Table 4. The chi-squared and K-S test results for the Student's t-distribution

Security	χ^2 Values for t-distributions				*K* Values for t-distributions			
	$\nu = 2.8$	$\nu = 3$	$\nu = 3.7$	$\nu = 4$	$\nu = 2.8$	$\nu = 3$	$\nu = 3.7$	$\nu = 4$
DJIA	398.84	**368.98**	684.72	928.20	**2.9055**	2.9779	5.9419	6.3446
SPY	**201.82**	260.49	231.18	314.93	**1.8412**	1.8815	2.9037	3.0419
GE	**377.32**	439.25	392.90	477.31	2.0909	2.0564	**1.9249**	2.0047
XOM	462.90	424.59	**238.97**	258.65	3.2684	3.2614	**1.2410**	1.2660

Note: the distribution that is the best fit according to the criterion is highlighted with bold.

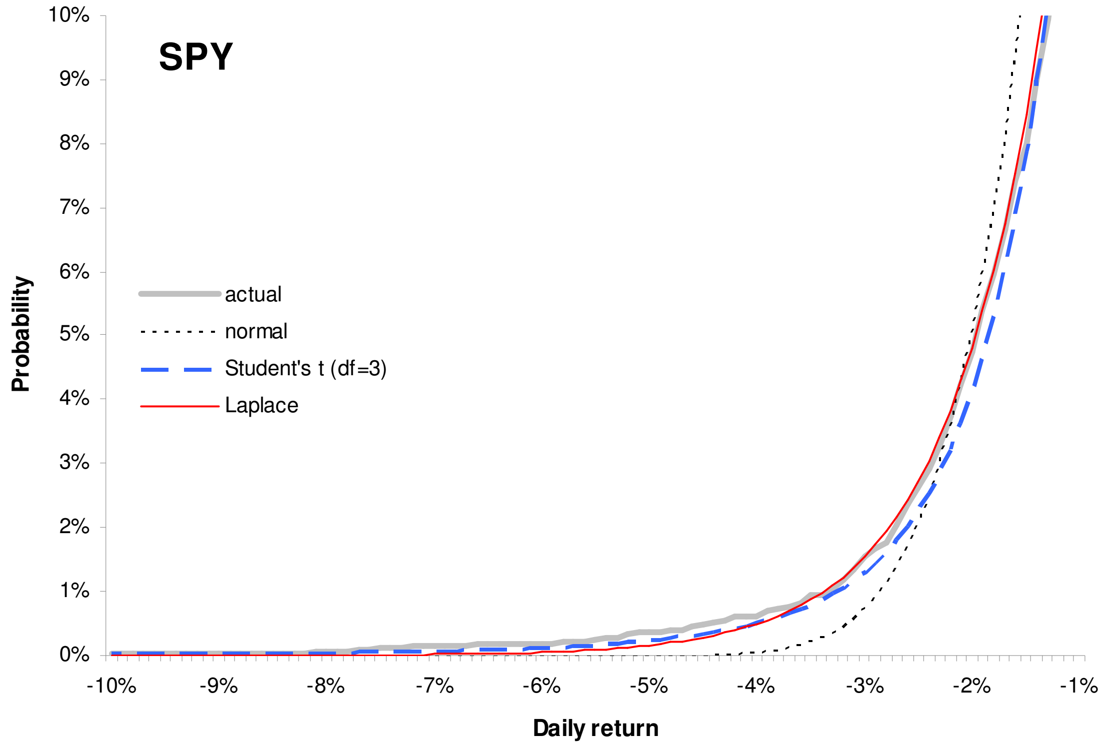

Figure 2. The left tail of the cumulative distribution functions for the SPY log-returns distributions

Test results for the Student's t-distributions with different degrees of freedom are given in Table 4. While they generally support the preliminary conclusion that the distributions with 3 to 4 degrees are freedom fit the actual data best, they do not indicate the sole best number of degrees of freedom to use. Moreover, the chi-squared and the K–S tests favor different numbers of degrees of freedom for some securities. The chi-squared test shows clear preference for $\nu = 2.8$ for modeling SPY and GE returns, $\nu = 3$ for DJIA and $\nu = 3.7$ for XOM. The K–S test shows that the t-distribution with $\nu = 2.8$ fits the actual distribution of DJIA and SPY the best, while for GE and XOM the preferred amount of degrees of freedom would be 3.7. If we limit ourselves with integer values of ν only, the most reasonable choice is $\nu = 3$, although for XOM $\nu = 4$ is much better. That may indicate that $\nu = 3$ is preferred for modeling high-kurtosis or higher-volatility assets, while $\nu = 4$ seems better for stocks with lower volatility.

Using the Laplace distribution for modeling daily returns

The Laplace distribution is an elliptic distribution that is overlooked by researchers and analysts. Unlike the normal the t-distributions, which both have a smooth top, the Laplace distribution has peaked top. Its excess kurtosis is 6, so it has fatter tails than the normal distribution. Unlike both the normal and t-distributions, the Laplace distribution has a simple analytical expression of PDF, CDF and the inverse CDF, which is quite essential for implementation of custom computer software.

The Laplace distribution, while not having any degrees of freedom, still allows discrepancy of choosing its parameters. The Laplace distribution PDF and CDF are expresses as

$$f(x) = \frac{1}{2b}\exp\left(-\frac{|x-\mu|}{b}\right); F(x) = \frac{1}{2}\left[1+\operatorname{sgn}(x-\mu)\left(1-\exp\left(-\frac{|x-\mu|}{b}\right)\right)\right]$$

where μ is the location parameter (the mode, median and mean are all equal to μ), b is the scale parameter (the variance is equal to $2b^2$). Hence we can estimate the parameters as follows:

1) L1: use the sample mean for μ and the mean absolute deviation from μ for b;
2) L2: use the sample mean for μ and the standard deviation for b: $\hat{b} = \hat{\sigma} / \sqrt{2}$;
3) L3: use the sample median for μ and the mean absolute deviation from μ for b;
4) L4: use the sample median for μ and the standard deviation for b: $\hat{b} = \hat{\sigma} / \sqrt{2}$.

Table 5. The chi-squared and K-S test results for the Laplace distribution

Security	χ^2 Values for t-distributions				*K* Values for t-distributions			
	L1	L2	L3	L4	L1	L2	L3	L4
DJIA	3,106.25	1,859.74	3,050.63	**1,810.45**	3.4750	4.4437	**2.9970**	5.0539
SPY	628.59	456.49	638.30	**454.74**	1.8229	1.7932	**1.2872**	1.9338
GE	518.93	**486.87**	524.15	491.21	2.1811	2.1413	**1.6900**	2.0458
XOM	**421.57**	431.24	433.49	443.12	**1.9703**	2.0497	2.6858	2.7475

Note: the distribution that is the best fit according to the criterion is highlighted with bold.

Test chi-squared test results for different sets of estimates are shown in Table 5. Surprisingly, the chi-squared and the K–S tests suggest using completely different sets of estimates for the Laplace distribution. The latter clearly prefers the sample median as the estimate for μ and the mean absolute deviation from it as the estimate for b (expect for XOM, for which the sample mean it the best estimate for μ), while the former typically favors using the sample standard deviation as the estimate for b.

We can also compare the test results for the Laplace and for the Student's t-distributions. While the chi-squared tests certainly indicate the superiority of the t-distribution, the K–S test favors the Laplace distribution more. Since the chi-squared test is based on the PDF and the K–S test is based on the CDF, we can conclude that it's the latter which should be given a higher weight for decision making. Therefore, using the Laplace distribution (variant L3 or L1) is the recommended way of modeling daily returns for securities.

Modeling periodic returns

As it has been shown in the previous sections, the Student's t or the Laplace distribution is much more appropriate for modeling daily log-returns both for individual stocks and for diversified equity portfolios than the normal distribution. But while for the latter it's true that the sum of normally distributed random variables adhere to the normal distribution, the same doesn't hold for both the t-distribution and the Laplace distribution. Thus, we can't rely on those distributions being appropriate for modeling periodic returns and must research this issue separately.

We use the Monte-Carlo simulation to generate a large sample (23,981) of periodic returns for different time intervals based on the assumption that the underlying daily returns follow either one of the Student's t or the Laplace distributions discussed in the previous sections. If the underlying daily rates follow the Student's t-distribution with $df = 3$, the corresponding periodic returns seem to adhere to a Student's t-distribution but with a different degrees of freedom. The longer is the period the larger is df, and starting with 63 trading days (3 months) we can reasonably model the periodic returns using the normal distribution, based on both chi-squared and Kolmogorov-Smirnov goodness of fit criteria (see Table 6).

Table 6. The chi-squared and K-S test results for a simulated periodic returns distribution when daily returns follow the Student's t-distribution with 3 degrees of freedom										
	Normal		t (df = 3)		t (df=4)		t (df=6)		Laplace	
Days	χ^2	K	χ^2	K	χ^2	K	χ^2	K	χ^2	K
1	2.1E+11	15.0288	**213.1**	1.9391	1,128.3	6.9815	2,569.8	10.5469	2,260.7	2.3186
2	1,526,509	10.0676	488.3	4.0981	280.9	2.0067	938.6	5.5339	828.9	3.2874
5	3162.6	12.3580	416.7	2.8462	566.1	4.2948	1,342.4	7.7448	820.2	3.8679
10	770.2	5.1604	2,094.0	9.0188	507.3	3.8093	**224.2**	1.2874	1,047.4	4.9328
21	456.0	3.9687	2,468.3	10.5040	691.9	5.1750	**220.1**	1.5964	1,058.5	5.4341
42	253.5	2.5762	2,223.9	11.7649	775.5	6.4506	304.9	2.8707	1,198.7	6.1310
63	**211.9**	2.1534	1,637.8	11.7318	598.7	6.4001	252.1	2.8814	933.7	5.9301
126	**206.5**	1.8200	1,372.1	12.1172	570.0	6.7929	272.6	3.4632	828.6	6.7649
252	**219.3**	1.8264	1,115.2	12.7168	488.3	7.3944	260.3	3.8230	721.4	6.8524

Note: the test statistics that are significant with 95% confidence level are highlighted with bold.

As you can see from Table 6, the Student's t-distribution with 3 degrees of freedom, which is used to model daily returns, is inappropriate for modeling periodic returns for periods over 1 week. For periods from 1 week to 1 month, it's best to use t-distributions with 6-7 degrees of freedom. For periods of 3 months and more it's quite appropriate to use the normal distribution, which passes chi-squared test and is quite good on Kolmogorov-Smirnov test, but using a t-distribution with 17-18 degrees of freedom (not shown in the table, but their χ^2 are 186 to 200 and their K are 0.60 to 0.80) provides even better fit. For intermediate periods, it's better to use a t-distribution with 11 degrees of freedom, which fits 2-3 months returns the best (not shown in the table, but its χ^2 is 173 to 200 and its K is 0.62 to 0.83).

When the Laplace distribution is used to model the underlying daily returns, the corresponding periodic returns have quite a different distribution. It's important to note than even for 2-days returns the Laplace distribution is inappropriate to use, and starting with monthly returns the normal distribution becomes appropriate to model the periodic returns both on chi-squared and on Kolmogorov-Smirnov criteria (see Table 7). For periods of 2-10 days a t-distribution seems to be reasonable, and the one with 6 degrees of freedom would not be a bad choice.

Table 7. The chi-squared and K-S test results for a simulated periodic returns distribution when daily returns follow the Laplace distribution										
	Normal		t (df = 3)		t (df=4)		t (df=6)		Laplace	
Days	χ^2	K	χ^2	K	χ^2	K	χ^2	K	χ^2	K
1	24,864,090	10.2436	1,032.2	4.8236	635.2	3.6801	1,161.0	6.3969	**116.8**	**0.5945**
2	172,293.1	5.8108	1,979.6	8.4924	446.6	3.5755	**196.6**	1.8128	496.6	3.4051
5	15,076.9	2.7287	3,576.5	11.4132	1,146.0	6.0663	336.7	2.5978	1,281.8	5.4930
10	315.5	1.2987	4,498.5	12.6291	1,634.7	7.4333	583.2	3.8856	1,576.8	6.3716
21	**214.8**	**1.1344**	4,698.7	13.0403	1,729.4	7.9301	634.7	4.3906	1,532.8	6.8451
42	**198.1**	**0.9066**	4,609.6	13.5848	1,795.5	8.2427	712.3	4.7357	1,632.5	7.0552
63	**239.5**	**0.8192**	3,696.0	13.4400	1,488.8	8.2363	638.8	4.7278	1,433.0	6.5204
126	**198.7**	**0.6970**	3,062.5	13.7043	1,251.7	8.3972	537.0	4.8118	1,134.2	6.8208
252		**0.5309**		13.8429		8.5740		4.9955		6.7149

Note: the test statistics that are significant with 95% confidence level are highlighted with bold.

Finally, let's check the hypothesis of the best distribution for modeling periodic returns on actual data. Since we can't rely on the fact that any 1-month holding period coincides with a calendar month or a 1-yeah holding period is always a calendar year, let's consider all possible daily return samples of specific length (21 trading days per month, 251 trading days per year). Using the same daily return samples for DJIA, SPY, GE and XOM we used before in this article, we can check if the goodness of fit criteria favors the same distributions as for the modeled data.

As you can see from Table 8, the actual historical distribution differs a great deal from the distribution of Monte-Carlo modeled data. Unlike the latter, it seems that the Student's t-distribution with 6 degrees of freedom doesn't fit well to the actual monthly returns distribution, and the t-distribution with 18 degrees of freedom doesn't fit at all the actual annual returns distribution. Chi-squared criterion doesn't favor any particular distribution, but according to the Kolmogorov-Smirnov criterion the Laplace distribution does a good job almost for all securities (with the exception of XOM, for which the normal distribution is even better). This result is much unexpected and clearly contradicts to result obtained using the Monte-Carlo simulation.

Table 8. The chi-squared and K-S test results for historical periodic returns distribution

Security	Normal χ^2	Normal *K*	t (df=4) χ^2	t (df=4) *K*	t (df=6) χ^2	t (df=6) *K*	t (df=18) χ^2	t (df=18) *K*	Laplace χ^2	Laplace *K*
	Panel a. Monthly returns (21 trading day)									
DJIA	2,228.6	11.1320	**684.5**	7.1325	1,011.2	7.6127			1,313.9	**5.5823**
SPY	546.5	5.2560	382.7	4.3644	**380.1**	4.6511			486.1	**2.9660**
GE	507.4	5.2840	**267.9**	1.8282	323.5	3.3446			306.7	**1.5693**
XOM	**240.9**	2.4817	405.1	3.7342	252.1	**2.2310**			433.9	2.9304
	Panel b. Annual returns (252 trading days)									
DJIA	509.2	10.2179			**402.8**	7.7456	449.7	9.3534	543.3	**6.5334**
SPY	564.6	8.7737			558.3	8.2382	559.4	8.5817	**498.6**	**5.1761**
GE	448.7	9.4501			347.4	7.9931	414.6	9.0437	**245.8**	**5.5579**
XOM	**290.4**	**2.9040**			365.8	4.6234	305.8	3.3664	388.7	4.0537

Note: the distribution that is the best fit according to the criterion is highlighted with bold.

Conclusions

The modern portfolio theory is not as bad as its critics say. While Nassim Taleb and others are partially right in their claims that the normal distribution is absolutely inappropriate for describing the daily log-returns, it doesn't mean the theory itself is flawed. As it has been shown in this article, the Student's t-distribution or the Laplace distribution can be used to describe the distribution of daily log-returns and address the famous problem of the extreme outliers ("Black Swans"). While the problem is not completely eliminated, the discrepancy between the actual probability and the theoretically predicted one becomes quite minor, thousands and even millions times less than it is when using the normal distribution. The modern portfolio theory allows both the Student's t-distribution and the Laplace distribution to be used for describing the returns without altering its assumptions. Based on my research, I recommend using either the t-distribution with 3 to 4 degrees of freedom, or the Laplace distribution (with the sample median and the mean absolute deviation from it as parameters) to model the daily log-returns.

Unlike daily returns, there's no clear preference for modeling periodic log-returns. While the t-distribution fits the Monte-Carlo simulated data better, the Laplace distribution seems to do a much better job when describing the historical returns. On the other hand, Black Swans problem is not such a big issue for periodic returns. The difference between the actual and the theoretic probability is moderate, and thus the risks related to investing over holding periods longer than 1 month should be manageable even if we assume the normal distribution of log-returns.

References

1. Taleb, Nassim. “The Black Swan: The Impact of the Highly Improbable.” Random House, 2007.
2. Chamberlain, Gary. “A characterization of the distributions that imply mean-variance utility functions.” Journal of Economic Theory, Vol. 29, Issue 1 (1983), pp. 185–201.
3. Owen, Joel and Ramon Rabinovitch. “On the Class of Elliptical Distributions and their Applications to the Theory of Portfolio Choice.” The Journal of Finance, Vol. 38, No. 3 (1983), pp. 745–752.