Generalizing the Sharpe Ratio and Infinite Divisibility in Financial Markets

Hal M. Switkay
Goldey-Beacom College

Investments with high Sharpe ratios tend to be fixed income investments, particularly short-term bonds. Such securities dominate income investment, but not growth investment. We generalize the Sharpe ratio to interpolate smoothly between risk-averse and risk-seeking behaviors, with the traditional Sharpe ratio as a special case. The generalized ratios may be used to optimize portfolios within a lifestyle strategy or a lifecycle strategy. We demonstrate that the Sharpe ratio has an interpretation connected to value at risk (VaR), if returns are normally distributed. We study broad market indices and indicators for US stocks, US corporate bonds, gold, and Japanese stocks. Using monthly data obtained from FRED, we show that these datasets are neither normal, nor log-normal, nor stable. We argue that none of these distributions is appropriate for market data. We argue instead for using infinitely divisible distributions, and fit Pearson distributions to the data.

Keywords: Sharpe ratio, value at risk, stable distribution, infinitely divisible distribution, Pearson distribution, portfolio theory

THE SHARPE RATIO AND ITS GENERALIZATION

The Sharpe ratio (Sharpe, 1970) is one of several attempts to quantify the tradeoff between return and risk. It equates risk with variability of return, defined by the standard deviation. It postulates a constant price for return, in the sense that for securities with the same Sharpe ratio, each percentage point increase in return is paid for by a constant percentage point increase in variability.

Issues are confused further by the postulation that a risk-free asset has zero standard deviation over a single period, when in fact such assets, like short-term US Treasury bills, are indeed subject to fluctuating returns over extended periods of time. Of course the Sharpe ratio concerns not absolute returns, but rather excess returns above the risk-free return. We could argue that the standard deviation of a set of zero excess returns is itself zero. However, the mean of that set would also be zero, but classic mean-variance analysis does not declare the mean return of the risk-free asset to be zero axiomatically.

In this paper, we are concerned most with the assumption of a constant price for return. We shall calculate Sharpe ratios for broad asset classes without taking account of the risk-free return (equivalently, setting the risk-free return to zero). We justify this by noting that the results, which heavily favor the most conservative asset classes (those with lowest mean return), would be even further exaggerated by subtracting out the risk-free rate.

Our data source was FRED, the research arm of the St. Louis Federal Reserve Bank (see FRED time series). We took monthly average prices from January 1973 through June 2021, calculating relative changes, for the following series: the Wilshire 5000 total market full cap US stock index; the gold fixing price in London; the ICE BofA US corporate bond total return index; and the Nikkei 225 stock index. For comparison, we included one-twelfth the yield on the US 1-year Treasury note. Based on this, we calculated the monthly Sharpe ratios in table 1, listed in decreasing order. Annualized Sharpe ratios are computed by multiplying monthly Sharpe ratios by $\sqrt{12}$.

TABLE 1
SHARPE RATIOS OF BROAD ASSET CLASSES

	mean	standard deviation	monthly Sharpe ratio	annualized Sharpe ratio
Treasury 1-year note	0.0042	0.0108	0.3842	1.3310
US Corporate Bond	0.0063	0.0182	0.3453	1.1962
Wilshire 5000 Stock	0.0095	0.0402	0.2370	0.8209
Gold	0.0070	0.0502	0.1389	0.4812
Nikkei 225	0.0040	0.0461	0.0868	0.3007

For completeness, we provide the correlations between monthly returns in table 2.

TABLE 2
CORRELATIONS BETWEEN BROAD ASSET CLASSES

	Treasury 1-year note	US Corp Bond	Wilshire 5000 Stock	Gold	Nikkei 225
Treasury 1-year note	1.0000	0.0323	-0.0276	-0.0286	-0.0096
US Corp Bond	0.0323	1.0000	0.2759	-0.0208	0.0882
Wilshire 5000 Stock	-0.0276	0.2759	1.0000	-0.0239	0.5150
Gold	-0.0286	-0.0208	-0.0239	1.0000	-0.0444
Nikkei 225	-0.0096	0.0882	0.5150	-0.0444	1.0000

We note for the record that only gold has negative correlations with each of the other four asset classes; 1-year Treasury notes have negative correlations with three of the other asset classes.

Based on Sharpe ratios alone, the prudent investor would be advised to focus mainly on fixed-income securities. A portfolio combining these assets should be able to increase the overall Sharpe ratio, but it would rely the most on fixed-income securities, together with gold.

We seek a way to incorporate information about mean return and standard deviation that also accounts for different risk tolerance profiles. Let μ be the mean return on a security, and let σ be the standard deviation of returns. An investor who is indifferent to risk will seek to maximize μ, no matter the value of σ. A cautious investor, for whom the price of return is constant, will seek to maximize the quotient μ/σ, which is the Sharpe ratio.

To interpolate smoothly between these behaviors, choose a value of λ between 0 and 1 inclusive; that is, $0 \leq \lambda \leq 1$. Then the generalized Sharpe ratio with exponent λ is defined as μ/σ^{λ}. When $\lambda = 0$, the ratio reduces to μ, the mean return itself. At the opposite extreme, when $\lambda = 1$, the ratio is the traditional Sharpe ratio. Note that because returns are dimensionless in currency units, the generalized Sharpe ratio is also dimensionless in currency units for every value of λ in $[0,1]$.

Using the values of the mean and standard deviation of returns of the five asset classes in table 1, it is easy to derive a graph of the monthly generalized Sharpe ratios for all exponents simultaneously. Annualized generalized Sharpe ratios are obtained by multiplying the monthly figures by $12^{1-\lambda/2}$. These exponential functions of λ have linear graphs when the ratios are graphed on a logarithmic scale, since $\ln(\mu/\sigma^{\lambda}) = \ln(\mu) - \lambda \ln(\sigma)$. See figure 1.

FIGURE 1
GENERALIZED SHARPE RATIOS BY ASSET CLASS

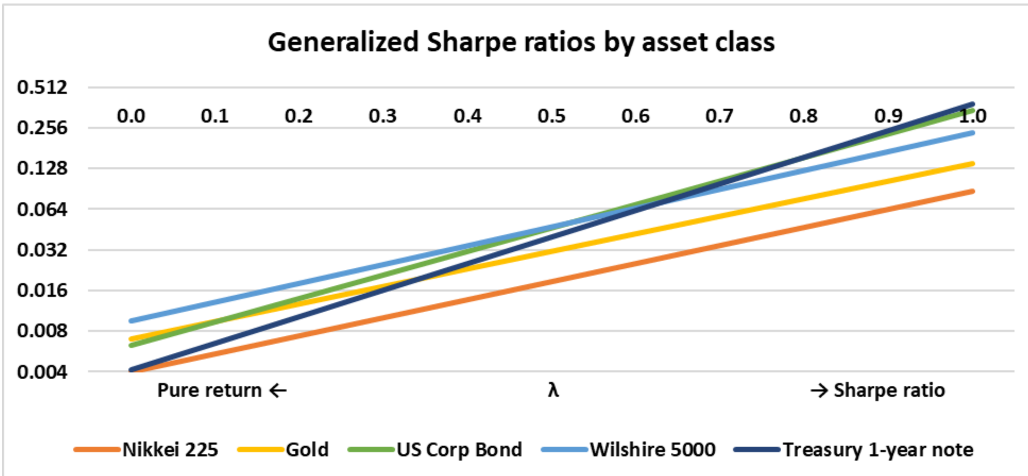

For comparison, we also investigated the 3-year mean return and standard deviation of the 500 largest ETFs, with data from the Fidelity Investments website (see Fidelity). For the values $\lambda = 0, 0.5, 1$, we enumerated the 10 ETFs among these 500 with the largest generalized Sharpe ratios. For both $\lambda = 0, 0.5$, the top 10 performers were technology sector equity funds. For $\lambda = 1$, the top 10 performers were fixed-income funds focusing on T-bills, short-term T-notes, and TIPS.

The investor aiming for moderate growth, or growth with income, might seek to maximize the generalized Sharpe ratio with an intermediate value of λ. The values for the five asset classes are given in table 3, ordered from largest to smallest, for $\lambda = 0.5$.

TABLE 3
THE GENERALIZED SHARPE RATIO FOR $\lambda = 0.5$

	Wilshire 5000	US Corp Bond	Treasury 1-year note	Gold	Nikkei 225
monthly	0.0475	0.0466	0.0400	0.0311	0.0186
annual	0.3065	0.3005	0.2580	0.2007	0.1201

A portfolio optimizing the generalized Sharpe ratio for an intermediate value of λ might well resemble a typical balanced portfolio, or a lifestyle moderate growth fund. A lifecycle fund manager might choose to vary λ throughout the customer's life, rebalancing at regular intervals. If the money is needed in Y years, where $0 \leq Y \leq 50$, the manager might choose to optimize the portfolio with $\lambda = 1 - Y/50$. An even better choice is $\lambda = (1 - Y/50)^2$, which allows the investment to grow more. Software exists to identify portfolios that optimize the traditional Sharpe ratio. The manager of this fund might run the existing software using the given values of μ, and merely substituting σ^{λ} for the standard deviation.

We used Excel Solver to determine portfolios that optimized the generalized Sharpe ratio for $\lambda = 0, 1/4, 1/3, 3/8, 2/5, 1/2, 2/3, 3/4, 1$. The permitted investments were the Wilshire 5000 index, the

US corporate bond index, gold, and the Nikkei 225 index. Allocations were fixed from January 1973 to June 2021, in an effort to create a "set it and forget it" portfolio. The recommended allocations appear in table 4 and in figure 2.

TABLE 4
ALLOCATIONS TO OPTIMIZE GENERALIZED SHARPE RATIO

λ	Nikkei 225	Gold	US corporate bond	Wilshire 5000	μ	σ
0.0000	0.0000	0.0000	0.0000	1.0000	0.0095	0.0402
0.2500	0.0000	0.0000	0.0000	1.0000	0.0095	0.0402
0.3333	0.0000	0.1489	0.0000	0.8511	0.0092	0.0348
0.3750	0.0000	0.1423	0.4580	0.3997	0.0077	0.0211
0.4000	0.0000	0.1392	0.5039	0.3569	0.0076	0.0201
0.5000	0.0000	0.1330	0.5941	0.2729	0.0073	0.0185
0.6667	0.0000	0.1288	0.6573	0.2140	0.0071	0.0176
0.7500	0.0000	0.1275	0.6756	0.1969	0.0070	0.0174
1.0000	0.0000	0.1252	0.7099	0.1649	0.0069	0.0171

FIGURE 2
OPTIMAL PORTFOLIO ALLOCATIONS

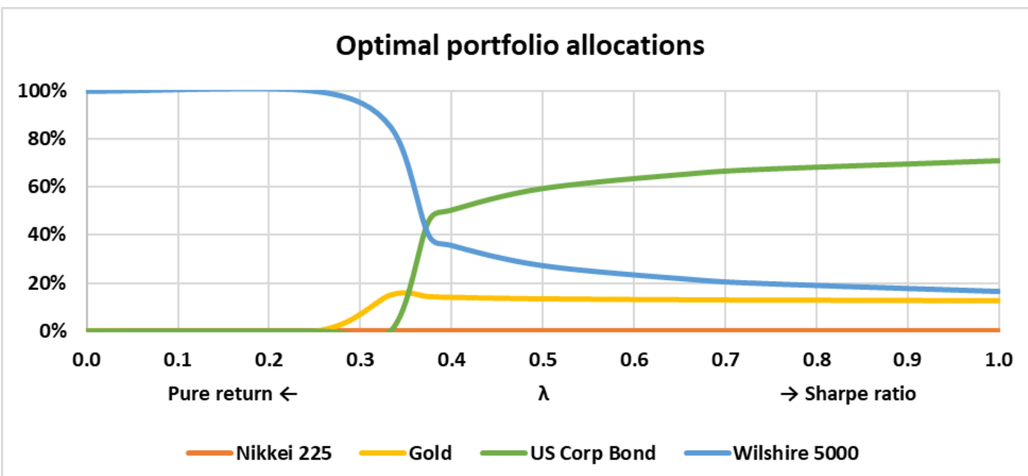

Figure 3 depicts the performance of each of the four asset classes starting in January 1973 at 100, together with the portfolio that optimizes the Sharpe ratio ($\lambda = 1$), the portfolio that optimizes the generalized Sharpe ratio with $\lambda = 0.5$, and two further economic indicators for comparison: the Consumer Price Index, and disposable personal income per capita.

FIGURE 3
PERFORMANCE OF ASSET CLASSES, PORTFOLIOS, AND ECONOMIC INDICATORS

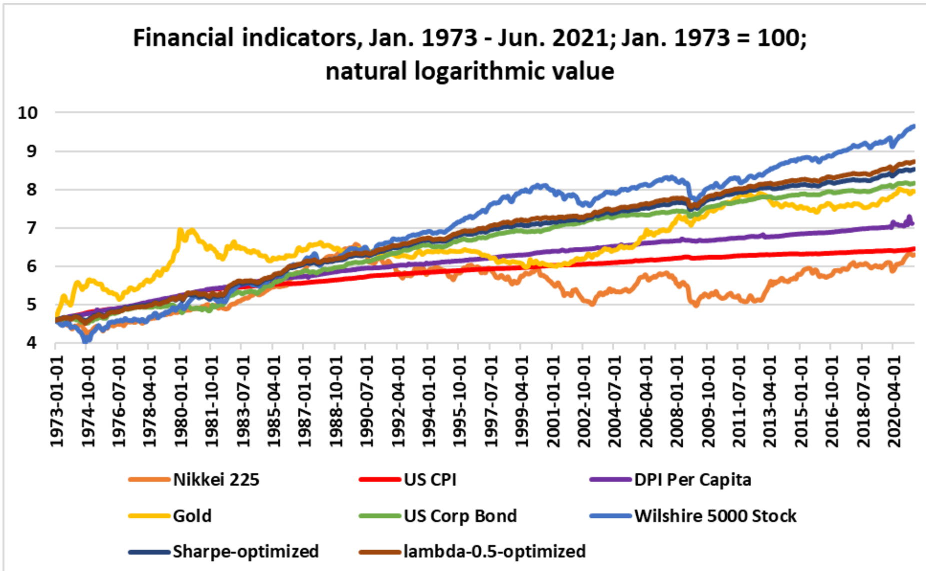

Figure 4 shows mean return vs. standard deviation for all of the optimal portfolios. As λ takes on all values in $[0,1]$, these points trace out a curve, although, unlike the Markowitz efficient frontier, there is no reason to believe this curve is hyperbolic in general. We tentatively identify optimal portfolios for varied values of λ with investment strategies as follows: 1/4, aggressive growth; 1/3, growth; 2/5, moderate growth or growth with income; 1, income. We note that the $\lambda = 2/5$ optimal portfolio has very little more variability than the Sharpe ratio optimal portfolio, but it shows increased return. Similarly, the $\lambda = 1/3$ optimal portfolio has very little less return than the highest performing asset (Wilshire 5000), but it shows decreased variability.

FIGURE 4
MEAN VS. STANDARD DEVIATION FOR PORTFOLIOS OPTIMIZING GENERALIZED SHARPE RATIOS

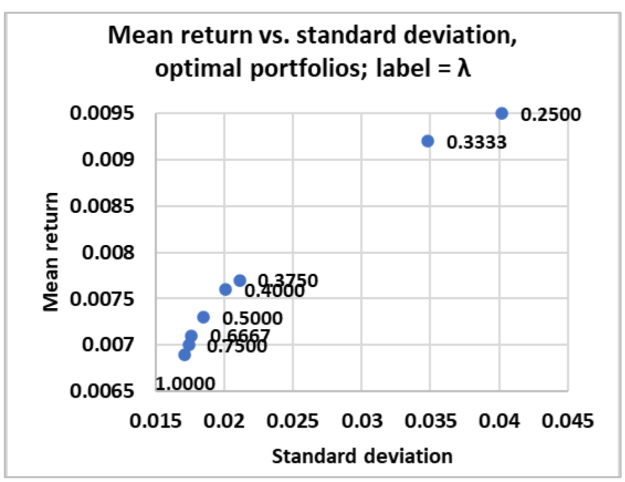

The terminal value for all portfolios optimized for the generalized Sharpe ratios in June 2021 is shown in figure 5, starting with a value of 1 in January 1973.

FIGURE 5
ENDING VALUE OF OPTIMAL PORTFOLIOS

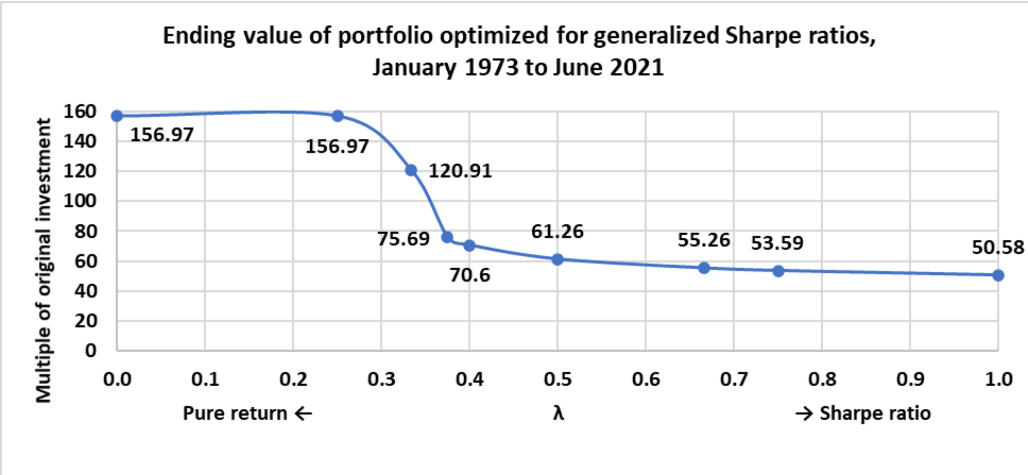

THE SHARPE RATIO AND VALUE AT RISK (VAR)

The traditional Sharpe ratio has an interesting interpretation connecting it to value at risk (VaR) when returns are normally distributed. Let a random variable x have mean μ and standard deviation σ. Define the z-score of x as follows: $z = (x - \mu)/\sigma$. If x signifies returns, or excess returns, on a security, then the z-score of zero is $(0 - \mu)/\sigma = -\mu/\sigma$, which is precisely the opposite of the traditional Sharpe ratio. Let us call the traditional Sharpe ratio SR.

Let F be the cumulative distribution function for x, where x signifies returns. For α in the interval $(0,1)$, we define the α-VaR for x as follows: $VaR_\alpha(x) = -\inf\{x|F(x) > \alpha\} = F^{-1}(\alpha)$ when F is continuous. We deduce that $VaR_\alpha(x) = 0 \leftrightarrow F^{-1}(\alpha) = 0 \leftrightarrow F(0) = \alpha$.

Let Φ be the cumulative distribution function of a standard normal random variable z. If x is normally distributed, then $F(x) = \Phi(z)$, so $F(0) = \Phi(-\mu/\sigma)$. Thus if x is normally distributed, $VaR_\alpha(x) = 0 \leftrightarrow \Phi(-SR) = \alpha$. More briefly, $VaR_{\Phi(-SR)}(x) = 0$.

This intriguing equation could be taken as a definition of the traditional Sharpe ratio. It implies that a fundamental metric of interest to investors is the fraction of the time that a security has negative returns, no matter what its distribution; perhaps this metric, or its complement, should be added to prospectuses and similar documents. This enables distribution-free analysis of security risk. Table 5 enumerates the fraction of the time that the major asset classes had negative monthly returns from January 1973 to June 2021.

TABLE 5
FRACTION OF MONTHLY RETURNS THAT ARE NEGATIVE

1-year Treasury note	US corporate bond	Wilshire 5000	Nikkei 225	Gold
0.0000	0.2940	0.3510	0.4510	0.4880

The ordering of asset classes does not match table 1 completely, and that is because returns are not normally distributed. Gold has a higher frequency of negative returns than does the Nikkei average, but its Sharpe ratio is larger, due to gold's far higher kurtosis.

DISTRIBUTIONS OF COMMON DATA SETS IN FINANCE

The choice of distributions to model real-world phenomena can have an enormous impact on one's conclusions. Incorrect choices in finance can lead to trillion-dollar economic catastrophes, as the likelihood of extreme events is not properly estimated. The financial world has relied on assumptions by one scholar after another that are not supported by reason or observation. It is time to correct these mistakes.

In modeling a random variable by a probability distribution, by far the most important point is to determine the theoretical support of the random variable: the maximum open interval in which the variable theoretically could take values. Return on an ordinary, non-leveraged security can never be less than $-100\% = -1$, and can attain arbitrarily large values. The proper support for such a distribution is the semi-infinite interval $(-1, \infty)$. Therefore, neither the normal distribution nor any other probability distribution supported on the full real number line is appropriate to model returns on investment.

The (shifted) log-normal distribution is not appropriate to model stock returns, which have negative skewness. The log-normal distribution has positive skewness and heavy kurtosis.

Since Mandelbrot (1963), much attention has been focused on stable distributions to model returns on investment. This is not justified either by data, by practice, or by theory. Using data from FRED, we examined the monthly price of cotton, as Mandelbrot did, from January 1990 to July 2021. Unlike Mandelbrot, we did not find a severe departure from normality. Skewness was 0.0915, almost completely symmetrical. Kurtosis was 3.1352, not terribly large, and comparable to a Student's *t*-distribution with 6 degrees of freedom. The histogram is shown in figure 6.

FIGURE 6
HISTOGRAM: RELATIVE MONTHLY COTTON PRICE CHANGES

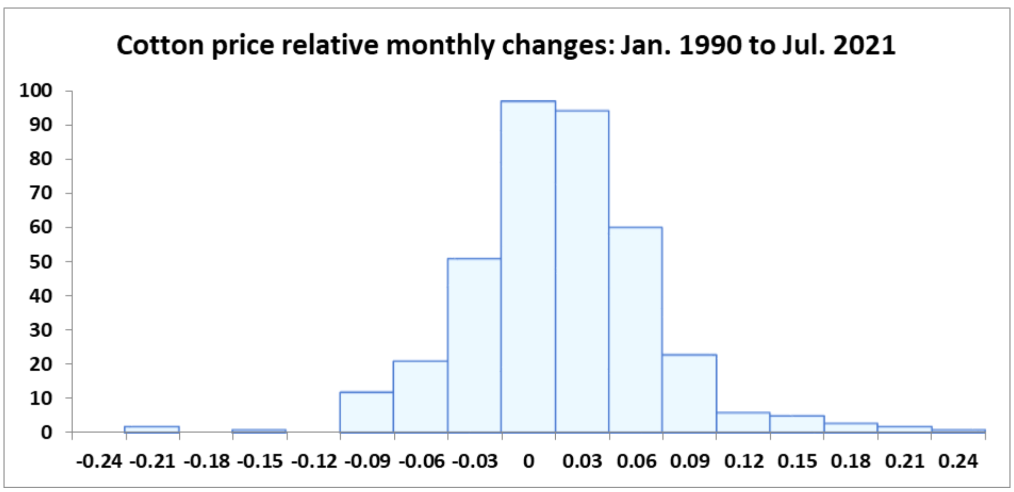

Non-normal stable distributions have infinite standard deviations. Finite samples from such distributions always have finite standard deviations. This raises the question of how to determine when the distribution from which one is sampling has an infinite population standard deviation.

We examined monthly price changes in the Wilshire 5000 index, and computed standard deviations over all rolling 12-month, 36-month, 120-month, and 360-month periods from January 1973 to June 2021. Next, we computed the average sample standard deviation for each of the four window lengths. The results were quite close: 0.0361 for 12 months, 0.0372 for 36 months, 0.0381 for 120 months, and 0.0379 for 360 months. The largest of these average standard deviations is only 5.54% more than the smallest. Moreover, the standard deviation of the sample standard deviations decreases as the window lengthens, showing that the estimates are converging. Figure 7 shows side-by-side boxplots of the natural logarithm of the rolling standard deviations for 1-year, 3-year, 10-year, and 30-year periods. Figure 8 shows the rapid convergence of the cumulative standard deviation of the monthly returns, typical only of distributions with finite standard deviations.

FIGURE 7
NATURAL LOGARITHM OF STANDARD DEVIATION OF MONTHLY RETURNS, WILSHIRE 5000 INDEX

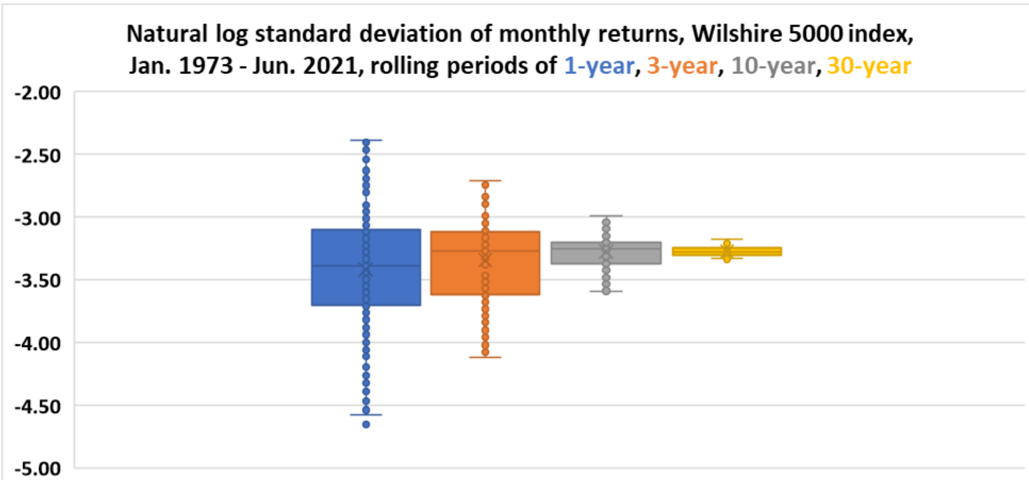

FIGURE 8
CUMULATIVE STANDARD DEVIATION OF MONTHLY RETURNS, WILSHIRE 5000 INDEX

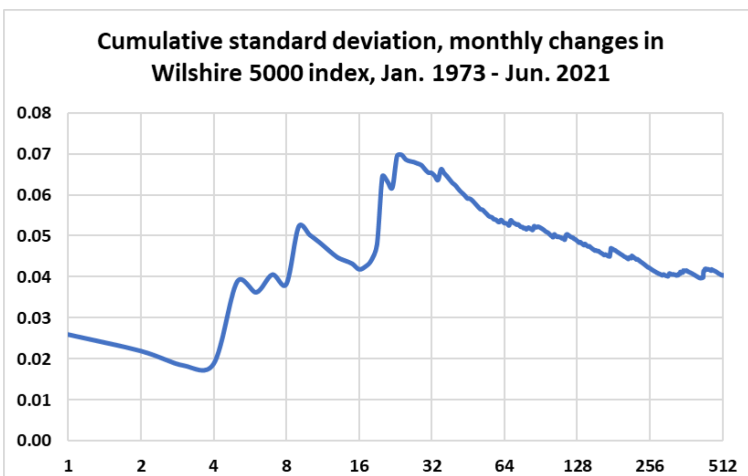

For further verification, we computed annual price changes in the Wilshire 5000 index, and computed standard deviations over rolling 10-year and 30-year periods. The average of such standard deviations was 0.1358 for 10 years, and 0.1398 for 30 years. Again, these values are quite close to each other and to the overall standard deviation of annual price changes, 0.1385. Finally, we observed that these values are close to the standard deviation of monthly changes, multiplied by $\sqrt{12}$, precisely what one would expect if the observations were chosen independently from the same distribution with finite standard deviation.

We do not find the claim of an infinite standard deviation credible for this distribution. Hence, we must reject the claim that stock index price changes follow a stable but non-normal distribution. We explained earlier why such price changes cannot be normally distributed. Therefore, we conclude that stock index price changes do not follow any stable distribution.

This conclusion is supported by the research of Xian Ning Yan. Yan studied the leading stock market indices of America, Asia, and Europe. Yan used statistical techniques to demonstrate that most of these indices have finite standard deviations, but possibly infinite kurtosis.

This raises the question of why it was ever believed that such return distributions must follow a stable distribution. This belief appears to be grounded in Mandelbrot's 1963 paper on cotton prices, and was reinforced decades later in his popular book on the subject (Mandelbrot and Hudson, 2004). In essence, Mandelbrot postulated that graphs of prices have a self-similar nature, an invariance as to scale that was the key feature of Mandelbrot's work on fractal analysis and geometry. In other words, with the axis scale removed, one could not distinguish between a one-day stock chart divided into minutes, and a one-decade stock chart divided into days. That assertion does not appear to be true. Furthermore, Mandelbrot's identification of the heavy tails of financial datasets is not unique to stable distributions; other distributions have heavy tails as well.

Stable distributions are defined below. Stable distributions have been completely classified. With only three exceptions, the general stable distribution lacks a probability density function that can be expressed in closed form. The infinite standard deviation of non-normal stable distributions makes them unsuitable for analysis of many financial markets. Ironically, a remark in Mandelbrot's paper on cotton prices identifies the central property that allows us to discover a broader, more flexible, and more useful family of probability distributions: infinite divisibility.

Let X be a random variable. Let $Y_1, \ldots, Y_n$ be independent, identically distributed random variables, with $Y_i \sim Y$. The conceptual definition of stability can be written symbolically as follows: X is stable if whenever $Y = X$, there exist constants a_n and b_n such that $\sum Y_i \sim a_n + b_n X$. Infinite divisibility is defined as follows: X is infinitely divisible if for every n, we can find Y such that $\sum Y_i \sim X$.

Evidently, every stable distribution is infinitely divisible; simply choose Y to be $(X - a_n/n)/b_n$. Other well-known distributions with finite standard deviations are infinitely divisible as well, including broad sub-families of the Pearson family: Student's t-distributions with any number of degrees of freedom; gamma and beta prime distributions with any combination of parameters. Note that some beta prime distributions have negative skewness, as we observe for many securities. Since infinitely divisible distributions are closed under limits, inverse gamma distributions are infinitely divisible too. Distributions with bounded support are not infinitely divisible. If x belongs to the support of a continuous, infinitely divisible distribution, so does kx for all positive k. Thus, the support of a continuous, infinitely divisible distribution can only be $(-\infty, 0)$, $(0, \infty)$, or $(-\infty, \infty)$. There is a sense in which the tails of an infinitely divisible distribution must be at least as heavy as those of a normal distribution. Steutel and Bose et al are two basic references on infinitely divisible distributions.

Figure 9 is a diagram of the Pearson distribution system, projected onto the skewness-kurtosis plane. Types II (symmetric: orange line) and I (yellow) have bounded support: beta distributions. Types III (gamma: red line), VI (beta prime: pink), and V (inverse-gamma: blue line) have semi-infinite support. Types IV (light turquoise) and VII (Student's t: turquoise line) have doubly infinite support. Types III, VI, V, and VII are infinitely divisible. It is not known whether type IV is infinitely divisible.

FIGURE 9
THE PEARSON DISTRIBUTION SYSTEM

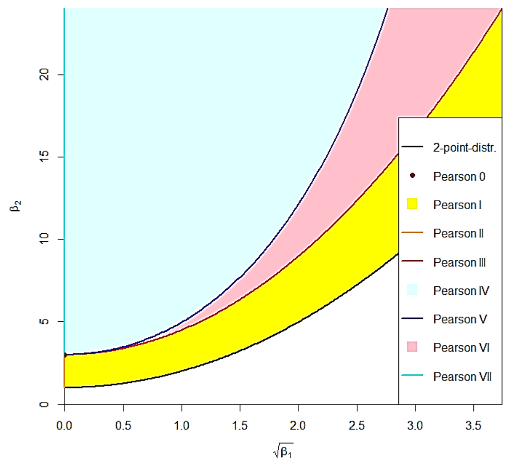

These observations suggest searching among the Pearson distributions, or perhaps related families like the generalized beta distributions, to model financial phenomena. We wish to test whether such distributions are good fits for individual securities, which might be expected to be far more variable and heavy-tailed than broad market indices. To that end, we study monthly price changes for two stocks: Exxon-Mobil (XOM) and Tesla (TSLA). We downloaded monthly price data from Yahoo Finance for the months July 2010 through August 2021. The price for each month was calculated as the geometric mean of the open, high, low, and close for that month. Relative price changes were calculated. For this dataset, $N = 133$. The correlation between the two variables was 0.0020, not significantly different from zero. Summary information for this dataset appears in table 6.

TABLE 6
SUMMARY STATISTICS FOR XOM AND TSLA

	XOM	TSLA
minimum	-0.2359	-0.2071
median	-0.0002	0.0243
mean	0.0010	0.0473
maximum	0.1433	0.5799
standard deviation	0.0510	0.1350
skewness	-0.6822	1.0420
kurtosis	3.4017	1.4173

TSLA is far more variable than XOM. We are particularly interested in the contrast in skewness; XOM is notably negatively skewed, while TSLA is notably positively skewed. Figure 10 provides the scatterplot matrix with histograms along the diagonal.

FIGURE 10
SCATTERPLOT MATRIX PLUS HISTOGRAMS FOR XOM AND TSLA

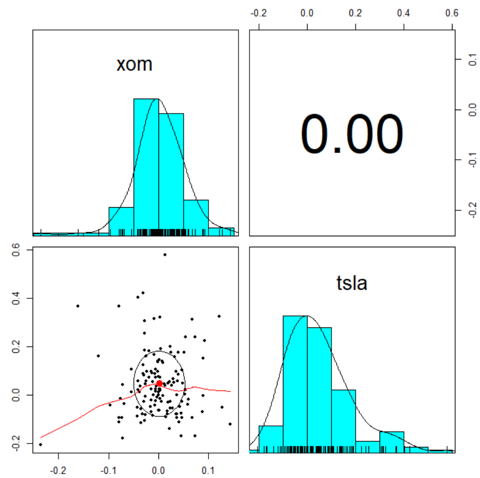

Next, we used the R package PearsonDS to find the best fit Pearson distribution for the two stocks' monthly returns. We focused only on Pearson distributions of type VI, the beta prime, on the grounds that returns theoretically have semi-infinite support $(-1, \infty)$, and that the limiting cases type III (gamma) and type V (inverse gamma) form a subset of measure zero.

To measure the quality of the fit, we constructed QQ-plots, and computed the correlation of observed quantiles to predicted quantiles. We also compared the extreme observations with the extremes of the fitted data. The QQ-plots are shown in figure 11.

FIGURE 11
QQ-PLOTS FOR XOM, TSLA AGAINST BEST-FIT PEARSON TYPE VI DISTRIBUTIONS

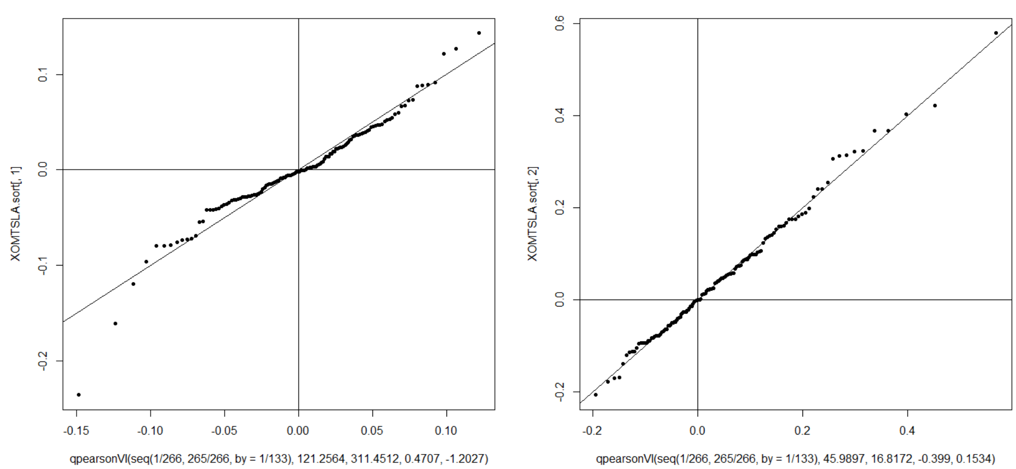

The correlations in the QQ-plots are 0.9754 for XOM and 0.9974 for TSLA. The fit for TSLA is excellent. The fit for XOM, however, is not good at the lower extremity. The maximum monthly decrease predicted by the model is 14.83%. In reality, the maximum monthly decrease was 23.59%, an error of 8.76% occurring one month in 11 years. Money managers and their clients must determine whether that size prediction error represents an acceptable risk, based on their risk tolerance.

Our use of Pearson distributions to model financial datasets was based on the knowledge that the Pearson distributions with semi-infinite support are known to be infinitely divisible. Two further things must be borne in mind. First, the infinitely divisible distributions have not been completely classified (although there is an abstract characterization of them). There could be other useful families of infinitely divisible distributions.

Second, observations supersede theory. The theory that financial datasets could be modeled well by stable distributions has been seen to be untenable. It may be that the assumption of infinite divisibility will need to be sacrificed one day as well. We need to describe what we observe, not what we prefer to see.

ACKNOWLEDGMENTS

The author wishes to thank his colleague Robert McConkie for lending him his copy of Sharpe's textbook, William Sharpe for a private conversation, and his wife Yan for her support and encouragement.

REFERENCES

Bose, A., Dagupta, A., & Rubin, H. (2002). A Contemporary Review and Bibliography of Infinitely Divisible Distributions and Processes. *Sankhya: The Indian Journal of Statistics*, *64*(3), 763–819.

Fidelity Investments ETF screener (member-only access). (n.d.). Retrieved August 2021, from https://research2.fidelity.com/pi/etf-screener

FRED, the research arm of the St. Louis Federal Reserve Bank. (n.d.). Time series. Retrieved August, 2021, from https://fred.stlouisfed.org/series/A229RX0, https://fred.stlouisfed.org/series/BAMLCC0A0CMTRIV, https://fred.stlouisfed.org/series/CPIAUCSL, https://fred.stlouisfed.org/series/GOLDPMGBD228NLBM, https://fred.stlouisfed.org/series/GS1, https://fred.stlouisfed.org/series/NIKKEI225, https://fred.stlouisfed.org/series/PCOTTINDUSDM, https://fred.stlouisfed.org/series/WILL5000INDFC

Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *The Journal of Business*, *36*(4), 394–419.

Mandelbrot, B., & Hudson, R.L. (2004). *The (Mis)behavior of Markets*. New York, NY: Basic Books.

Sharpe, W.F. (1970). *Portfolio Theory and Capital Markets*. New York, NY: Mc-Graw Hill.

Steutel, F.W. (1979). Infinite Divisibility in Theory and in Practice. *Scandinavian Journal of Statistics*, *6*(2), 57–64.

Yahoo Finance. (2021). Time series. Retrieved August 2021, from https://finance.yahoo.com/quote/TSLA?p=TSLA, https://finance.yahoo.com/quote/XOM?p=XOM

Yan, X.N. (2008). *Test for Infinite Variance in Stock Returns*. Retrieved September 10, 2021, from https://ink.library.smu.edu.sg/cgi/viewcontent.cgi?article=1038&context=etd_coll