

Proprietary trading: truth and fiction

Peter Muller, who has spent eight years building what he describes as a ‘reasonably successful proprietary trading group’, introduces some of the issues behind his strategies—without giving his game away.

My group’s approach to trading is model based. We spend considerable time researching and back-testing trading strategies before we implement them. If you read a journal article on an asset-pricing anomaly, the chances are fairly high that we have read it too, probably verified the research, and occasionally used it in a modified form in one of our strategies. For competitive reasons I won’t describe much about what we trade, how much we manage or what our track record looks like. After all, I hope my group continues to generate abnormal trading profits for quite some time.

I feel fairly confident that our results are not entirely due to luck (although I sometimes wonder if we’re that different from the lucky monkey who has managed to tap out *Hamlet* on his typewriter). To my knowledge our results are not related to market direction (we are market neutral), liquidity premiums (our positions are liquid), skewness premiums or other forms of optionality (our upside and downside risks are symmetric). Nor do I believe they are related to the lower trading and financing costs that may be available within an investment bank (we calculate our results assuming we pay commissions and financing spreads). And of course, although we are housed in an investment bank, we have no access to any information about our bank’s clients. But there is always the possibility that we are exposed to risk factors that we don’t know about.

Managing risk

We manage risk in two ways. First, for each strategy, we impose limits on some or all of the following: capital usage, expected portfolio standard deviation, value-at-risk (VaR), position liquidity, and exposure to various pre-identified risk factors (for example, a rho limit caps the first derivative of portfolio return with

respect to a parallel shift in interest rates). Incidentally, our client, the bank, imposes overall limits on us for many of these variables and for some of their own.

But the most important risk is the possibility of our models not working correctly. To minimize that risk, we set loss targets for strategies—if we lose more money than the pre-specified target then the strategy is re-evaluated and shut down for a while (perhaps forever). This is not that different from the old school of proprietary trader management: ‘Go ahead and trade, don’t do anything too risky, and if you lose more than \$x we’re going to shut you down.’

Our strategies are evaluated by looking at reward/risk measures. For symmetric, market-neutral strategies without significant tail events, the Sharpe Ratio (SR) is probably the best *ex ante* measure. SR is defined as the portfolio annual excess return divided by the annualized standard deviation of that return. Our benchmark is cash, hence measuring excess returns is appropriate for our portfolio. For long-only managers, the Information Ratio—which measures excess returns relative to a benchmark—is more appropriate.

When we evaluate past performance, we also look at peak-to-trough drawdowns (a measure of the maximum drop between consecutive maximum and minimum values of return over the life of the strategy) as an additional risk variable. This can help pick up serial correlation in portfolio returns that the Sharpe Ratio doesn’t capture. Also of interest is the fraction of expected gross profits consumed by expected transaction costs. The higher this number, the more money we expect to lose if our model stops working.

At least some of our edge comes from opportunities that are created in the market by institutional managers who trade too much. Their trading is usually based on either an exaggerated view of how well they can predict investment returns, or a misunderstanding of how trading costs increase with size. The strategies of institutional managers can still be perfectly rational despite providing us with opportunities through over-trading, simply because of the huge agency issues in portfolio management.

Incentives versus performance

Investment strategies have fixed capacities. As I increase the money invested in a strategy, my expected transaction costs increase while my pre-transaction cost estimate of expected return stays constant. Figure 1 shows an example of this—once my marginal expected return, net of transaction costs, crosses zero, increased investment in a strategy only loses money.

Unfortunately for most investors who have delegated fiduciary responsibility to investment managers, investor and investment manager incentives are not well aligned. Almost all investment managers are paid a percentage of the funds under management. In the case of mutual funds and most institutional portfolios, this means their primary incentive is managing more funds. Performance does also help determine reward, but only because it helps to keep assets or bring in more.

I’m not saying that investment managers don’t try hard to deliver superior returns, rather that economic incentives move many successful managers to the right side of figure 1. The problem is compounded by the significant errors involved in estimating the shape of the graph: most managers grossly overestimate their ability to forecast asset returns. It’s hard to turn people down



tioned agency issues, but not entirely free of them. Asset management firms usually only collect management fees and trade at 8–12 times earnings; hedge fund managers also collect management fees and can still fetch 2–3 times earnings if they sell their firm.

By contrast, proprietary traders typically only earn a percentage of their trading profits (as in our case). There is no reward for holding more assets than we need, since we are charged financing for all positions. Assuming the bank has good risk management controls in place, incentives are well aligned. Excellent risk management is essential to avoid giving too much value to a trader’s free option (think Barings). Another way of mitigating the value of the option that a pay-out structure presents to the proprietary trader or hedge fund manager is to hold back compensation as a reserve against possible future drawdowns. This further aligns incentives and removes some of the asset substitution problem.

A simple study is suggested: rank investment-manager categories by the risk-adjusted return of all managers in that category. It would be surprising if the underlying economic incentives did not determine the relative investment performance for each category. Proprietary traders have the purest incentive and would do the best, hedge fund managers would come next and institutional portfolio managers last. I believe this is the case even if you include the well-publicized disasters that have occurred over the years in hedge fund management and proprietary trading. Getting accurate data on proprietary trading performance is probably too difficult to make such a study feasible, but the literature on hedge fund performance that tries to estimate the performance of the group as a whole is consistent with these assertions. (See ‘Further reading’ below.)

More simply: if your investment firm has a marketing department, you’re probably not that good an investor. See figure 2.

Creating a high Sharpe Ratio strategy

How does one go about creating a high Sharpe Ratio strategy? Well, you’re not going to get much advice from me—sorry! I will, however, address two issues. First, should one try to build one really great strategy or put together a bunch of good strategies? Second, how do shorter and longer horizon strategies differ?

In Grinold and Kahn’s book on *Active Portfolio Management* (see ‘Further reading’), the authors describe the ‘Fundamental Law of Active Management’: a strategy’s Sharpe Ratio is proportional to the number of independent bets taken by the strategy multiplied by the correlation of those bets with their outcome (figure 3). To get a higher SR, you need to increase the number of your bets or increase the strength of your forecasts.

In my opinion it is far better to refine an individual strategy by increasing both the number of bets within the strategy and the strength of the forecasts made in the strategy, than to attempt to put together lots of weaker strategies. Depth is more important than breadth for investment strategies.

As a strategy develops, betting opportunities increase and returns for each bet increase. But a huge transaction-cost barrier must be overcome before a strategy becomes profitable. Once this is overcome, additional improvements will leverage profit much more efficiently than initial research will. Even though these improvements are harder to come by, the work is worth the effort. I know a proprietary trader who was offered extra compensation by a smart investor if he focused solely on improving his original

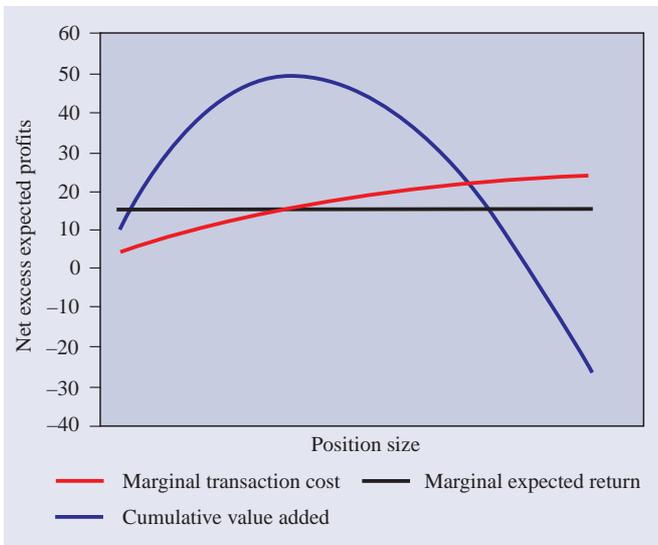


Figure 1. Investment level versus excess profits earned.

who want to invest money with you, while it’s easy to believe that one is further left on the graph in figure 1 than is actually the case. (N.B. As an asset management firm grows, expected return may initially increase as a function of assets under management because additional compensation can be paid to improve one’s investment process.)

Hedge fund managers are incentivized both by a percentage of the profits they make and by how much money they have under management. They are therefore less susceptible to the aforemen-

Sharpe Ratio	Marketing department
≤ 0	Runs the firm
0.25	Very important; involved in all investment decisions; major focus on asset gathering
0.5–1.0	Secondary
1.0–2.0	Almost superfluous
≥ 2.0	What marketing department?

Figure 2. Investment firm politics.

SR: Sharpe Ratio. N : Number of independent bets.
 IC: Information coefficient (correlation of bet with outcome).

$$SR^2 \approx N \cdot \left(\frac{IC^2}{1-IC^2} \right)$$

Figure 3. The fundamental law of active management (Grinold and Kahn 1999).

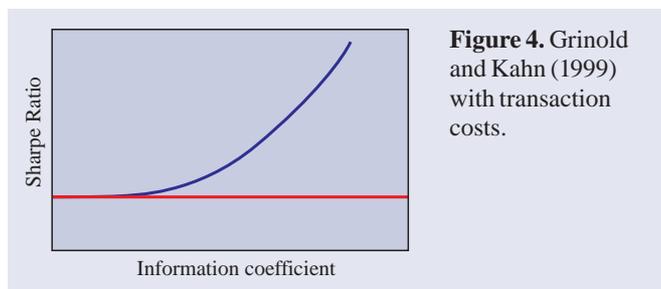


Figure 4. Grinold and Kahn (1999) with transaction costs.

strategy instead of developing a strategy in a different area. He ended up figuring how to increase the returns to his original strategy many times over.

If trading costs are added to the equation in figure 3, the relationship between forecast strength and Sharpe Ratio stops being linear and looks more like figure 4. (Interestingly, the experience of many of the proprietary traders and hedge funds managers I know looks just like that graph, only with ‘effort’ on the x-axis and ‘reward’ on the y-axis.)

I would much rather have a single strategy with an expected Sharpe Ratio of 2 than a strategy that has an expected Sharpe Ratio of 2.5 formed by putting together five supposedly uncorrelated strategies each with an expected Sharpe Ratio of 1. In the latter case you’re faced with the risk that the strategies are more correlated than you realize (think Long Term Capital). There is also the increased effort of ascertaining whether each individual strategy really has a Sharpe Ratio of 1.

Of course, choosing where to dig is important. There are well-established model-driven groups doing convertible arbitrage, mortgage-backed arbitrage, futures trading, long–short equity statistical arbitrage and option arbitrage. Some use more than one of these strategies, but to make significant reliable profits in these areas, you need to put in enough work to become one of the best players in each one.

Short versus long horizon strategies

An important choice for many proprietary traders is whether to focus on shorter or longer horizon strategies. Typically, shorter horizon strategies get their edge from providing temporal liquidity to a market place or predicting short-term trends that arise from

inefficient trading. Longer-term models focus on asset pricing inefficiencies. How does implementation of these strategies compare?

Shorter-horizon investment strategies are desirable because they tend to create higher Sharpe ratios. If your average holding period is a day or a month, you have the opportunity of placing many more bets than if you hold positions for three months to a year or longer. On the flip side, shorter horizon strategies tend to have capacity issues (it’s easy to make a small amount of money with them, harder to make a lot of money). Shorter horizon strategies also require serious investments in trading infrastructure, since quick and inexpensive execution is much more important than for longer horizon strategies.

Risk management for shorter horizon strategies tends to occur through position trading rather than portfolio construction. Assets are not held for long periods of time and portfolio characteristics change quickly. The biggest risk for shorter horizon strategies is model risk, or the risk that the trading strategy deployed has stopped working. Since even the best trading strategies experience periodic drawdowns, the hardest challenge for the short-term model-based trader is to figure out whether his model is going through a regular drawdown or has stopped working altogether.

Longer-horizon model-driven investment strategies have different issues. Since assets are held for longer periods of time, execution costs (although still important) are not the primary focus. Statistical inference becomes more difficult and the danger of overfitting or mining data becomes larger. Risk management for longer-term strategies happens in portfolio construction: since rebalancing occurs less frequently, more care needs to be taken to ensure the portfolio is not exposed to unintended sources of risk. Because they tend to have lower Sharpe ratios, longer horizon strategies have a different kind of capacity issue—the manager’s capacity for pain. However, there is one advantage: because trading occurs less frequently it’s possible to lead a much better lifestyle than if you’re running shorter horizon strategies!

Conclusion

My aim in this article was to be informative (and occasionally entertaining) while not telling you anything that my competitors or potential competitors would find useful. Unfortunately, the mere knowledge that it is possible to beat the market consistently may increase competition and make our type of trading more difficult. So why did I write this article? Well, one of the editors is a friend of mine and asked nicely. Plus, chances are you won’t believe everything I’m telling you. And if you do, well, I’ve always liked a challenge.

Further reading

Grinold R C and Kahn R N 1999 *Active Portfolio Management: A Quantitative Approach for Producing Superior Returns and Controlling Risk* (Probus)
 Goetzmann W N, Ibbotson R G and Brown S J Offshore hedge funds: survival and performance 1989–1995 *J. Business* 72 91–117
 Why hedge funds make sense *Global Equity and Derivative Markets* (Morgan Stanley Dean Witter) November 2000

Peter Muller is Managing Director, Head of Process Driven Trading, at Morgan Stanley Dean Witter, New York.