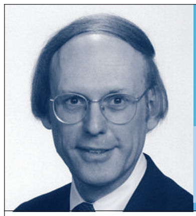

Daniel Fylstra

Introducing Convex and Conic Optimization for the Quantitative Finance Professional

Few people are aware of a quiet revolution that has taken place in optimization methods over the last decade

Optimization has played an important role in quantitative finance ever since Markowitz published his original paper on portfolio selection in 1952[1]. Most "quants" have some knowledge of linear programming as used in bond duration matching, quadratic programming as used in equity portfolio optimization, and nonlinear optimization used with portfolios of derivatives. The finance industry and technical literature is home to many different, competing approaches that apply these standard optimization methods to different models and risk measures, in an effort to achieve better investment results.

But fewer people are aware of a quiet revolution that has taken place in the optimization methods themselves over the last decade. A better understanding of the properties of optimization models, and new algorithms – notably interior point or barrier methods – have led to a changed view of the whole field of optimization. It is not an exaggeration to say that linear and quadratic programming are being replaced – or more properly subsumed – by more powerful and general methods of convex and conic optimization.

What this means for quantitative finance is a relaxing of restrictions on optimization models – for example, using quadratic constraints as easily as a quadratic objective – and new ways to deal with important problems such as "unnatural" portfolios from optimization that are due to "noise" in the return and covariance or factor parameters of a portfolio model. This article will introduce the ideas of convex and conic optimization, and their applications in quantitative finance.

It is not an exaggeration to say that linear and quadratic programming are being replaced – or more properly subsumed – by more powerful and general methods of convex and conic optimization

Linear and quadratic programming

An optimization problem consists of *decision variables*, an *objective function* to be maximized or minimized, and *constraints* that place limits on other functions of the variables. In *linear programming*, the objective and constraint functions are all linear – hence the simple form:

$$\begin{array}{ll} \text{max/min} & cx \\ \text{subject to} & b_l \leq Ax \leq b_u \\ & x_l \leq x \leq x_u \end{array}$$

where x is a vector of decision variables, cx is the objective function, A is a matrix of coefficients and Ax computes the constraints, b_l and x_l are lower bounds, and b_u and x_u are upper bounds. In *quadratic programming*, the objective is generalized to a quadratic function of the form:

$$\begin{array}{ll} \text{max/min} & x^TQx + cx \\ \text{subject to} & b_l \leq Ax \leq b_u \\ & x_l \leq x \leq x_u \end{array}$$

where Q is a matrix of coefficients. In the simplest formulation of the classic Markowitz portfolio optimization problem, Q is a covariance matrix, the objective x^TQx is portfolio variance to be minimized, and A has just two rows: A *budget constraint* $1x = 1$ and a *portfolio return threshold* $ax \geq b$ where a is the expected return of each security and b is the minimum portfolio return. A factor model that expresses the "beta" or sensitivity of each security to one or more market factors also leads to a quadratic programming model.

Classical quadratic programming still requires that all constraints are linear; quadratic or more general constraints would put the problem in the domain of nonlinear optimization. Quantitative finance professionals have worked hard to create models that "fit" within the domain of linear and quadratic programming, and avoid models that require nonlinear optimization methods. Why?

Convex and non-convex problems

Nonlinear optimization is a well-developed field, with many solution algorithms for problems of the general form:

$$\begin{array}{ll} \text{max/min} & f(x) \\ \text{subject to} & b_l \leq G(x) \leq b_u \\ & x_l \leq x \leq x_u \end{array}$$

where $f(x)$ is a smooth function, and $G(x)$ is a vector of smooth functions of the variables x. Linear and quadratic programming problems are special cases of this form, where $f(x) = cx$ or $x^TQx + cx$, and $G(x) = Ax$. But solution algorithms for this general problem can find only "locally optimal" solutions that may not be "globally optimal," and they may fail to find a feasible solution even though one exists. Unlike linear and quadratic programming, the time taken to find a globally optimal solution can rise exponentially with the number of variables and constraints. This severely limits the size of problems that can be solved to global optimality. Hence the aversion to nonlinear optimization for practical quantitative finance problems seems well founded. Or is it?

In an observation now famous for its prescience among optimization researchers, mathematician R. Tyrrell Rockafellar wrote in *SIAM Review* in 1993[2]:

> ". . . in fact, the great watershed in optimization isn't between linearity and nonlinearity, but convexity and nonconvexity."

What makes the general optimization problem so hard to solve? It is *not* the fact that $f(x)$ or $G(x)$ may be nonlinear! It is the fact that $f(x)$ or $G(x)$ may be *non-convex*.

A **convex optimization** problem is one where all of the constraints are convex functions, and the objective is a convex function if minimizing, or a concave function if maximizing. A **non-convex optimization** problem is any case where the objective or any of the constraints are non-convex functions. The difference is dramatic: Convex optimization problems can be efficiently solved to global optimality with up to tens or even hundreds of thousands of variables. In contrast, the best methods for global optimization on modern computers usually can solve non-convex problems of only a few hundred variables to global optimality.

Even a quadratic programming problem may be "impossibly" hard to solve (in mathematical terms, *NP-hard*) if the objective function x^TQx is non-convex, which happens when the matrix Q is indefinite[3]. Fortunately for portfolio optimization, the matrix Q will be positive definite if the covariance terms are computed from historical data where the number of observations exceeds the number of securities.

Convex functions

What do we mean by a convex function? Geometrically, a function is **convex** if a line segment drawn from any point $(x, f(x))$ to another point $(y, f(y))$ – called the *chord* from x to y – lies *on or above* the graph of f, as in the picture below:

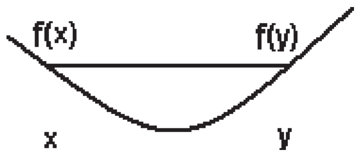

Algebraically, f is convex if, for any x and y, and any t between 0 and 1, $f(tx + (1-t)y) \leq tf(x) + (1-t)f(y)$. A function is **concave** if $-f$ is convex – i.e. if the chord from x to y lies *on or below* the graph of f. It is easy to see that every linear function – whose graph is a straight line – is both convex and concave. A quadratic function x^TQx is convex if Q is positive semi-definite, or concave if Q is negative semi-definite. A **non-convex** function "curves up and down" – it is neither convex nor concave. A familiar example is the sine function; perhaps less familiar is the quadratic x^TQx where Q is indefinite, for example $x_1^2 - 2x_1x_2 - 1/2(x_2^2 - 1)$ which is plotted below:

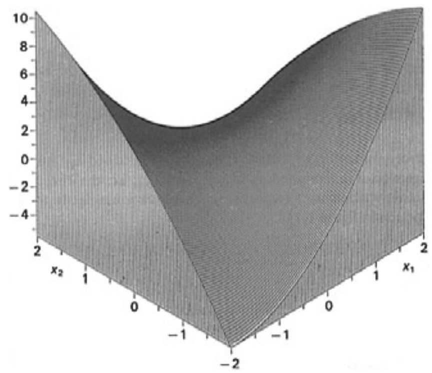

General nonlinear functions may be convex or non-convex. A problem with all convex nonlinear functions can be solved efficiently, to global optimality, to very large size; but these favorable features are lost if even one problem function is non-convex.

Quantitative finance professionals have worked hard to create models that 'fit' within the domain of linear and quadratic programming, avoiding models that require nonlinear optimization methods. Why?

Interior point methods and software

Rockafellar's observation moved from theoretically interesting to practically significant with the development of interior point methods – first for linear programming, pioneered by Khachian and Karmarkar[4] in the 1980s, then for more general convex optimization problems, pioneered by Nesterov and Nemirovskii[5] 1994. Where the simplex method relies on the fact that the constraints are linear, and moves from vertex to vertex where constraints intersect, interior point methods rely only on the fact that the constraints are convex – they move along the so-called central path defined by a nonlinear barrier function, even for LP problems.

This means, for example, that an interior point method doesn't "care" whether it is dealing with a quadratic objective or (one or many) quadratic constraints – it will take essentially the same number of steps to find the optimal solution.

Commercial interior point optimization software that handles quadratic constraints has just begun to appear; examples are the Barrier component of the CPLEX optimizer, the MOSEK optimizer, and the Barrier solver in Frontline Systems' Premium Solver Platform for Excel.

Efficient portfolios and quadratic constraints

As mentioned earlier, quantitative finance professionals have worked hard to make their models "fit" within the domain of linear and quadratic programming. For example, in Robert Fernholz's well-regarded work applying "stochastic portfolio theory" to equity portfolio management[6], the linear portfolio return threshold constraint is replaced by a portfolio growth threshold constraint that includes a quadratic term. Fernholz remarks that "this constraint is nonlinear, and conventional quadratic programs cannot be used in this case." He goes on to show that, for enhanced index ("tracking") portfolios, the quadratic term in the growth constraint is small in magnitude, and therefore can be neglected – yielding a linear constraint and a conventional QP problem. But with modern interior point methods, this simplification is no longer necessary – the original problem including the quadratic constraint can be solved just as reliably, and in about the same time, as the simplified QP problem.

Conic optimization

Nesterov and Nemirovskii's seminal work also opened up a natural generalization of linear programming, called **conic programming**, that shares the favorable characteristics of LP but encompasses a wider range of problems. And conic quadratic programming – now known as **second order cone programming** – shares the favorable characteristics of QP but encompasses a wider range problems relevant for portfolio optimization.

A conic optimization problem can be written as an LP – with a linear objective and linear constraints – plus one or more cone constraints. A **cone constraint** specifies that the vector formed by a set of decision variables is constrained to lie within a **closed convex pointed cone**. The simplest example of such a cone is the *non-negative orthant*, the region where all variables are non-negative – the normal situation in an LP. But conic optimization allows for more general cones, that can express all the elements of a portfolio optimization problem, and much more.

A simple type of closed convex pointed cone that captures many optimization problems of interest is the **second order cone**, also called the Lorentz cone or "ice cream cone." Geometrically it looks like the picture below, in three dimensions:

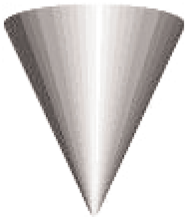

A **second order cone (SOC) constraint** of dimension n specifies that the vector formed by a set of n decision variables must belong to this cone. Algebraically, the L2-norm of $n - 1$ variables must be less than or equal to the magnitude of the *nth* variable. Any convex quadratic constraint can be reformulated as an SOC constraint, though this requires several steps of linear algebra. A quadratic objective $x^T Qx$ can be handled by introducing a new variable t, making the objective "minimize t", adding the constraint $x^T Qx \leq t$, and converting this constraint to SOC form.

A problem with a linear objective and linear plus SOC constraints is called a **second order cone programming (SOCP)** problem. Such a problem can be efficiently solved via specialized interior point methods, in about the same time as a QP problem of the same size. Enhancement of the simplex method to solve SOCP problems is an area of active research, but there are not yet commercial quality implementations of the proposed methods.

SOCP problems and portfolio optimization

One of the major problems with classical Markowitz portfolio optimization is the presence of "noise" or sampling errors in the input data used to estimate the terms of the covariance matrix, or the factor loadings computed for a factor model. Michaud drew attention to this issue in a 1998 book[7], referring to portfolio optimization as "an error prone procedure that often results in error-maximized and investment-irrelevant portfolios." This occurs because optimization, by its nature, maximally exploits the diversification potential of the data, including the estimation errors. Resampling of the efficient frontier, proposed by Michaud, and various scenario-based approaches have been developed to cope with this problem, but these methods do not explicitly model parameter uncertainty or provide

An interior point method doesn't 'care' whether it is dealing with a quadratic objective or (one of many) quadratic constraints – it will take essentially the same number of steps to find the optimal solution

>

any performance guarantees on the computed portfolio.

Goldfarb and Iyengar[8] have shown how conic optimization can be applied to the sampling error problem in portfolio optimization. They begin with a factor model, but they treat the mean return estimates and the factor loadings for each security as "noisy," belonging to parameterized uncertainty sets. They show how the parameters of the uncertainty sets can be obtained through the normal process of linear regression used to estimate historical returns and factor sensitivities, and how to obtain confidence intervals for the error in these estimated parameters. Then they solve variations of the portfolio optimization problem (minimizing variance, maximizing Sharpe ratio, minimizing VaR, etc.) by reformulating the constraints that include the uncertain parameters as second order cone (SOC) constraints, with a user-specified confidence threshold. The result is a portfolio that does offer probabilistic guarantees on risk-return performance, and that requires about the same amount of computing time as conventional portfolio optimization.

Software for convex and conic problems

Conic optimization, including second order cone programming and a related methodology called semidefinite programming (SDP), has been a very active area for optimization researchers in recent years. Some codes developed by academics, for example Boyd and Vandenberghe's SOCP code and Sturm's SeDuMi code, can be found on the World Wide Web.

For quantitative finance professionals who are comfortable working in Excel, the simplest way to explore conic optimization is to download a free trial version of Frontline Systems' Premium Solver Platform V6.0 from www.solver.com. This software is upward compatible from the standard Excel Solver (that Frontline originally developed for Microsoft) and now includes a built-in SOCP Barrier Solver. It automatically handles the linear algebra needed to transform quadratic constraints and objectives into SOC constraints, sparing the user this effort. Cone constraints can be entered by simply selecting "soc" from a drop-down list.

The limitations of nonlinear optimization – heavily used in portfolios of derivatives – is not due to non-linearity, but rather is due to non-convexity

Hence, one can concentrate on the problem of interest rather than details of the software.

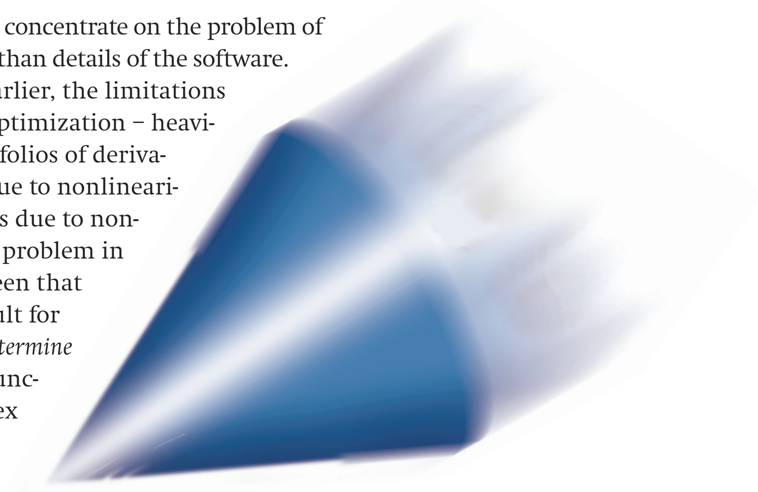

As noted earlier, the limitations of nonlinear optimization – heavily used in portfolios of derivatives – is not due to nonlinearity, but rather is due to non-convexity. The problem in practice has been that it's very difficult for modelers to *determine whether* their functions are convex or non-convex. It takes considerable mathematical expertise and time to prove convexity, and software to assist in this task has not been available. But a first-generation facility[9] to automatically determine whether a nonlinear model is convex or non-convex, by simply clicking a button, is included in the Premium Solver Platform V6.0. This makes it possible to explore convex optimization using any of the variety of nonlinear optimizers available for this platform.

Conclusions

There is growing excitement in the optimization research community about the new developments in convex and conic optimization – but this is just beginning to "spill over" into the quantitative finance community, where many users of optimization are found. Linear and quadratic programming – now over fifty years old – are being transformed into methods more powerful and general than ever before. Quantitative finance applications of the new technology of convex and conic optimization are sure to grow in importance over the next few years.

REFERENCES

1. H. Markowitz. Portfolio Selection. *Journal of Finance* 7 1: 77–91, 1952.
2. T. Rockafellar. Lagrange multipliers and optimality. *SIAM Review* 35:183–238, 1993.
3. R. Horst, P. Pardalos, N. Thoai. *Introduction to Global Optimization*. Kluwer, 1995.
4. N. Karmarkar. A new polynomial-time algorithm for linear programming. *Combinatorica* 4: 373–395, 1984.
5. Y. Nesterov and A. Nemirovskii. *Interior-Point Polynomial Algorithms in Convex Programming*. SIAM, 1994.
6. R. Fernholz. The Application of Stochastic Portfolio Theory to Equity Management. INTECH research report, 2003.
7. R. Michaud. *Efficient Asset Management*. Harvard Business School Press, 1998.
8. D. Goldfarb and G. Iyengar. Robust portfolio selection problems. CORC Technical Report TR-2002–03, Columbia University.
9. I. Nenov, D. Fylstra, L. Kolev. Convexity Determination in the Microsoft Excel Solver Using Automatic Differentiation Techniques. Submitted for publication.