

Equity Portfolio Construction

Demystifying optimisation

PATRICK BRAUN, SCOTT LIU, MSCI BARRA

If ever asked to characterise the word ‘optimisation’, many of us would draw a black box with the secret hope of never having to look into it.

Yet every day, without noticing, we are exposed to the predictability of optimisation tools. Have you ever queued for a lift, hoped for mobile phone reception in a busy airport, or used a web search engine? These are just a few examples of every day processes that use optimisation.

While the hedge fund community talks about ‘alpha’ and ‘leverage level’ in its different forms, the world of optimisation has its very own language, such as ‘objective functions’, ‘solvers’ and ‘constraints’. Without falling into another jargon trap, this article looks into specific areas where non-quantitative long/short (L/S) fund managers, and prop traders could benefit from exploring optimisation techniques. Given the ongoing financial challenge to deliver ‘pure alpha’, the few language and technical hurdles can be worth overcoming.

The constant presence of beta

For hedge fund managers and prop traders running long/short equity books, a key to delivering their specific value-added lays first in their ability to neutralise directional market exposures, or in general to minimise the overall costs of their own beta.

In the case of an absolute return hedge fund, a ‘safe harbour’ portfolio with a low directional market tilt and low leverage may represent this beta component of the strategy. The fund manager may seek to minimise the volatility of this component of the book, subject to certain constraints such as no directional bias, and a reduced gearing. The construction of such a component will make sense to many quantitative traders. It would however prove hard to build and rebalance consistently without the assistance of optimisation tools.

Pure stock pickers, who seek to focus on the idiosyncrasy of their individual investment positions, can also use optimisation tools to maintain dynamic overlays onto their initial selection bets. This hedging component is typically a long/short equity basket built either to reduce the book’s net exposures to market factors, or to avoid unwanted style drift.

In these two situations, the use of optimisation in the portfolio construction process enables traders to build books more aligned with their initial market, or company-specific views. This leads to a reduction of volatility coming from beta territories, and a clearer allocation of the risk budget into areas of expertise.

Some alpha-centric problems

At times when small teams of specialists look for alpha generation ideas across an expanding universe

of equity data, analysts can use optimisation to systematically search and segment their original opportunity set. For example, those confident with the market equilibrium concept may generate insightful asset-level analysis, ranking stocks by implied returns under different weighting schemes.

When relying on bottom-up qualitative research, long/short equity hedge fund managers may favour equal-weighted books, built around some strong individual company views. However, whether such a balanced approach delivers all the potential alphas generated by their own research process would need to be confirmed, and if possible, back-tested over a reasonable time period. Such a simulation exercise may bring good news: for similar or lower levels of risk, the proportion of expected returns passed into the L/S portfolio may increase if the hedge fund manager softens this ‘equal weight’ constraint. For strategies with few positions, allowing for uneven weights on both long and short sides may enable managers to reach higher expected Sharpe Ratios, and distribute risk more evenly across their entire book.

“For many hedge fund managers, the struggle is with the practical aspects of implementation”

Fig.1 illustrates such a situation. In this particular instance, the hedge fund manager uses a proprietary scoring process to screen and rank a universe of 100 highly traded stocks. Based on these forecasts, he or she constructs a L/S equity portfolio of 30 stocks, with 20 equal-weighted names on the long side, and 10 on the short side. As shown in Fig.1, this 300%-leveraged portfolio is intuitive in its structure but its volatility distribution is unbalanced: shorts (-100% of the net value) absorb 64% of the risk budget, versus 36% to the long positions (200%). In the presence of such a linear alpha model, same forecasting ability for longs and shorts, the hedge fund manager should reach a better risk balance across the book.

He or she now amends the portfolio by adding an overlay basket to the L/S fund, and by using a portfolio optimisation tool with his or her own stock-level constraints. In this example, the overlay universe comprises 30 large cap stocks, with no specific return forecast or transaction cost estimates for these.

The optimised portfolio the manager selects on the efficient frontier is a book with an expected return equivalent to that of the initial fund, and a lower risk forecast. Alternatives include the book with maximum expected Sharpe ratio, ie., the highest expected return/risk combination, or one with higher expected returns, which would absorb the overlay implementation costs.

As shown in Fig.2 (overleaf), this new optimisation-enhanced book still reflects the initial stock convictions. Only seven of the previous longs (lowest positive scores) now turn into zero-weight positions. Risk results differ from the original book, with a lower volatility (18% vs 30%) for an equivalent expected return. The risk distribution is much more balanced than before. The long positions now consume 60% of the risk budget (vs the previous 36%), for 35% attributed to the shorts (vs 64%). The new overlay basket provides the expected volatility dampener to the overall strategy, and takes only 5% of the total risk budget.

At the stock level, risk contributions from individual long and short positions have now converged. Overlay aside, the average risk contribution is 3% (vs the previous 1.8%) for a long position, and 3.5% (vs 6.3%) for a short one. These new figures are much more in line with the symmetric scoring model originally put in place.

Be it through volatility management, or more consistent portfolio construction, the risk-adjusted performance of L/S equity portfolios may benefit from optimisation-assisted enhancements.

Refining the implementation stages

Even those hedge fund managers and prop traders who resist the idea of segmenting expensive alpha and cheap beta will find that optimisation tools are probably already playing a part in the market implementation of their current strategies.

In times of fierce competition, the ability to identify the next investment idea (on both longs and shorts) and convert hypothetical alpha into real profits are equally important. The distinction between these is one reason why many quantitative investment teams continue to develop sophisticated implementation programs using optimisation engines, and proprietary transaction cost models. Other, more fundamental, hedge fund managers are also daily users of the broad family of implementation analytics; often via the prime brokers they directly trade with, or through internal expert groups who deal centrally with trade execution. The tools used in these cases focus more specifically on the management of execution risk. This is a different optimisation problem from the one described earlier in this article. Here, the challenge is in executing a

pre-defined series of trades with either the lowest possible market impact (with additional market risk), the highest trading speed (with additional market impact-related costs), or with an acceptable trade-off between these. Robert Almgren and Nobel Prize winner Robert Engle have extensively documented the theory surrounding this finance area.

For many hedge fund managers however, the struggle is with the practical aspects of implementation. It starts with the simple gathering of reliable historical data to assess transaction costs. Although all hedge fund managers will


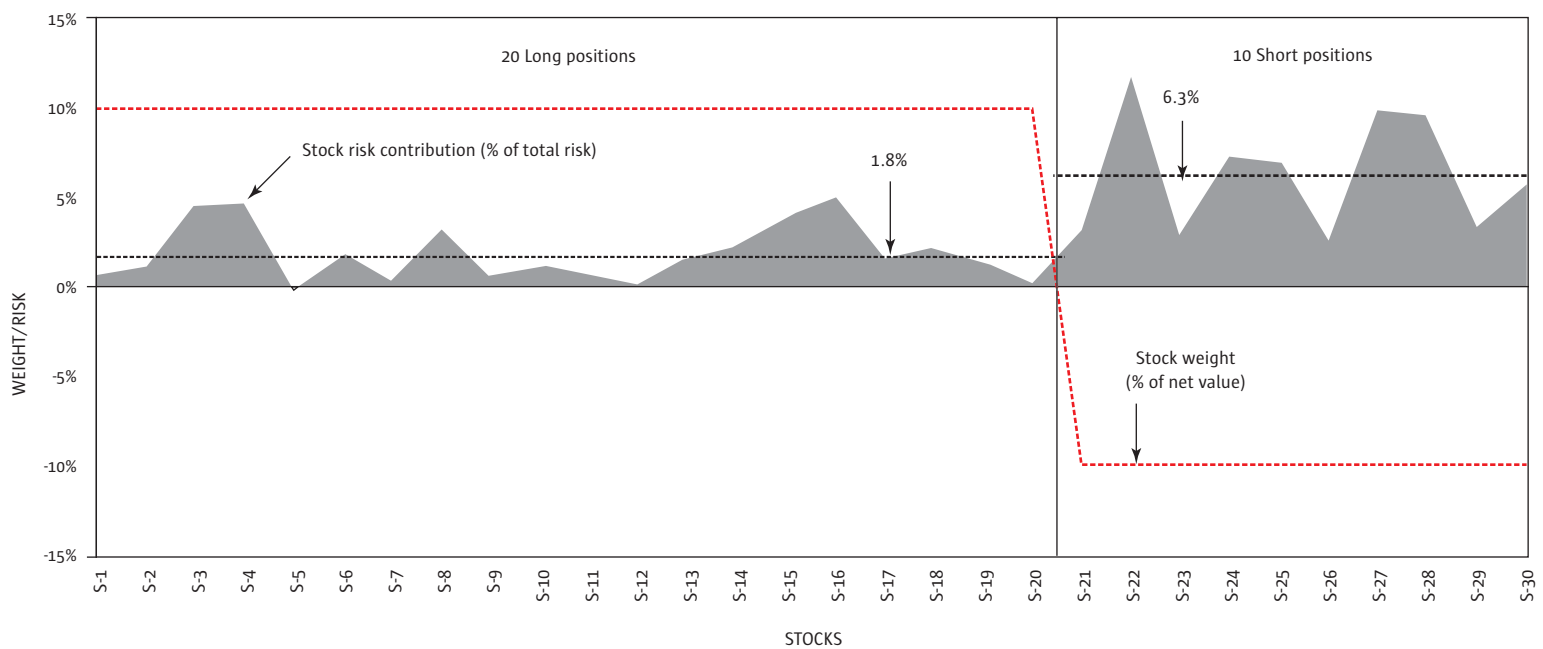
know the price they paid earlier today for their own positions in stock A, B or C, only a few may have recorded data showing the broader market bid-offer levels and volumes at the very time their trades took place. For managers who can estimate different transaction cost functions for their assets and can easily source stock-level 'repo' rates, it is possible to use a portfolio optimisation tool to fine-tune simultaneously the implementation of their alphas on the long and short side. Based upon the hedge fund manager's own return expectations, this optimisation process would seek to form, for example, the L/S book maximising risk-adjusted

returns under user-defined constraints on net exposures, and custom liquidity flags.

The successful market implementation of original research ideas is now so important that some fund managers consider it a source of alpha in itself, ie. a team-specific mark of value added. Whether it deserves this label or not is an open debate around the broad access to innovation. In both cases though, solving complex, practical portfolio implementation problems will remain a critical area for future developments of financial optimisation tools.

Fig.1 Weights and risk in the initial equal-weighted L/S fund

Source: MSCI Barra



Example analysed using the Barra Aegis System, Barra Optimizer and risk models

Solving with style, if possible

As shown previously in this article, hedge fund managers and prop traders may benefit from using optimisation tools throughout their investment processes. This is applicable not only to hard quants, but also to more fundamental investors managing pair trading books, or focusing on companies' activism.

The enhancement areas illustrated above range from the systematic screening of stocks (daily generation of implied returns), to the L/S portfolio construction (alignment of risk distribution and alpha model, with

overlay), and the market implementation of the L/S book (optimisation with transaction cost models, and 'repo' rates).


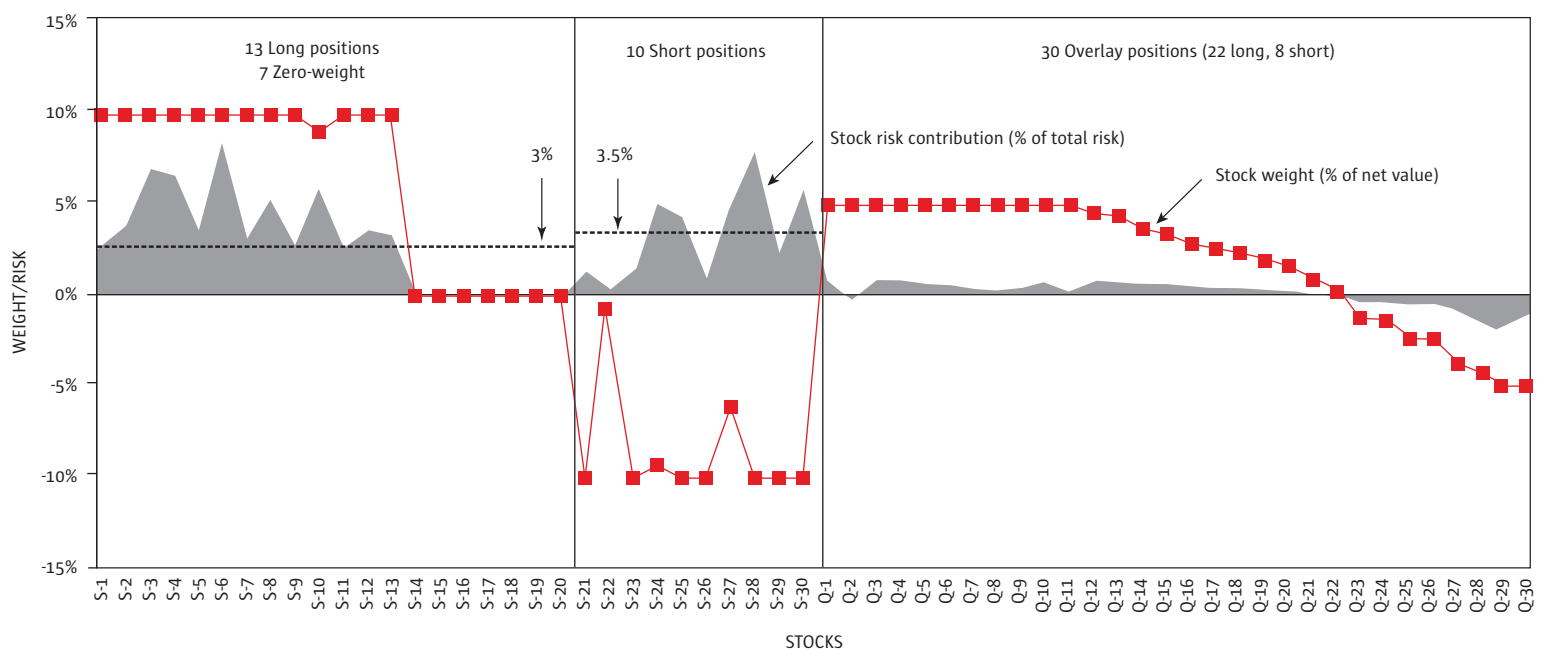
Although a good knowledge of automotive engineering is not a prerequisite for holding a driving licence, a general understanding of what affects an engine's performance can often guide our commuting habits and road behaviour. The same applies to the world of portfolio optimisation. For a hedge fund manager or trader, a good understanding of the optimisation problem in the first place can lead to a more informed selection of solvers (the engines),

and a better usage of the algorithms currently available. It can also accelerate innovation through the development of new engines.

Purists and more casual drivers may argue as to what best describes the performance of an engine. One may express it in the form of low fuel consumption and low emission figures, the other in terms of high torque and speed. The choice of engine (and vehicle) will vary, depending on how one defines and measures this important 'performance' function. What is intuitive in the world of car engines is again applicable to the world of financial optimisation

Fig.2 Weights and risk in the optimisation-enhanced L/S fund

Source: MSCI Barra



Example analysed using the Barra Aegis System, Barra Optimizer and risk models

engines. The financial functions the solvers will seek to maximise/minimise will of course differ from the previous car-related 'performance' ones. In both cases however, the adoption of a multi-purpose engine may fail to please high-end users. In the particular context of L/S portfolio optimisation, objective functions can become complex as users translate intuitive statements like "100 assets max" into hard constraints, or try to include short rebates. Across the different solvers that will return solutions to these types of problems, few will form an integrated optimal long/short portfolio. Many will take two-step approaches, and blend an optimal long portfolio, with a separate optimal short. When acceptable solutions to these complex problems exist, only specifically-developed portfolio optimisation tools will return these solutions within a reasonable timeframe.

In simple terms, optimisation engines and car

engines have many things in common. Within their own product families, they differ a lot. While all may work fine in general circumstances, under hard usage conditions such as dealing with the specifics of L/S portfolio construction and implementation, the number of suitable engines rapidly shrinks.

One way towards industrial customisation?

The different examples of L/S portfolio enhancements described within this article may have by now confirmed that, for many non-quant traders and hedge fund managers, the few language and technical hurdles surrounding optimisation can be worth overcoming.

The growing number of institutional mandates allocated to the hedge fund world could also add to these motivations. When these mandates include elaborate constraints for the fund manager to adhere to, a good understanding of their

impacts on the strategies may help discussions with sponsors. Can all processes really deliver a good performance under pre-listed constraints such as a limited leverage, or a maximum downside risk? Using an optimisation tool to back-test broad strategies under a set of mandate-specific requirements can also help provide these answers. **THFJ**

Biography

PATRICK BRAUN

Patrick Braun is a Vice President, Product Management, at MSCI Barra in London.

SCOTT LIU

Scott Liu (PhD) is a Vice President, Operations Research, at MSCI Barra in Berkeley, California.