# Efficient Portfolio Management — Designing Optimal Strategies with Evolutionary Algorithms

**L. Grajeda**[1]

**Abstract**

This work combines Markowitz–Sharpe and Black–Scholes models in such a way as to take advantage of the best of both worlds, by introducing a multi-asset continuous time framework for Monte Carlo simulations. Such a framework takes into account any costs, including taxation, and allows for shocks in both return and volatility, thus becoming a convenient background for the search of quasi-optimal investment strategies via genetic algorithms. This allows for an efficient portfolio management, which keeps the overall risk low due to a Markowitz diversification, while taking advantage of the market fluctuations perceived by the Black–Scholes model. Numerical results, based on real data from 2005 to 2009 for a "blue chips" portfolio in the BOVESPA stock exchange, show a definite improvement over a CAPM "buy-and-hold" pure strategy.

## Introduction

Modern portfolio theory provides optimal asset diversification from the risk–return viewpoint. The cornerstone of this theory was laid down by Harry Markowitz (Markowitz, 1952) in a seminal work that was soon followed, expanded and complemented by many other authors. A major contribution was made by William Sharpe (Sharpe, 1964), whose Capital Asset Pricing Model (CAPM) is now a market standard for massive investors on the long run. A multifactor generalization of CAPM, the Arbitrage Pricing Theory (APT), was later introduced by Stephen Ross (Ross, 1976).

[1] Professor do Instituto de Ciências Exatas da Universidade Federal de Minas Gerais, Belo Horizonte, Brasil.
Professor do Ibmec — Minas Gerais, Belo Horizonte, Brasil.
Pesquisador do Instituto de Pesquisa em Finanças & Investimentos, Belo Horizonte, Brasil.

However, all classical models of portfolio theory consider only one single time period and rely implicitly on a “buy and hold” approach. They provide no help for asset managers over time, as no rebalancing strategy is produced. Indeed, it was quickly pointed out that Markowitz model is not suitable for repeated investments (Kelly Jr., 1956). Moreover, for any given risky asset, it is almost impossible to have reliable return expectations and its volatility changes quite often. Even worse, the one fund that would lead to optimal investment is computationally unstable, so that the smallest discrepancy in the risk-return estimates of any of the assets in the portfolio results in a completely different optimal portfolio. Therefore, despite its philosophical soundness and mathematical elegance, Markowitz model is not useful at all for most traders, as it exposes them to estimation risks that cannot be avoided and are hard to evaluate properly. Nevertheless, the model shed a new light to financial theory, as it points out the importance of risk management in investment decisions.

As for the CAPM, many restrictions do apply. To begin with, the hypotheses of the model are really hard to believe. Even worse, the model does not allow for any fundamentalist analysis, as it simply prescribes the purchase of a little bit of every available asset — this does not seem very wise, as one puts the good, the bad and the ugly companies all in the same basket. Moreover, the beta of any risky asset cannot be consistently used by traders for investment decisions, since its behavior is not statistically reliable. Most curiously, CAPM turns out to make sense for huge investors, for whom the model hypotheses sound pretty fair and who cannot easily get too far away from the market portfolio, which does not depend on the betas anyway. Also, empirical data shows that, in practice, CAPM is not at all easy to outperform. Together with its computational simplicity and the deep insight on the risk structure of interest rates, this is one of the reasons why CAPM is by far the most important model of modern finance.

The APT model generalizes CAPM to a multi-index framework. Though the idea sounds great, the model is not popular at all, as it requires a much larger econometric and computational effort, while resulting in little (if any) improvement over CAPM.

On the other hand, the Black–Scholes model, based on arbitrage theory and stochastic calculus, makes an extremely rich description of an asset price over time as a geometric Brownian motion, though it is not particularly useful for diversification. Further generalization, using Lévy processes, allows for the handling of non-lognormal assets. The model is the cornerstone of continuous time finance theory, which also received major contributions by Merton. However, the model applies only to one single asset and derivatives over it.

Anyway, both approaches are extremely sound, not only philosophically but also in practice, as they also give traders much qualitative insight, defining the parameters on which they operate. Moreover, they set the modern standards on which investment is conceived, made and analyzed, not only theoretically but also in terms of regulation, thus becoming landmarks of modern finance. With much merit, Markowitz and Sharpe were awarded the prestigious Nobel Prize in Economics in 1990, while Scholes and Merton received the same distinction in 1997.

Many attempts to combine these models were made along the past decades, however with little progress. So far, the mathematical complexity of the probabilistic structure of multidimensional Brownian motion has not allowed for any relevant theoretical or numerical improvements. A remarkable initiative was the Log-Optimal model, developed by David Luenberger in the early 1990's (Luenberger, 1998), which imbued Markowitz model with some long-term wisdom from Black-Scholes analysis. Unfortunately, Luenberger's model turns out to be equivalent to Markowitz model for most traders that know better, giving them no additional insight, nor any help for rebalancing portfolios on the long run.

As for the use of evolutionary algorithms in financial optimization, it was pioneered by Richard Bauer in the early 1990's (Bauer Jr., 1994). His work, however, focuses on the determination of optimal parameters for a prescribed trading strategy — the strategy itself remains the same all along — and does not consider dynamics of asset prices, taking into account only historical data.

One last advance in the use of genetic algorithms in financial optimization that must be mentioned is the work of Yang (Yang, 2006), which, in addition to historical information, incorporates future uncertainty by using binomial trees to enhance the accuracy of returns estimations. The method is very expensive in computational terms, not allowing for real time applications. Again, the investment strategy remains the same the whole time, just with optimal parameters, such as in Bauer's work.

This work significantly improves Yang's results by replacing binomial trees by the appropriate multidimensional stochastic differential equations system, which allows for simulations that take into account historical data, fundamentalist analysis and future uncertainty, resulting in a combination of Markowitz and Black–Scholes models in such a way as to take advantage of the best of both worlds, by introducing a multi-asset continuous time framework for Monte Carlo simulations. This sets the perfect background for the search of quasi-optimal investment strategies via genetic algorithms, resulting in an efficient asset management, which keeps the overall risk low with a Markowitz diversification while taking advantage of market fluctuations perceived by the Black–Scholes model. Numerical results for a "blue chips" portfolio in the BOVESPA stock exchange show a definite improvement over "buy-and-hold".

## Multidimensional Brownian Motion

Black-Scholes analysis assumes that the price $S(t)$ of an asset follows an Ito process:

$$\frac{dS}{S} = \mu\, dt + \sigma\varepsilon\sqrt{dt}$$

Here $\mu$ is the expected return of the asset and $\sigma$ is its volatility, whereas $\varepsilon$ is a normally distributed random variable. All the analysis is based upon this stochastic differential equation, which is easily generalized by considering that $S$ follows a Lévy process instead, i.e., that $\varepsilon$ has any stable distribution.

Lévy processes have a considerably more elaborate mathematical structure, which makes them hard to treat from the algebraic viewpoint, though it makes little difference on any numerical treatment. Therefore, the stochastic differential equation is well fit for Monte Carlo simulation, as it becomes the receipt to generate random paths of the Lévy process:

$$S_{k+1} = S_k + \mu S_k\, \Delta t + \sigma S_k \varepsilon\sqrt{\Delta t}$$

Realizations of the stochastic process produced this way (FIGURE 1) can then be used to test trading strategies including both the asset and derivatives over it.

Now, if a portfolio of $n$ assets is considered, with prices $S_1(t)$, $S_2(t)$, $S_3(t)$, … and $S_n(t)$, then those prices follow a system of stochastic differential equations:

$$\begin{cases} \dfrac{dS_1}{S_1} = \mu_1\, dt + \sigma_1\varepsilon_1\sqrt{dt} \\ \dfrac{dS_2}{S_2} = \mu_2\, dt + \sigma_2\varepsilon_2\sqrt{dt} \\ \qquad\vdots \\ \dfrac{dS_n}{S_n} = \mu_n\, dt + \sigma_n\varepsilon_n\sqrt{dt} \end{cases}$$

The corresponding difference equations system is a receipt to generate multidimensional random realizations of the combined process:

$$\begin{cases} S_{1,k+1} = S_{1,k} + \mu_1 S_{1,k}\, \Delta t + \sigma_1 S_{1,k}\varepsilon_1\sqrt{\Delta t} \\ S_{2,k+1} = S_{2,k} + \mu_2 S_{2,k}\, \Delta t + \sigma_2 S_{2,k}\varepsilon_2\sqrt{\Delta t} \\ \qquad\vdots \\ S_{n,k+1} = S_{n,k} + \mu_n S_{n,k}\, \Delta t + \sigma_n S_{n,k}\varepsilon_n\sqrt{\Delta t} \end{cases}$$

The method devised in this paper relies upon such equations to generate Monte Carlo simulations. This solution is noticeably simple, so that the total computing time is fairly small, allowing for an overhead optimization process of trading strategies by genetic algorithms to happen even in real time.

## The Correlation Matrix

The major difficulty in handling such a system is that the random variables $\varepsilon_1$, $\varepsilon_2$, $\varepsilon_3$, … and $\varepsilon_n$ are not independent. Thus the prices realizations cannot be generated independently. (Note that even if the random variables $\varepsilon_1$, $\varepsilon_2$, $\varepsilon_3$, … and $\varepsilon_n$ are all normally distributed, the combined distribution is not a multidimensional normal distribution.)

Indeed, the asset prices are almost always correlated, as shown in Markowitz model. All in all, the correlation matrix is the key point of Markowitz model. It is important to point out that, while the covariance matrix is computationally unstable, the correlation matrix is not. The reason for this is that the covariance reflects the instability on the assets' volatilities, while the correlation correlation expresses structural connections among companies and assets (the sectors on which companies operate, the commodities their business rely on, the currencies they trade with, the interest rates they are subject to, and so on). Thus, though the correlation matrix also changes over time, those changes are slow, small, often predictable and their source is usually easily identified in the real world as facts that affects the companies involved (accidents or frauds in a given company, deals among companies, etc.), the sectors on which they operate (new technologies, new regulations, etc.), the countries they trade in (interest rates, inflation, recession, etc.) and the international scenario (global crisis, wars, etc.).

Therefore, in any reliable simulation of the prices $S_1(t)$, $S_2(t)$, $S_3(t)$, … and $S_n(t)$ the individual stochastic processes of the assets of the portfolio must preserve the correlation matrix:

$$Q = \begin{pmatrix} \rho_{1,1} & \rho_{1,2} & \cdots & \rho_{1,n} \\ \rho_{2,1} & \rho_{2,2} & \cdots & \rho_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{n,1} & \rho_{n,2} & \cdots & \rho_{n,n} \end{pmatrix}$$

This work gives an effective solution for the problem of generating normally distributed random variables $\varepsilon_1$, $\varepsilon_2$, $\varepsilon_3$, … and $\varepsilon_n$ that will lead to the properly correlated stochastic processes. This is achieved by generating auxiliary normally distributed random variables $\eta_1$, $\eta_2$, $\eta_3$, … and $\eta_n$ and applying an appropriate linear transformation based upon the correlation matrix obtained from Markowitz model.

For instance, according to historical data obtained from Economática, ordinary stocks of Petrobras (PETR3) traded in the BOVESPA and spot gold contracts (OZ1DIS) traded in the BM&F have the following correlation matrix:

$$Q = \begin{pmatrix} 1 & -0{,}1867 \\ -0{,}1867 & 1 \end{pmatrix}$$

Let $\eta_1$ and $\eta_2$ be two normally distributed independent random variables the correlation matrix of the corresponding Ito process will just be the identity matrix, instead of $Q$. Now, take:

$$\begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \end{pmatrix} = \begin{pmatrix} -0{,}09376307 & 0{,}99559454 \\ 0{,}99559454 & -0{,}09376307 \end{pmatrix} \begin{pmatrix} \eta_1 \\ \eta_2 \end{pmatrix}$$

It is easy to see that both $\varepsilon_1$ and $\varepsilon_2$ are normally distributed random variables that have $Q$ as their correlation matrix. These random variables generate realizations of the Ito processes for both assets that are perfectly fit for the desired Monte Carlo simulations.

The situation in $n$-dimensional Brownian motion is similar that of the example above, a result that can be summarized in the following theorem (the proof is omitted for sake of brevity):

THEOREM: Given any correlation matrix $Q$, there is a linear transformation $L$ such that, if $\eta_1, \eta_2, \eta_3, \ldots$ and $\eta_n$ are normally distributed random variables, then

$$\begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \\ \vdots \\ \varepsilon_n \end{pmatrix} = L \begin{pmatrix} \eta_1 \\ \eta_2 \\ \vdots \\ \eta_n \end{pmatrix}$$

are normally distributed random variables that have $Q$ as their correlation matrix.

It is important to remark that the transformation $L$ is not unique and that different choices may be preferable in different situations. In the example above, $L$ is symmetric, but in general it can always be taken as a lower triangular, saving almost half the time consumed in the matrix multiplication.

In any case, this solution is noticeably simple to compute (once $L$ is evaluated, the additional computational cost is just one matrix multiplication for each step), so that the total computing time of the Monte Carlo simulations is fairly small, allowing for an overhead optimization process of trading strategies by genetic algorithms to happen even in real time.

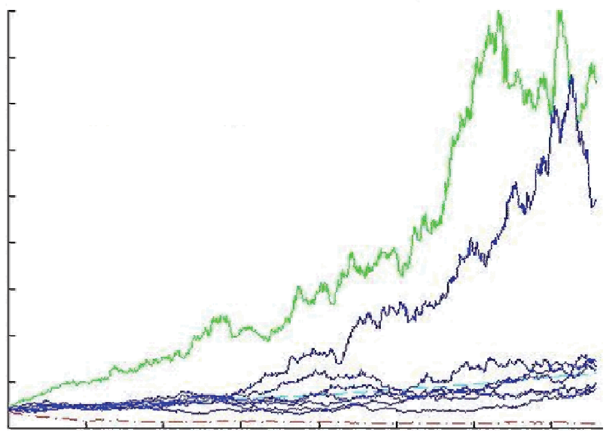

FIGURE 1 — Monte Carlo simulations for PETR3.

## Risk and Return Estimates

Once the random variables $\varepsilon_1$, $\varepsilon_2$, $\varepsilon_3$, … and $\varepsilon_n$ are properly generated, it is necessary to set the appropriate returns and volatilities. One of the major advantages of the method devise here is that such a framework allows for shocks in both return and volatility. This is done by letting both the $\mu$'s and the $\sigma$'s vary over time, sometimes rather abruptly.

Unlike other models where the $\mu$'s and the $\sigma$'s are set constant and equal to their historical averages, this model considers historical market data only to estimate the returns' and volatilities' range and elasticity, as well as to set the initial values of the $\sigma$'s as the current market volatilities.

On the other hand, the initial $\mu$'s must be set using both historical data and parameters from fundamentalist analysis, such as the Tobin's $q$'s. As noted by Dixit and Pindyck (Dixit, 1994), the principle underlying Tobin's $q$ is the orthodox net present value theory and, as Romer analysis clearly shows (Romer, 2001), "$q$ summarizes all information about the future that is relevant to a firm's investment decision". Other fundamentalist indicators are also taken into account, not only those concerning each company, but also those related to the sectors and countries in question.

After the initial setting of risk and return parameters, $\mu$'s and the $\sigma$'s will be altered throughout the simulation according to specific rules that consider both the fundamentalist indicators available and random fluctuations, including random shocks.

The shocks are all treated from a CAPM point of view and are considered to be of one of two kinds: specific (diversifiable), affecting only one asset, or systematic (nondiversifiable), affecting a whole sector, country or the world. Empirical results show that the all-time average $\beta$'s are reliable, easy-to-use parameters to spread systematic risk shocks throughout simulations.

It is important to note that fundamentalist analysis is implicitly implicated from the start in the model devised here, as it is supposed to be used in the initial selection of assets to be considered for investment. Indeed, the approach assumed here is plain orthodox and has conservative investors in mind.

The method's capacity to imbue highly complex numerical procedures with fundamentalist, micro- and macroeconomic analysis is certainly one of its most distinctive features. Nevertheless, the method works just fine if all that information is simply neglected, though the strategies produced will lack the ability to profit from a few high quality analyses (on the other hand, they will not be jeopardized by a few bad analyses either).

## Utility Function for Investment Strategies

Of course, the idea of testing strategies over past real world data is irresistible, but any optimization method over those tests will eventually lead to a "crystal ball for the past" strategy that will probably prove no so good for the future. Indeed, optimization over one scenario only will always result in the best strategy for that one scenario only and, unless the future repeats the past perfectly, this strategy is likely to be no better than any random strategy.

Therefore, it is highly desirable to set a simulation framework that allows for the generation of thousands of possible, reasonable scenarios that can be used to test strategies. And it is exactly this that was achieved above.

Now that a framework has been established for testing investment strategies, it is time to define what makes a strategy better than other. In other words, it is time to define a utility function that will be used for rating strategies. Many choices are possible, according to each investor risk profile. In this work, the long-term average return over all simulated scenarios was taken as the utility function:

$$U(\xi) = \bar{\mu}(\xi) - \frac{\bar{\sigma}^2(\xi)}{2}$$

where $\xi$ is any given trading strategy.

This utility function is so because, on the long run, risk affects return (negatively). The corrected long-term return is the geometric average of the individual one-period returns, corresponding to the lognormal arithmetic average of the return. This is very clear from stochastic calculus, as noted by Luenberger in his Log-Optimal model(Luenberger, 1998).

It is important to point out that the method devised here can easily accommodate a multi-objective optimization so that several different utility functions can be used simultaneously. Although maximum long-term return is sure to be an objective for all wise investors, it is likely that asset managers will also set out short-term goals, maximum risk tolerances and other objectives.

It is also noteworthy that some conditions can be embedded in the testing algorithm itself, such as taxation, transaction costs and trading restrictions. For instance, most investment clubs at the BOVESPA are required to have at least half their money on stocks at any given time, so that only those strategies may be considered from the start.

## The Genetic Algorithm

A genetic algorithm is an optimization technique that emulates a natural evolutionary process that, over the generation, selects the individuals best adapted to their environment. The algorithm starts with an initial population, which is ranked by a selection criterion; then the best fit individuals pass their characteristics on to their offspring through reproduction and a new generation is formed; the process repeats itself over and over until a stop condition is reached. There are no *a priori* limits on how many factors to take into account, on how large the populations should be or on how many generations to produce, as long as one is willing to take the programming effort to encode all the details and wait for the algorithm to run.

In this work, individuals are trading strategies. Those strategies have definite characteristics such as preset conditions that precisely establish when to trade (starts, stops, etc.), what to trade (asset selection), what to do (buy, sell, short, long, etc.), how much to trade, etc. The steps are as follows:

| | |
|---|---|
| STEP 1: | INITIAL POPULATION<br>An initial population of 200 strategies, that includes traditional trading strategies (such as "buy-and-hold"), strategies from selected traders and randomly generated strategies. |

| | |
|---|---|
| STEP 2: | SELECTION CRITERIA<br>The section criterion is the utility function mentioned in the previous section, using the Monte Carlo simulation framework devised above. The best adapted individuals are those strategies that score the highest values of the utility function. |

| | |
|---|---|
| STEP 3: | REPRODUCTION<br>Reproduction is made both by crossover (sex) and cloning, with a few random mutations. A new generation is formed as follows: the top ranking 20 strategies were simply cloned; the other 180 strategies were obtained by combining the whole population in pairs, and mixing their characteristics. (The probability of coupling was taken proportionally to the strategy ranking.) Finally, random mutations were introduced in 10 randomly selected individuals. |

| | |
|---|---|
| STEP 4: | ITERATION<br>The new generation is then taken back to STEP 1 and the process is repeated, up to 30 generations or until no improvement whatsoever is made for at least 3 generations. |

## Results

The genetic algorithm runs with great efficiency, even on small personal computers. Simulations started with the generation of 10,000 scenarios and then run the genetic algorithm for up to 30 generations. The whole process usually took no longer than 2 minutes, as most of the populations converged not much after 20 generations.

One interesting result to mention is that "buy-and-hold" almost always ranks within the top 10 trading strategies throughout the generations when only one asset is considered. However, when two or more assets are used, "buy-and-hold" hardly ever stood up till the end of the process. Instead, a Markowitz portfolio with a little pumping, so that the weights of all assets remain constant over time, appeared quite often, though the best strategies outperformed it up to 30% a year.

In a few cases, such as the one shown bellow (FIGURE 2) an optimized strategy led to a portfolio that even surpassed the return of the best performing stock in the period. It is important to mention that the strategy was obtained using only data and analyses from the previous time period.

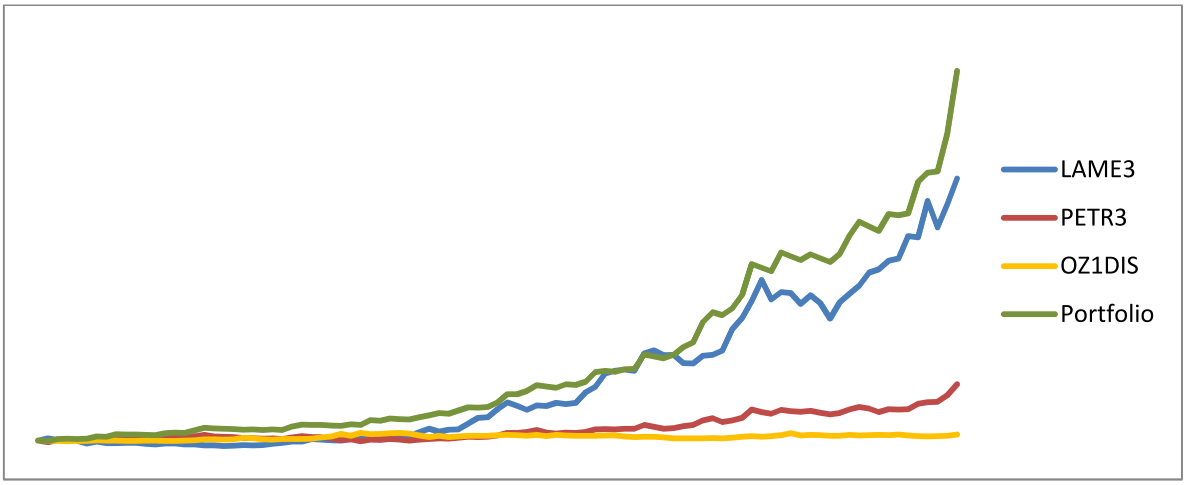


FIGURE 2 — The performance of a quasi-optimal strategy for 3 assets.

Finally, some of the strategies obtained by using this method for one asset alone are already being effectively tested in the BOVESPA. For instance, since October 29, 2009, a strategy for trading PETR3 is in use by Clube de Investimento Savassi, an investment club whose portfolio is managed by the author. Ever since, the club has earned a net profit of 1,67% on that stock, despite of a price drop of 19,47% for that asset. Preliminary analyses show that the actual results really confirm the computer simulations.

## Conclusion

The method devised here significantly improves previous results on portfolio management. The simulations base upon multidimensional stochastic differential equations system works well, allowing for a promising combination of the mainstream models of modern finance and economic theory.

Numerical results are currently being confirmed by actual market operations, achieving a definite improvement over “buy-and-hold” for small and mid-sized portfolios. Such success is not expected to apply to huge investors, for which CAPM is still the way to go.

The best feature of the method is that it allows investors to trade the way they like, as they can simply input their current trading strategy and have it enhanced by setting optimal operational parameters and perhaps also profiting from sound analyses.

Further enhanced applications of this method are expected for the near future as new parameters are yet to be programmed and scientific cooperation with world-class economists can still improve the model a lot.

Theoretical advances towards the combination of Markowitz and Black–Scholes models are also expected for the future.

## References


- Bauer Jr., R. J.: *Genetic Algorithms and Investment Strategies*. John Wiley & Sons, Inc., 1994.
- Dixit, A. K., & Pindyck, R. S.: *Investment under Uncertainty*. Princeton University Press, 1994.
- Kelly Jr., J. L.: *A New Interpretation of Information Rate*. Bell System Technical Journal , 35, pp. 917-926, 1956.
- Luenberger, D. G.: *Investment Science*. Oxford University Press, 1998.
- Markowitz, H. M.: *Portfolio Selection*. Journal of Finance , 7 No. 1, pp. 77-91, 1952.
- Romer, D.: *Advanced Macroeconomics*. 2nd ed., McGraw-Hill, 2001.
- Ross, S. A.: *The Arbitrage Theory of Capital Asset Pricing*. Journal of Economic Theory , 13, pp. 341-360, 1976.
- Sharpe, W. F.: *Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk*. Journal of Finance , 19, pp. 425-442, 1964.
- Yang, X.: *Improving Portfolio Efficiency: A Genetic Algorithm Approach*. Computational Economics , 28, pp. 1-14, 2006.