POWER-LOG OPTIMIZATION FOR OPTIONS AND STOCKS VS. MEAN-VARIANCE OPTIMIZATION

Jivendra K. Kale, School of Economics and Bus. Ad., Saint Mary's College of California
1928 Saint Mary's Road, Moraga, CA 94556. (925) 631-4128, jkale@stmarys-ca.edu
Arnav Sheth, Graduate Business Programs, Saint Mary's College of California
Arnav Sheth, School of Economics and Bus. Ad., Saint Mary's College of California
1928 Saint Mary's Road, Moraga, CA 94556. (925) 631-8707, aas3@stmarys-ca.edu

ABSTRACT

Kale (2006) introduced a portfolio selection method based on Power-Log utility functions, and demonstrated its effectiveness simulated returns. We use empirical data from 1997 through 2009 for a treasury security, the S&P500 index and a call option on the index to show that optimal Power-Log portfolios have consistently low tail risk, and are also very effective in protecting against large negative surprises in returns. The geometric average realized return over this period is always positive for optimal Power-Log portfolios for all levels of risk, while it is negative for the riskier M-V efficient portfolios and -100% for the riskiest one.

INTRODUCTION

Mean-variance analysis has been the standard method for portfolio selection since Harry Markowitz (1953) introduced the idea of a mean-variance efficient portfolio. It imposes strong assumptions of either normality on asset returns, or quadratic utility for investor preferences when asset returns are non-normal. Many assets, such as bonds, real estate and options, have non-normal returns, and for portfolios containing such assets mean-variance analysis is questionable. Other portfolio selection methods that are used when asset returns are non-normal, such as mean-semivariance analysis suggested by Markowitz (1959, 1991), or mean conditional VaR optimization described by Xiong and Idzorek (2011), are explicitly designed to control downside risk, but these methods are not grounded in utility theory. Kale (2006) introduced Power-Log utility functions, which can be used for portfolio construction when asset returns are normal, or non-normal, and they satisfy the Friedman-Savage (1948) axioms that define risk-averse utility functions, thus putting them on a firm theoretical foundation. Power-Log utility functions combine multiperiod portfolio theory and the tenets of behavioral finance that treat gains and losses differently to represent investor preferences. While Kale (2006) laid the conceptual foundation for Power-Log optimization with simulated asset returns, here we use empirical data to test this portfolio selection methodology in the real world.

METHODOLOGY

The Markowitz mean-variance methodology for constructing investment portfolios is well known, so we will describe the methodology for constructing optimal portfolios with Power-Log utility functions. Portfolio selection with a utility function uses the expected utility criterion developed by Von Neumann and Morgenstern (1944) and Savage (1964). The following description of the log, power and Power-Log utility functions, the methodology for constructing the joint return distribution, and the resulting optimal portfolios is based on Kale (2006). Each Power-Log utility function is a two-segment utility function, where the utility of gains is modeled with a log utility function and the utility of losses is modeled with a power utility function with power less than or equal to zero. It combines the maximum growth characteristics of the log utility function on the upside, with the scalable downside protection characteristics of the power function on the downside. It is defined as,

$$
\begin{aligned}
U &= \ln(1+r) && \text{for } r \geq 0 \\
&= \frac{1}{\gamma}(1+r)^{\gamma} && \text{for } r < 0
\end{aligned}
\tag{1}
$$

where,

r	portfolio return
γ	power, is less than or equal to 0

Power-Log utility functions conform to the Kahneman and Tversky postulates of reference dependence, loss aversion, and diminishing sensitivity for gains. Investors can vary the level of downside protection they build into their portfolios by changing the downside power. Selecting a downside power of zero is equivalent to using a log utility function for losses, which will result in the construction of the maximum growth portfolio, since the utility function for gains is always a log utility function. Lower values of the downside power represent greater loss aversion since the penalty for losses increases, while the value associated with gains is left unchanged.

Kale (2006) showed the effectiveness of this portfolio construction method by simulating returns for three assets with very different types of return distributions, a riskless asset, a stock index and a call option on the stock index, for a single period. Our empirical test extends Kale's simulation study by using market data and annual rebalancing of a portfolio that contains a treasury security, the S&P500 index, and the closest to the money call option on the S&P500 index with approximately one year to expiration. The data is for the years 1996-2009, which includes periods with wide swings in the stock market as shown in Table I. This data allows us to build optimal portfolios at the beginning of each of 13 one-year holding periods and evaluate their realized returns. Portfolio rebalancing and the time to maturity for the treasury are synchronized with the expiration dates of the call options. For each year in the sample period, the portfolio is purchased on the December expiration date of the call option and sold on the December expiration date of the following year, which is the rebalancing date. For example, for the first holding period an optimal portfolio is constructed and purchased on December 20, 1996, the December expiration date in 1996; it is held for about a year and sold on December 19, 1997, the December expiration date in 1997, when a new optimal portfolio is constructed and purchased for the following year. To construct the optimal portfolio at the beginning of this one-year holding period, we select the closest to the money call option that expires on December 19, 1997, for inclusion in the portfolio. The yield on a treasury that matures close to the December 19, 1997 expiration date is used to calculate the riskless return for the holding period. To calculate portfolio values we use market prices for the S&P500 index at the beginning and end of the holding period. For the call option we use market price at the beginning of the holding period, and expiration value at the end of the holding period, because the expiration value gives us a more reliable valuation of the call than the market price on the expiration date. The realized asset returns for all the holding periods are shown in Table I.

Table I

Holding Period	Start Date	End Date	Riskless Return (%)	S&P500 Return (%)	Call Option Return (%)	B-S Implied Volatility (%)

1	12/20/1996	12/19/1997	5.01	28.43	202.74	18.75
2	12/19/1997	12/18/1998	5.32	27.09	135.67	23.22
3	12/18/1998	12/17/1999	4.48	20.97	67.61	25.84
4	12/17/1999	12/15/2000	5.42	-6.51	-100.00	24.58
5	12/15/2000	12/21/2001	6.12	-11.50	-100.00	24.21
6	12/21/2001	12/20/2002	1.71	-20.39	-100.00	22.39
7	12/20/2002	12/19/2003	1.21	23.32	118.86	25.79
8	12/19/2003	12/17/2004	0.88	11.30	59.66	15.81
9	12/17/2004	12/16/2005	2.19	7.74	-5.71	15.14
10	12/16/2005	12/15/2006	3.92	14.36	80.90	15.14
11	12/15/2006	12/21/2007	4.99	5.81	-42.50	14.22
12	12/21/2007	12/19/2008	2.99	-38.32	-100.00	23.37
13	12/19/2008	12/18/2009	0.02	27.39	67.89	40.09
		Geometric Average	3.39	4.75	-100.00	

To construct an optimal portfolio with a Power-Log utility function, we need the joint return distribution for all assets that can be included in the portfolio. The key distribution for our empirical test is the distribution of returns for the S&P500 index. We examined the shape of the distribution of annual S&P500 returns with the goal of using it as a basis for simulating future returns at the beginning of each holding period. We performed three separate EDF tests on annual S&P 500 returns from 1950 through 2011, Lilliefors, Jarque-Bera and χ^2 with the null hypothesis that the return distribution is lognormal, and got p-values of 0.7107, 0.3371 and 0.6797 respectively. Clearly, we cannot reject the hypothesis of lognormality at the 5% level of significance. Assuming that the distribution of annual S&P500 returns is approximately lognormal, allows us to use the Black-Scholes option pricing model to calculate the implied volatility for the index from the price of the closest to the money call option at the beginning of each holding period. This implied volatility is shown in Table I for all holding periods, and serves as our estimate for future volatility of the S&P500 index for each year.

Given that annual S&P 500 returns are approximately lognormal, we simulate the distribution for the log return as a normal distribution for each one-year holding period. The Black-Scholes implied volatility shown in Table I gives us a market forecast for the standard deviation of return for the holding period. Since there is no universally accepted forecast for the mean of the distribution, we use 10%, which is approximately the mean log return for the S&P500 index during the postwar period. In addition, to avoid the problem of isolated randomly generated points in the tails of the simulated return distribution for a holding period, we use deterministic simulation instead of Monte Carlo simulation. .

To generate the call option returns that correspond to the simulated S&P500 index returns for a given holding period, we use the market value for the index at the beginning of the holding period, and the 10,000 simulated index returns to generate 10,000 index values for the end of the holding period. Next we calculate the expiration values based on these index values and the strike price, and then use these expiration values and the call price at the beginning of the holding period to calculate the 10,000 call option returns that correspond to the 10,000 S&P500 returns. For the call price at the beginning of the holding period, we use the average of the closing bid and ask prices for the call option[5].

We combine the return distributions for the treasury, the S&P500 index and the call option on the index, to create a joint distribution for the three assets with 10,000 observations for each holding period. We use this joint return distribution for portfolio optimization with Power-Log utility functions, and its first two moments for mean-variance optimization. Although the optimization algorithms we have used for this study can create portfolios with long or short positions in any of the assets, we put in a "no short sales" constraint on the S&P500 index and the call option in order to make it easier to interpret the results.

OPTIMAL PORTFOLIOS

For downside powers ranging from 0 to -50, we construct optimal portfolios using Power-Log utility functions, and then construct corresponding mean-variance efficient portfolios with identical expected returns. Table II shows optimal Power-Log portfolios for downside power 0, and the matching mean-variance efficient portfolios. The downside power of 0 corresponds to the log utility function, which in theory produces the maximum growth portfolio, the riskiest in our set of portfolios. The expected return for each holding period is calculated for the optimal Power-Log portfolio from the 10,000 point joint return distribution for the three assets for that period, and then we use that expected return to construct the matching mean-variance (M-V) efficient portfolio.

Table II

		------ Optimal Power-Log ------ Portfolios with Downside Power 0				----- Mean-Variance Efficient ----- Portfolios			
Holding Period	Expected Return (%)	Riskless Weight (%)	S&P500 Weight (%)	Call Option Weight (%)	Realized Return (%)	Riskless Weight (%)	S&P500 Weight (%)	Call Option Weight (%)	Realized Return (%)
1	29.18	67.37	0.00	32.63	69.52	67.37	0.00	32.63	69.52
2	21.26	74.16	0.00	25.84	39.00	-32.18	122.61	9.57	44.48
3	20.80	36.18	45.79	18.03	23.41	-67.32	167.32	0.00	32.07
4	18.60	41.33	41.39	17.28	-17.74	-55.98	155.98	0.00	-13.19
5	18.36	62.72	16.00	21.28	-19.28	-33.47	129.53	3.94	-20.89
6	34.17	-0.55	75.88	24.67	-40.15	-180.27	280.27	0.00	-60.22
7	32.89	15.08	60.77	24.15	43.05	-143.41	243.41	0.00	55.03
8	69.60	-24.29	83.99	40.30	33.32	-524.86	624.86	0.00	66.00
9	59.81	30.34	25.30	44.35	0.09	-502.45	602.45	0.00	35.60
10	44.51	57.18	0.00	42.82	36.88	-95.81	166.79	29.02	43.68
11	41.41	57.33	0.00	42.67	-15.27	57.33	0.00	42.67	-15.27
12	30.10	56.72	12.51	30.77	-33.87	-156.91	256.91	0.00	-100.00
13	29.98	40.30	39.03	20.67	24.73	-51.98	151.98	0.00	41.62
Geo. Avg.	33.87				5.91				-100.00

Every optimal Power-Log portfolio in Table II has a large investment in the call option in every holding period. In contrast, the M-V efficient portfolios had an investment in the call option for only five out of thirteen holding periods, since the positive skewness of the call option's return is not valued in a mean-variance optimization. As expected, the realized returns for the optimal Power-Log and M-V efficient portfolios turned out to be very different from the expected returns. Every expected return is positive, but several realized returns are negative. The worst realized return for the optimal Power-Log portfolios was -40.15% in holding period 6, the year 2002, while the worst return for mean-variance efficient portfolios was a fatal -100% in period 12, the year 2008, that destroyed the mean-variance efficient portfolio. The second worst realized return for the optimal Power-Log portfolios was -33.87% in holding period 12, the year 2008, while the second worst realized return for mean-variance efficient portfolios was -60.22 % in period 6, the year 2002. The worst return for optimal Power-Log portfolios is better than both the worst and second worst returns for the M-V efficient portfolios. The market suffered large unanticipated losses 2002 and 2008, and both the optimal Power-Log and M-V efficient portfolios performed poorly, but the optimal Power-Log portfolios provided far better downside protection against these large unanticipated market declines. Clearly, Power-Log optimization does a much better job of controlling tail risk than mean-variance optimization for these risky portfolios, and also does a better job protecting against large unanticipated market declines.

While some investors, most notably Paul Samuelson (1971), might be willing to accept the risk associated with portfolios constructed with the log utility function, the risk is unacceptable for the vast majority of investors. As shown in Table II the smallest investment in the call option was 17.28% for the optimal Power-Log portfolio in holding period 4, and the largest investment in the call option was 44.35% in holding period 9. These are very large investments in a derivative, and carry a lot of risk. We construct less risky portfolios by reducing the downside power, making it more negative, which increases the penalty for losses and the resulting optimal Power-Log portfolios have less tail risk. Table III summarizes the realized returns for optimal Power-Log portfolios for five downside powers, 0 through -50, and the matching M-V efficient portfolios. As the downside power decreases, the optimal Power-Log portfolios become more conservative. This is reflected in the worst and second worst returns for optimal Power-Log portfolios in Table IV. The worst return increases from -40.15% for the downside power of 0, to -2.15% for a downside power of -50, and the second worst return increases from -33.87% to -1.94%. Tail risk drops dramatically for optimal Power-Log portfolios as the downside power decreases from 0 to -50.

Table III

			Optimal Power-Log Portfolio Returns (%)				M-V Efficient Portfolio Returns (%)			
No.	Downside Power	Expected Return (%)	Worst	Second Worst	Geo. Avg.	Best	Worst	Second Worst	Geo. Avg.	Best
1	0.00	33.87	-40.15	-33.87	5.91	69.52	-100.00	-60.22	-100.00	69.52
2	-0.60	28.72	-31.21	-24.71	6.28	58.85	-85.25	-49.41	-2.35	58.85
4	-3.00	19.68	-17.45	-14.78	5.70	40.61	-53.61	-30.40	4.32	40.61
4	-9.00	13.18	-8.75	-7.74	4.93	27.96	-31.21	-16.73	4.93	27.96
5	-50.00	7.70	-2.15	-1.94	4.04	17.76	-12.70	-5.24	4.11	17.76

Comparing the optimal Power-Log and M-V efficient portfolios in Table III, we see that the worst realized return for the optimal Power-Log portfolios is better than both the worst and second worst realized returns for the matching M-V efficient portfolios in every row. The optimal Power-Log portfolios have lower tail risk across the board. On the upside, both sets of portfolios had their best returns in period 1, the year 1997, for each downside power, and the compositions were identical for each matched pair. While Kale (2006) found significantly higher returns for optimal Power-Log portfolios than M-V efficient portfolios in the upper tail with simulated returns, our limited data did not give us enough discrimination to reproduce that result empirically.

The geometric average realized return shown in Table III is positive for the optimal Power-Log portfolios for each downside power. Interestingly, initially the geometric average returns increased for the optimal Power-Log portfolios as downside power decreased, rising from 5.91% for a downside power of 0, to 6.28% for a downside power of -0.6, and then declined steadily to 4.04% for a downside power of -50. According to theory, the downside power of 0, which corresponds to the log utility function, should produce the maximum growth portfolio, i.e., the portfolio with the highest geometric average return. That would have been true if the realized distribution for joint returns is the same as the distribution used to construct the optimal portfolios, but that is not the case here since there are always unanticipated changes in the economic environment. Large unanticipated negative returns are likely to have lowered the geometric average return for downside power 0 to less than that for downside power -0.6, which provides better protection against large losses.

Comparing the geometric average return for optimal Power-Log and M-V efficient portfolios in Table III, we see that for medium and high risk portfolios the return is far better for optimal Power-Log portfolios, but about the same for the most conservative portfolios. It turns out that for downside powers of -0.8 and higher, which correspond to expected returns of 27.46% and higher, the returns for M-V efficient portfolios are negative and give the appearance of falling off a cliff, ending up at a -100% as the downside power gets close to 0.

CONCLUSION

Kale (2006) introduced Power-Log utility functions to represent investor preferences that combine multiperiod portfolio theory based on log and power utility functions, and the tenets of behavioral finance for constructing portfolios that range from high-risk maximum growth portfolios to conservative portfolios with very low tail risk. Here we have extended Kale's simulation results with empirical data from 1997 through 2009 and annual rebalancing for portfolios to show that the realized returns for optimal Power-Log portfolios have lower tail risk than mean-variance efficient portfolios, and also have higher geometric average returns by and large. The geometric average return was positive for all optimal Power-log portfolios, while it was negative for the riskier M-V efficient portfolios and -100% for the riskiest one. We also found that the optimal Power-Log portfolios provided much better downside protection against large negative return surprises than the corresponding M-V efficient portfolios in 2002 and 2008, when the market was down substantially.

REFERENCES

1. For a list of references please contact the authors.