Arbitrage Research and Trading

ARTicles: Optimisation and P&L – Part 1

Feb 11, 2000

ARTicles are a series of notes produced by the staff at *ART*. They are intended to share (as we believe) some interesting thoughts, and supplied solely for awareness and education. *ART* makes every effort to produce accurate and reliable materials, but can make no warranties whatsoever as to the content, or any results or interpretation of these materials.

ARTicles are opinions expressed by our staff and as such make no attempt to conform to practices observed in traditional publications.

42 Brook Street, London, UK, W2 2LZ Tel: +44 (0) 870 750-1410 Fax: +44 (0) 870 750-1411 http://www.arbitrage-trading.com

Optimisation and P&L: Part 1

Optimisation and P&L is at once one of the simplest and also one of the most difficult problems in trading. The notion of optimal P&L is instinctive since we all know where we would like to end up: rich and happy. However, it is the getting there (the implementation and execution), and importantly the reconciliation of differences between expected and realised P&Ls that are tricky.

To this end, we here at *ART,* have produced this series of notes to help demystify terminology and concepts, and most importantly to provide, in a pedestrian language, explanations of the issues and shortcomings as applied to real world risk/return. This note is the first part in the series, and its job is to lay the foundation of some of the concepts for both "theory" and "reality".

Part 2 in this series reviews the simplifying assumptions embedded in some market convention portfolio optimisation and introduces a *free* Efficient Frontier Optimiser (downloadable from our*ARTWeb* at http://www.arbitrage-trading.com sometime in May), that is an interactive tool to help illustrate the basic workings of "one type" of portfolio allocation technique. Part 3 of the series goes on to discuss methods that overcome the difficulties of the severe assumptions in many techniques and to lay the foundation for a relatively *foolproof*, albeit expensive, approach to real world portfolio risk/return optimisation.

Before we begin, we pay homage to our God of "do I buy or sell it". We feel that it is easy to demonstrate that if your investment horizon is long (5-10 years or more), and your risk profile favours a buy and hold strategy, then very little if any of this analysis machinery will be critical to you. On the other hand if you are seeking to outperform the market average by clever utilisation of market dynamics, then you (as us) will need all the help you can get.

In order to assist us in this journey, we first present two anecdotes that we feel are most appropriate to many aspects of quantitative trading and have been born of some experience (we are not certain who should get credit for coining these, and we would be grateful for your feedback):

> *For every complicated problem, there is always a simple solution, but usually its wrong.*

And

> *During a stroll on a particularly dark night, I came upon a man under the beam of a streetlight – frantically searching for car keys on the ground. I offered my assistance and to help expedite the search I asked:*
>
> *"Do you recall approximately where you dropped them?"*

He indicated that it was off in the distance somewhere in the dark, and motioned in the general direction. At which point I asked:

"Then why are you searching over here?"

Whence he promptly straightened me out and replied:

"Because this is where the light is"

As we shall see these two themes repeat themselves with astonishing frequency

1 Of Happiness and Lemmings

The ideal result is to be "happiest". The general notion of happiness (or Utility Theory) is far beyond the scope of this discussion, and we will restrict the current definition of "happiest" as that of maximal P&L (with financial risk adjustment to be introduced latter).

Even with this restricted objective the realities of the trading environment (e.g. job security and "plausible deniability" when things go badly – as they often do) bring into play complex human traits. A very common theme is that

"it is better to fail conventionally, then to succeed unconventionally"

Resulting with the predictable and well-documented market convention of "the herd mentality".

By our definition this must produce sub-optimal behaviour, and will figure into the discussion of P&L optimisation further below. However, at the end of the day what is truly optimal is subjective, and as such regardless of how we or anyone feels about optimality, "you" are the only one who could possible decide what is "best" for you. We hope, that at least you have the right tools and understanding to do so.

2 Forecasting as the trader's *rasion d'etre*

Whether we like it or not (or know it or not), all trading and investment activities involve forecasting. Unfortunately, this forecasting is performed under uncertainty, and thus quite tricky[1].

So there has been (and there is) a great deal of effort put into improving forecasting. The vast majority of such efforts arise via quantitative methods. So, first and foremost, it essential to understand that any such endeavour is made-up of at least two (and we think three) parts:

[1] It is noteworthy, that "market neutral" trading such a pure market making is also a forecasting problem as embedded in every traded price is a reflection of future expectations about risk and returns.

1) Modelling: this is the step wherein we formulate the properties of the process to be implemented. This is the step that requires keen insight into the "nature of things". This step is often quantitative, but importantly it is the step that captures the "physics" of the target problem. For example, the "binomial model" in options trading is not really so much of a model as one implementation of a "model" of drift and volatility in a manner that aims to mimic real world patterns. Similarly, in optimisation the modelling is the point where we decide what questions need to be answered, what are the properties of the problem, and what (if any) simplifying assumptions are reasonable. Present value theory is a model, with a relatively simple solution.

   Often, modelling can be a subjective process, since "you" must decide what the correct properties are for your circumstances (business mandate, share holder requirements etc ultimately dictate restrictions – e.g. it is pointless to model arbitrage behaviour that may require holding positions to maturity if management *et al* will not permit holding products to maturity, or insist on rebalancing limits that conflict with the arbitrage dynamics).

2) Solving the Model: this is generally a purely mechanical (albeit often complex) step. This is the step that takes the model and converts/implements it into a "machine" for generating the results of the model. It is not subjective. However, it may be (and is) intractable for many interesting real world problems; thus leading to simplifying assumptions. It is frequently the case that the simplifying assumptions are too unrealistic, and thus lead to arbitragible model/solution errors. Even worse some simplifying assumptions rear their (ugly) heads infrequently (for example some quant models may be OK'ish when there is liquidity, but ….).

   Implementing present value theory (e.g. *future value = (1+r) * present value*) is a solution/implementation of that model.

3) Audit: perhaps the most important step of all, and in our experience one all to rarely done. This comes in many forms including backtesting, frequent and careful checking of outcomes against expectations etc, with a view to confirm or update (or abdicate) as necessary. This is the anchor that holds the "control loop" together.

3 Optimisation: the "touchy-feely" version

At some stage in any quantitative analyses maths are required. Here we will use the "picture" equivalents to illustrate (some of) the maths. The bulk of what follows belongs in Step 2 of the process (Solving the model), but we will interject some "reality impact" addressing a few modelling and audit issues as well.

For simplicity consider optimisation problems in two categories: those that are amenable to mathematical manipulation, and those that are not. The first step will be to introduce some ideas for the mathematically treatable ones.

3.1.1 A Mean/Variance Portfolio Optimisation Example

Optimisation is actually a very broad set of problems. However, on lucky occasions we find problems that are also (easily) mathematically tractable. Another words, we dropped our car keys just underneath the streetlamp, and they are easily locatable. We stress that at this stage we are dealing with the "solution step", and that the "model" has already been decided on.

As an example consider the case of Mean/Variance Portfolio Optimisation[2] as popularised my Markowitz in 1952. This is only one form of optimisation, and many other tools exist to perform portfolio risk/return optimisation. This is one technique that assists in choosing the "optimal" composition of a portfolio. Though, we shall see further below that the mathematical meaning of "optimal" may not always coincide with our instinctive or real world definition.

Suppose that we are only interested in two (risky) investments[3] and wish to determine the "best" composition of them in a mean/variance sense. The return on the portfolio is given by:

$$R_{Port} = r_1 x_1 + r_2 x_2 \qquad (0.1)$$

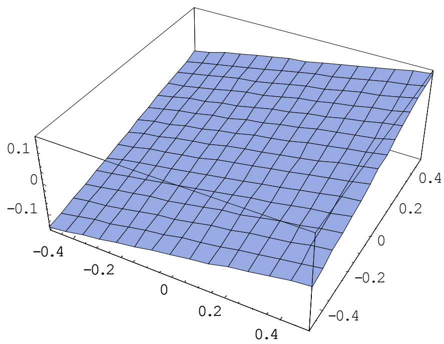

where the r_i's are the individual returns and the x_i's are the compositions. Figure 1 illustrates the possible portfolio returns for any given composition. By itself this graph is not sufficient to choose the "best" compositions on a risk-adjusted basis.

To make the connection to "risk", this approach <u>assumes</u> that risk may be represented in term of the portfolios variance (this is quite a controversial assumption, but one that simultaneously makes the problem "easy", and also one that pushes it away from the real world). The usual formulation of this is:

[2] There are many other "optimal" portfolio selection methodologies as well.
[3] Exactly the same principles apply for more securities (technically referred to as higher dimensions), but we don't know how draw those "hyper planes".

$$Var_{Port} = \sigma_1^2 x_1^2 + 2\rho_{1,2}\sigma_1\sigma_2 x_1 x_2 + \sigma_2^2 x_2^2 \qquad (0.2)$$

where the *σ's* are the standard deviation of the returns, and *ρ* is the correlation coefficient between the two returns process. Figure 2 illustrates the portfolio variance for both standard deviation being 0.3, and zero correlation. This is exactly a two dimensional parabola (a paraboloid).

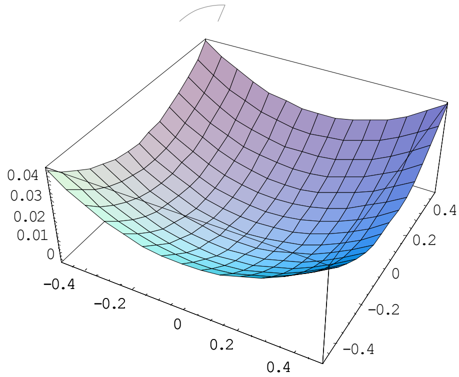

This is a "non-linear" optimisation problem since at least one of the surfaces is "curved". Specifically, this is referred to as a "quadratic programming" problem, since the nature of the curvature is exactly quadratic (parabolic).

The optimal solution to this grossly simplified problem might be thought of as the compositions that give rise to the maximum return for a "given" variance along the intersection of the two surfaces, or obtain the minimum variance solution for a "given" desired return (Figure 3). In most problems there may be many planes that intersect (representing conditions such as the level of variance (risk) that the investor may wish to take and so forth). Thus the picture of fully specified problem can be quite complicated.

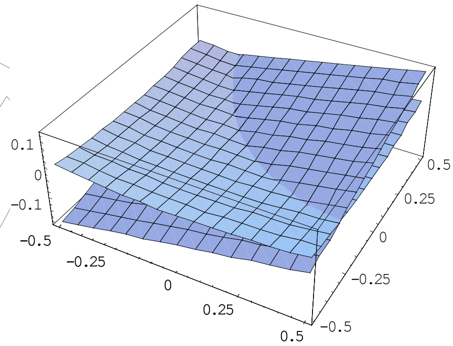

In order to assist us with an understanding of "mathematical optimisation", let us alter the example and consider the variance surface in isolation. The "bowl" shape that we see is a direct consequence of the "chosen model" (i.e. the variance equation (0.2)).

The minimum of this problem is easy to "see" qualitatively: its just the bottom of the bowl. Mathematical methods that "see" the bottom of the bowl vary considerable in implementation, but tend to employ effectively similar principles[4]. The mathematical equivalent to seeing the bottom of the bowl can be complicated by the result that the equations do not always "see" the entire bowl sufficiently clearly to decide the location of the bottom immediately. Rather, many of these methods are only capable of seeing one point on the surface at a time, and then employ some form of "search" and "am I at the bottom test" process to move around the bowl until the "test" returns the "yes".

If "you" could not see the entire bowl, one technique you might use is to place a ball bearing at any point that you can see, and let it go. Gravity would pull the ball bearing to the minimum. The mathematical idealisation of this type of approach is to analyse the

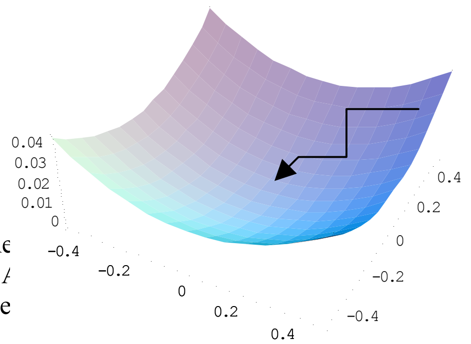

[4] In fact, this problem can be solved by very much more direct mathe rudimentary training in first derivatives and simple matrix algebra. A not actually be used in a practical setting. Rather, the example is use appropriate examples would dilute the thrust of this presentation.

"slope" of the point you are at. By analogy to "gravity" we would seek the direction in which the slope is steepest (the greatest gravitational pull) as the search direction for the next points. In the case of quadratic problems this is very easy and also has the special favourable circumstance that the direction in which the slope is steeper is generally the direction towards the minimum (this property may not hold in general problems).

So in this class of methods the mathematics does not actually search for the minimum directly, but rather searches for a specific "implied" property (shallowest slope), by repeatedly taking "steps" in the direction of the steepest slope.

Derivatives are the mathematical notion of slopes. In most mathematical settings the surface is expressed in a standard coordinate system (most often the (x,y) type Cartesian coordinate system). In this setting, the slope used in the search for the minimum is then expressed combinations of slopes in the x, y etc directions (e.g. if the desired direction if North-East, one could equivalently say so far North plus so far East). If there is more than one principal direction (in this case there are two, one for each of the portfolio components), the derivatives are called *partial* derivatives. It is for this reason that many optimisers require that the slopes "exist" and that they are defined in some mathematical sense.

3.1.2 So is it as simple as all that?

Not in the real world, and not even in many not so real worlds. Keeping in mind that all of the discussion so far is focused on purely the mathematical problem following once the model has been accepted (and thus the "foundation" is assumed to be solid), we are still exposed to a wide variety of difficulties. We introduce a few important mathematical issues here (by way of more pictures), and later we introduce some of the concerns about the model. By and large the problems introduced immediately below are still quite tractable, and would be analogous to dropping our car keys in the light, but not in our line of sight, or perhaps in the light but now in the grass as well (we will find the keys but extra effort/technique is required).

3.1.3 Correlation

Even a small change to the problem above can lead to challenges. Figure 4 shows the same variance problem from before, but here the two investments here are positively correlated (and thus the parabaloid has become "squished"). If we were to employ a technique relying on the slopes, difficulties could arise. The reason for this is that the direction of steepness along the "flattened" sides does not point towards the bottom in quite the clear cut way that the maths would like. Thus, the search direction would take steps in a sub-optimal direction. In many cases this would simply slow the search down as extra steps were made. Alternatively, there are improved methods that are "smart enough" not to be fooled this way.

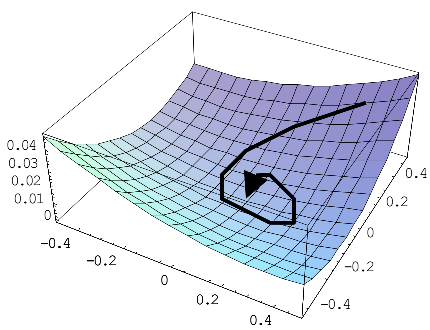

However, even some of the smarter methods can run into trouble. If there was very high correlation and large differences in the standard deviations, and thus a very "squished" bowl, then the first step of the search could point the next test point completely outside of the area of interest as per Figure 5. Once again, additionally clever machinery can be employed to prevent such difficulties.

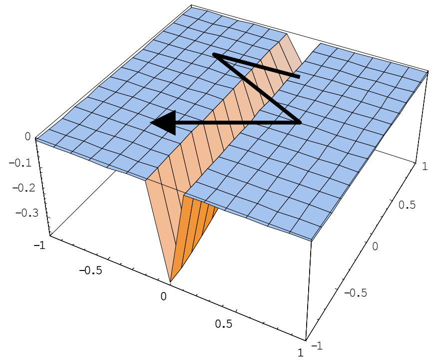

3.1.4 Constraints

Another of type of challenge arises in the case of so-called "constrained optimisation". In making our selection for an optimal portfolio, we may be constrained in various ways. For example, are you allowed to short a position, do you have a pre-set limit as to the maximum or minimum that may be invested in one or more instruments, and so forth. These types of constraints complicate the "bottom searching" process, since the search steps may wish to move into a "valid" portion of the bowl, but one that is disallowed by the constraints.

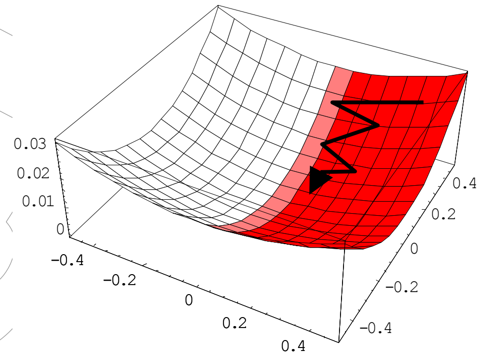

In general, there is a great deal of well researched and proven mathematics to deal with constrained optimisation, but as always their applicability is problem dependent. In the case of our mean/variance problem, and even for a large number of normal and even "wild" constraints, there is no problem in finding an algorithm most of the time (although the purchase price increases dramatically with increasing sophistication of mathematical wizardry).

3.1.5 Many investments

Many investors and traders need to compose their portfolios by selecting from a very large number of investment alternatives (thousand of stocks, bonds, currencies etc). As it happens moving from an admissible universe of a few instruments to one of very many introduces new difficulties. Some of the difficulties are purely technical, while others are subtle.

The mathematical process of optimisation almost always involves matrix algebra and in particular so-called matrix inversion. The number of calculations required to invert a matrix grow disproportionately with the size of the matrix, and the number of instruments and the number of constraints determines the size of the matrix. The specific increase in the calculation costs is dependent on the exact problem, but it can easily be a quadratic process (e.g. including twice as instruments will cost four times as much in calculations). This means that for very large problems that need to be solved in a practical way (e.g. sometime today as opposed to a week/month from now) one needs to obtain either or both of more clever optimisation formulations, and or faster computers.

However, there are more subtle types of difficulties that arise when many instruments are considered. One of these arises as every time we add an additionally investment into the admissible pool, we are also adding a kind of informational noise[5]. Thus, when the search process tries to move through an increasingly noisy forest of possible investments it (although mathematically correct) may choose compositions that are not sensible in the real world. For example the properties (average return and standard deviation) of an investment may prompt the calculation to suggest selling an astronomical amount of it, while buying an astronomical amount of another (something that would not even be practically possible much less sensible). There are techniques for dealing with such problems as well, but now we are starting to rely on "human post-processing" to check for the investment sensibility that the machine lacks (a kind of audit process we should do anyway).

3.1.6 My world is wobbly

The inputs to the mean/variance problem (returns and their standard deviations) are calculated from historical performance. It is easily demonstrated that these numbers are non-stationary (they change with time). Moreover, even by changing the sample period[6] looking backwards these inputs change again (even though the histories have not changed). This leads to uncertainty over the inputs, and perhaps you would like the optimisation process to account for this uncertainty.

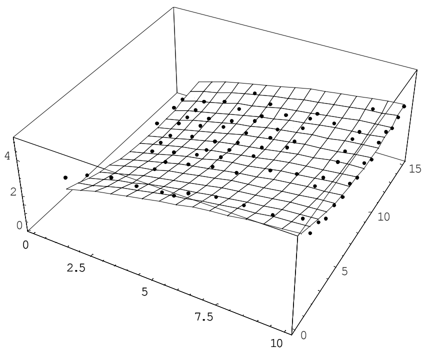

Figure 7 illustrates such a problem. The exact position of the surface of the bowl in not known precisely. Only some "average" surface is known. There are a number of ways of dealing with this problem. First, you could spend the extra money on a "stochastic optimiser". Keeping in mind that as usual the mathematicians producing such machines do not concerns themselves with the portfolio problem as such, but rather the properties of surfaces and how to "search". In this light, we hope that it is easy to guess that even more "human post-processing" is required to ensure that sensible asset allocation results (regardless of what the optimiser says).

In some cases, experienced market professionals are able to make (possibly heuristic) adjustments to the basic (non-stochastic) optimiser to account for this issue. In this case, keen insight into the physics (i.e. a modelling effort) is used to overcome the mathematical limitations.

[5] Technically, "noise" can arise for a variety of reasons, some of which are no fault of the end-user (trader, investor). For example, the covariance matrix can be near ill-conditioned due to some real, but pathological, combination of market inputs. Additionally, this may be exacerbated by round-off errors, which get worse with number of inputs. A near ill-conditioned matrix can result in "wild" answers on inversion.
[6] The science (or art) of dealing with calculations on histories is a giant can of worms onto itself and the subject of a future ***ARTicle***. In some ways it is more difficult than the optimisation problems.

3.1.7 I can't buy 1.5 contracts

In all of the examples so far, the optimiser has not been told that trading occurs in "blocks[7]". Thus, it would quite happily, and quite correctly for its given rules, return an answer involving the buying and selling of fractional amounts. In some cases this may not bother a trader too much. For example, rounding the answer from a buy of 573.2 T-Bond contracts to 573 contracts would (in many cases) not be too serious, although it would no longer be at the mathematical optimum.

The mathematically "correct" technique to overcome this problem is to use so-called "integer programming". This method forces all of the results to be integers, and via some manipulations one could obtain "block integer" (for example trade DEM swaps in blocks of 25 million notional).

There may be circumstances (especially for complex models) wherein a restriction to "blocks", would make a material difference to the machines view of the optimal solution. For example, it may decide that if integer rules applied then it would not suggest 573 T-Bonds, but rather 230 T-Bonds, and simultaneously alter other components equally dramatically. This does (and should) scare traders, and once again the first line of defence is a good understanding of how the optimiser thinks, and then make appropriate adjustments.

Nevertheless, there are plenty of integer-enabled optimisers about (historically used in linear programming but there is reasonably broad availability for quadratic and for some more general non-linear problems), and as always for an extra fee you can have this as well.

3.1.8 Transactions Costs

This segment should have a rather more general heading as it applies to a very much broader scope of issues, and some very important. However, these concerns are well illustrated by some of the challenges arising from accounting for transactions.

As at the outset, if you are a long-term buy-hold investor (or rebalance infrequently) then some of these issues may seem (at first) irrelevant.

In the simplest case, one could simply apply a constant cost to these (and remembering to make consistent transformation from the "cash" units of costs to the "returns" units of our prototypical mean/variance problem – and which in some realistic costs models may not be so easy) to yield a slight variation on the original equations:

$$R_{Port} = a_1 r_1 x_1 + a_2 r_2 x_2 \qquad (0.3)$$

[7] Even for OTC contracts, market conventions compel trading in minimum amounts. For example, try to deal in GBP 100,000 notional of a Sterling swap.

$$Var_{Port} = b_1^2\sigma_1^2 x_1^2 + b_{1,2} 2\rho_{1,2}\sigma_1\sigma_2 x_1 x_2 + b_2^2\sigma_2^2 x_2^2 \quad (0.4)$$

where the a_i's and b_i's capture the cost influences.

In the simple case of constant costs, the optimisation difficulties are not altered very much. However, modelling for realistic costs can have dramatic effects. Consider a couple of possibilities:

1) The costs vary linearly or (smoothly) non-linearly with traded size. For example, there is a fixed bid/offer spread or a smoothly[8] increasing bid/offer spread. In this case, the portfolio returns function is also non-linear and the variance function may become highly non-linear. Thus, intersection of the two surfaces can become quite complicated (and it is the optima of the intersection that we need).

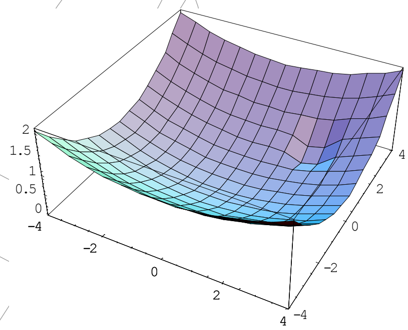

The types of extra complications that arise vary, but as an example, consider the possibility of a "bumpy" surface. That is, the ball bearing is now set adrift on a surface that has "mini bowls" on it, as in Figure 8. Technically, we have local minima. It is quite possible that the ball bearing lands in a local minimum thinking to be at the global minima, and clearly it would be mistaken, as would the trades based on those results. Not surprisingly there are, for additional fees, vendors of optimisation engines that aim to deal with mistaken local minima.

2) A potentially much more difficult real world effect may occur if we model the costs in a stepwise fashion. For example, and as is often observed, bid/offer spreads widen suddenly with increasing volume (as opposed to widening smoothly). Figure 9 illustrates a surface that has such stepwise character. Clearly, there is no "slope" at the edges of the "plateaus", so any technique relying on slopes would fail. Again, there are optimisation engines that have a chance at dealing with such discontinuities as they use search procedures that do not need any slopes or possibly only a "few" slopes (yes, these are extra fees as well).

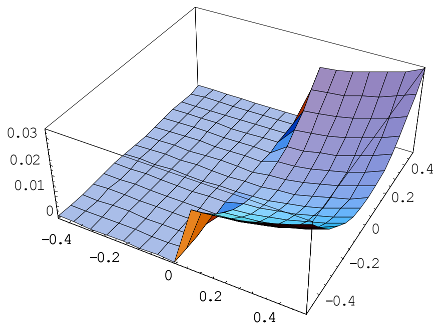

[8] Mathematically, smooth can have a slightly different interpretation than the obvious. Here, for example, we will see that the derivatives of the functions need also be smooth if we are to rely on search methods that employ slopes, since a non-smooth function in mathematics means that it may not have a useable slope.

3.1.9 The story so far

Trading is a complex business requiring forecasting, and a compulsion to achieve the best possible "results". A widely accepted view of "best possible result" is maximal risk/adjusted returns, although in reality there are human factors that necessarily restrict this. Nevertheless, focusing in on maximum risk/adjusted returns, one may appeal to quantitative methods for assistance.

Such techniques are comprised of the three steps of modelling, solving, and auditing the strategies/methods. We have in this section focused primarily on the second step of "solution" methods used in mathematical optimisation. This is only one category of optimisation, and later we will talk of optimisation problems not easily or at all mathematically treatable.

Taking the classic case of Markowitz mean/variance optimisation as a prototypical example of classical optimisation, we could learn about the qualitative aspects of optimisation, and its pros/cons. Very little discussion was given to the underlying model problem (some of which follows below), and for the illustrations we (temporarily) accepted its validity.

For problems like (simple) mean/variance optimisation, mathematical methods are reasonably well understood and so long as the number of real world features is not severe or are not simultaneous, it is possible to obtain (at additional expense) suitable algorithms. In some cases, real world insight is required to augment the optimiser. We note, that if many or all of the features that were raised as "challenges" need to be addressed simultaneously, then there may not be a suitable off-the-shelve package, and independent research may be necessary. We also note that in some cases heuristic trading experience can be used to overcome subtle mathematical difficulties (although these may not enjoy the support of mathematicians).

Philosophically and practically, we always insist on linking methods/strategy to actual performance.. If we had to make a choice to do excellent mathematics, or to obtain excellent P&L, well …. we are traders … , but at least we hope that we made progress on some insight into one family of optimisation problems.

4 I live on planet Earth

If you prefer to hide your shares under the mattress until retirement, then no more needs to be said. However, most professional investors/traders are governed by self-imposed or mandated limits, instrument dependent constraints, exit conditions, etc. Even, if you do not have such limits, you will (almost surely) still find that a month[9] from now the optimiser recommends an optimal portfolio that is very different from that one it suggested (and you traded) earlier: will you rebalance?

For any of the dynamic conditions above, all optimisers of the type in discussion here will be meaningless. There a variety of reasons for this, and some of these are discussed next, but the primary issue is that the “core” model that is at the foundation of the optimiser is inconsistent with either or both of real market dynamics and or real world human behaviour. In this section we are focusing on “step1” or the model. We will find that we have dropped our car keys way off in the dark, and man-eating clams populate the dark. Sometimes, we will even find car keys that will not start/open our car (so it looks right, but it wont work).

4.1.1 Market Non-Stationarity

As already alluded to above, the markets vary in and the way in which they vary also changes the input parameters (returns and their standard deviations) over time[10] and thus are not “stationary”. This means that such “static[11]” techniques could easily recommend very different “optimal” trading from minute to minute. It is important to recognise that the machinery is working correctly for the “world” that it assumes, but that is simply not world many of us live in.

It is possible to (largely) account for non-stationarity with “adjusted” conventional methods or techniques that are based on non-stationary methods from the outset. In either case, it is our experience and opinion that the machinery needs to be considerably more sophisticated.

At the very least, the wise trader would need to take the “standard” optimiser result with a grain of salt (and sometimes a very large one).

4.1.2 Holding Period Activities

Static optimisers assume no activity during the holding period, since as far as they are concerned there is no holding period, but only this instant in time. Most of us, however, will almost surely need to alter the portfolio for many important reasons. If these reasons are

[9] This is true on virtually any time scale as the motion of the markets and non-stationarity of the parameters can change dramatically by the minute; you don’t need to see a GKO default or LTCM bailout for this experience.
[10] Also, as alluded to above, the assessment of these numbers is a bit of art anyway, and may be the subject of a future ***ARTicle***.
[11] Here “static” refers to the notion that forecasted portfolio does account for any changes over the holding period: a crucial disadvantage.

unknown to the optimiser (and for most optimisers they are unknown), then its recommendation cannot be consistent (except by pure chance on the odd occasion) with the world you live in.

To see this consider that static optimisers assume that the inputs fully capture the parameterisation of the dynamics. For example, in the case of mean/variance optimisers it assumed that the variance (input as standard deviation) captures "wobbliness" of the returns. But, your portfolio returns wobble for reasons outside of pure market dynamics. Rebalancing the portfolio for say limit constraints immediately locks-in a P&L (and hence return) effect that is not accounted for in the market behaviour standard deviation. Had the optimiser know about such possibilities it would almost surely have recommended a different "optimal" portfolio at the outset. So you found "some" car keys, but they just won't work properly.

An optimiser that accounts for holding periods and for holding period activity would need to know about all the dynamics of the markets as well as all of the actions applied to the portfolio for "management" reasons (strategy and limits). Such things are possible and exist, but it should be obvious that building this type of machine is non-trivial. So you found the "right" car keys, but you had to defeat the man-eating clams.

4.1.3 Market Uncertainty and "Risk"

Almost exclusively[12] market uncertainty is handled via statistical/stochastic methods. While we do not have the room here, it can be shown that the vast bulk of human knowledge in respect of stochastic methods are tractable for only a small selection of problems, and importantly not for many interesting problems in trading. There are instances where we "shoe-horn" the understood (but inappropriate) maths into a semblance that may or may not be reasonably close to the real dynamics, but all to often these are so complicated that no non-academic will grasp (or have the time to grasp) the many important and subtle shortcomings[13]. So the car keys are off in the dark, but some insist on looking only where the light is.

As it happens, even in the cases where the maths are understood and appropriate, traders have an uncanny capacity to abuse and pervert the usage and meaning of the results.

In order to illustrate this point, let us review a few concepts of risk/return. A common first approach is the convert a history P&Ls into a history of returns ($ROI = (P_2 - P_1)/P_1$) as per the Figures 10 and 11 below:

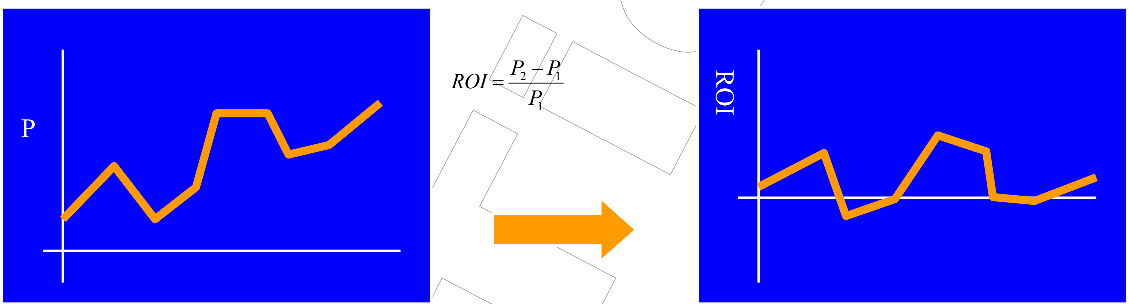

The returns (or for that matter the P&Ls) can now be analysed statistically. The most common "next step" is to get your software package to tell you the standard deviation of the returns. Which in itself is all right, but the following step is where things usually get dicey.

Having calculated the standard deviation, it is market practice to "assume" at least two things. First, it is common practice to assume that the standard deviation is equivalent to risk (which is quite silly since an up-swing in your P&L is not really risk). Second, it is common practice to assume that the standard deviation by itself fully captures the variability of the returns.

[12] We are aware of alternative methods, and even use a few our selves, but the statistical/stochastic approaches are by and far the most common.
[13] Of course this is a happy state of affairs for arbitrageurs.

This second assumption can be fatal. Using a statistical approach what we should have done

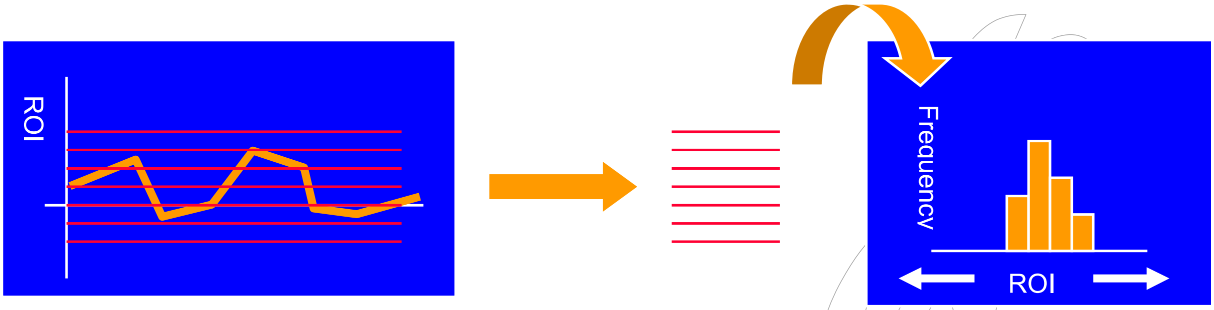

is to create the distribution of returns (a histogram) as per Figures 12 and 13. The distribution may or may not be Gaussian (i.e. nice bell shaped symmetric etc). If the distribution is Gaussian (or at least some second order process), then we have a credible shot at using the standard deviation as a measure of "variability". However, this is rarely the case. Indeed, while analyses of price behaviour may yield approximately Gaussian distribution in some cases, analyses of (dynamic) portfolio distributions hardly ever do.

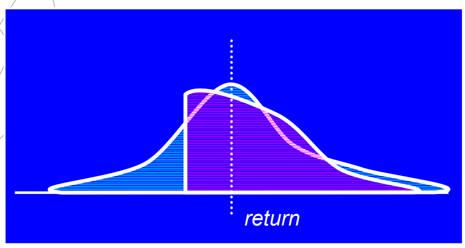

Moreover, it is precisely the aim of structuring and dynamic strategies to alter the shape of the portfolio's distribution to one that is more "optimal" for the given risk preference of the investor(s) such as might be in Figure 14. The truncated distribution could arise from a desire to limit downside risk and be implemented by dynamic trading strategies or options.

If any of these aspects apply to you, then standard deviation by itself does not fully capture the variability of your portfolio returns. Thus, relying solely on standard deviations as a measure of variability will necessarily misguide optimisers that make the same assumptions (this includes the bulk of the mean/variance methods in common usage), and thus to make sub-optimal asset allocations[14].

So the first question that needs to be answered is "what are the properties of the markets/dynamics that are important to you?" Knowing this (and this may require some research and analyses), at least you have a chance at employing the right tools, or at least knowing when not to rely on dodgy tools. As traders we are less worried about using imperfect tools (and hopefully know their characteristics), as compared to not knowing that we are using imperfect tools.

4.1.4 But I make money with these tools

An important notion frequently omitted from such considerations is that of a precise understanding of how trading profits were made. Does this sound presumptuous? In fact it is

[14] Incidentally, all of this applies equally well to any and all other investment considerations where such statistics and stochastic methods are used (e.g. options trading, term-structure models etc).

often the case that while a big picture understanding of the origins of the profits are known, in many cases there are crevices that are not revealed, and market effects that are not admitted.

As an illustration, consider that a pure buy/hold strategy in say the S&P or mainstream bonds would have consistently made profits for even relatively short-term investment horizons (e.g. 1-2 years or more) during the bulk of the last 70 years. So how much excess profits can be attributed to the use of an "optimiser" (or for that matter any technique). For this question to have meaning, some form of historical analyses is required (say a moving average/variability of the excess returns of some sort). While all of us must show this type of history for the "over all" P&Ls, very few bother with analysing the components of the excess returns. As such very few traders know (or reveal) specifically how the profits were made, and not surprisingly we prefer to put it down to our "skill" instead.

In this sense, it is possible to make money while employing an "optimiser", but it is clearly possible that sub-optimal optimisers by themselves may not be responsible for the profits. Of course unless the analyses is done, no one will know.

To be clear, we would not presume to tell anyone how to trade or what techniques to use out of hand. Indeed, it is a deep and fundamental theme here that risk/return decisions are subjective and it is up to the individual to decide what is "best" for them. However, we are advocates of "knowing".

4.1.5 What does "optimal" actually mean?

In all of the cases so far discussed, it is important to emphasise once more that a machine that is called an "optimiser" may very well have a different definition of "optimal" than the one you hold. Hopefully, we have already belaboured the point that the mathematical optimisers' only job is to "find the bottom of the (possibly complicated) bowl", and even this is usually accomplished indirectly. The shape of the bowl reflects some idealised model of the problem faced in the real world. The model almost surely does not account for many important real world effects.

We believe that there is a very real danger of transposing one's imagined or expected definition of "optimiser" to that built into some machine. This often arises since the technical details can be overwhelming, and in any case the salesmen/women will apply "lets blind them with science".

Not surprisingly our view is that the first line of defence is to "know" enough about these machines to understand what they do not do.

4.2 Non-Mathematical Optimisation Problems

At the outset we had presented optimisation problems as two families. One family (those covered above) are those that are directly amenable to mathematical methods (e.g. can be completely expressed as some sort of "bowl"). However, there at least two important (practical) cases where mathematics will fail us.

First, some problems cannot (practically) be expressed mathematically. To illustrate this, consider that we have so far glossed over the critical requirement that using, say, mean/variance optimisation the results are incomplete (even if the model is accurate). This is because the trader must decide what level of returns to demand for a given level of "risk". A perfect mean/variance optimiser can only tell you the composition of the "optimal" portfolio for either a given desired maximum return, or for a given desired minimum variance. However, it cannot decide what either of these "given" values should be. These levels are purely subjective and a reflection of individual risk preferences.

In this example, the complete optimisation process might require that the trader/investor "picks" the desired risk/return characteristic from say the efficient frontier (or the securities market line). This in itself is a form of optimisation, but one done by "eye ball", since attempting to program a computer to make such decisions seems intractable.

Second, some types of desired machinery may be too expensive to construct. For example, there are on the market "back test optimisers". These (usually) are restricted to backtesting simple trade ideas with the hope that the "optimiser" will select the trading strategy parameters that are the "best"[15]. One might consider taking some of the mathematical machinery from above and embedding it into such a process. Unfortunately this may well require many man-years of programming and thus be (practically) intractable. Moreover, even armed with a budget for such, there may be cases where the traditional optimisation mathematics would fail anyway. A hint of this type of problem was raised during the discussion of including many and stochastic variables in the mean/variance optimiser. Once the problem is sufficiently large and noisy, the traditional methods can very well begin to suffer from subtly numerical mishaps, leading to incorrect or nonsensical results[16].

[15] We use a very sophisticated implementation of this idea that does permit full accounting of all behaviour and market processes.
[16] Our own approach to these problems is based on a completely different (and unique) fundamental mathematical basis to avoid such difficulties.

5 Summary

The desire to be happiest has been approached from the perspective of optimising portfolio risk/return based on subjective preferences. This involves forecasting, and hence quantitative methods for those who are primarily dynamic traders (as opposed to long term buy-hold).

The complete optimisation problem is built on a model, and then sent to a "solver" machine. These models and machines need to be audited and benchmarked against the real world.

The optimisation problem falls into mathematically tractable, vs. heuristic families. The mathematical methods can be viewed as the expression of the problem as some sort of (possibly complicated) bowl, with a search engine that looks for the "bottom". The prototypical mathematical portfolio optimisation machine was taken as the Markowitz-like mean/variance process. It was initially assumed that the underlying model assumptions were valid. This subset of optimisation problems is a relatively straightforward mathematical problem for many cases.

However, the underlying model has some deficiencies. Some of the important shortcomings are that stationarity, Gaussian behaviour, and zero-length holding periods (i.e. static) belie the assumptions. Thus, while the mathematical solvers can easily produce an "answer", it is likely that many traders would not consider these answers as optimal or even in some cases sensible.

There are both heuristic and more formal methods for overcoming these deficiencies. Care should be taken in constructing these as they may be expected to be complex and expensive, but you get what you pay for.

Ultimately, the value of any such device is subjective and must be compared to individual needs and circumstances.

What's next?

In "***ARTicles**: Optimisation and P&L – Part 2*", due out in May, we present a note on using and interpreting portfolio optimisation via mean/variance methods, and include a *free* downloadable Efficient Frontier Optimiser at http://www.arbitrage-trading.com.

ARTicles