

Portfolio Cramér-Rao bounds

(why bad things happen to good quants)

Steven E. Pav
steven@cerebellumcapital.com

Cerebellum Capital

May 29, 2015

Why does this happen?



(courtesy Automated Trader magazine)

Culprits

The usual suspects:

- Plain old overfit (“p-value hacking”).
- Broken backtests: lookahead bias, survivorship bias, *etc.*
- Bad understanding of trade costs.
- Bad execution.

But also perhaps:

- Bad luck?
- A fundamental bound?

The setup

Consider portfolio construction technique $\hat{\nu}(\cdot)$:

- Takes historical data, X , a $T \times p$ matrix.
- Produces a portfolio, $\hat{\nu} = \hat{\nu}(X)$, a p vector.

The (population) Sharpe ratio is a function of X :

$$q(\hat{\nu}(X)) = \frac{\text{expected return of } \hat{\nu}(X)}{\text{volatility of } \hat{\nu}(X)}.$$

The bound

Then

Cramér-Rao + geometry \Rightarrow bound on expected Sharpe ratio

Expectation is over draws of the historical data, X .

Under a wide range of conditions the bound is:

$$E_X [q(\hat{\nu}(X))] \leq \sqrt{\frac{\text{effect size}}{\# \text{ knobs} + \text{effect size}}} \text{maximal Sharpe ratio,}$$

with effect size = # years (maximal Sharpe ratio)².

Control degrees of freedom

This is why portfolio optimization is not performed on 100's of unknowns:
If maximal Sharpe is 1.1, observing 5yr of data:

- for 10 stocks, the bound is 0.7.
- for 40 stocks, the bound is 0.4.
- for 160 stocks, the bound is 0.21.

However, maximal Sharpe should grow with universe size.
Can it grow fast enough?

Diversification and universe size




Figure : Bound vs. p for $\zeta_* = \zeta_0 p^\gamma$, with γ between 0.15 and 0.35. Maximal Sharpe must grow at rate faster than $\gamma = 1/4$, otherwise bound will decrease.

Diversification in S&P 100 (there is none)




Figure : Estimated maximal Sharpe versus p for the S&P 100 Index names, over 1000 permutations of stock order.

What's next?

- Read the paper:
<http://arxiv.org/abs/1409.5936>
(Many interesting theoretical questions remain.)
- Browse code on github:
<http://github.com/shabbychef/rfin2015>
- Try the shiny app:

```
library(shiny)
runGist("75e1066b3db6b0f9e981")
```

- Ask me questions.
- Thank you.

Common Questions I

The proof assumes *normal* returns?

Indeed, however I suspect that a stronger upper bound holds for the case of more fat-tailed distributions, though I do not have a proof yet.

What about the $1/n$ allocation?

Like all Cramér-Rao results, this one requires a 'stopped clock' proviso (a stopped clock is correct twice a day), which precludes us from considering the $1/n$ allocation. The $1/n$ allocation is utterly useless in some situations, e.g., a long/short portfolio.

Common Questions II

What kinds of portfolio estimators *are* covered?

The technical condition is that the expected Sharpe of the estimator must be only a function of maximal Sharpe. This is implied by *directional equivariance*: the estimator is not 'fooled' by an invertible transform of the returns vector.

Nobody uses simple returns estimates!

There is a similar result for the case of linearly conditional expectation, hedged portfolios, and subspace constraints. Probably there is a much more general result.

Common Questions III

This bound uses an unknown population parameter. Can you do better?

Not at the moment. This is a particularly interesting question: how to construct confidence intervals on the Sharpe of the Markowitz portfolio. It is different than the typical statistical analysis, which performs inference on the maximal population Sharpe ratio.

What does this bound say about *my* portfolio?

Very little. It only gives a bound on the *expected* Sharpe based on repeated draws of the historical data. You only got one draw of that historical data.

Common Questions IV

I don't like this result: it seems I can be harmed by performing more backtests.

It is commonly stipulated that a perfectly rational agent cannot be harmed by the addition of new information or new optional courses of action. Barring the fact that humans are not perfectly rational agents, a quantitative trading scheme that can only improve with the addition of new information sounds like the Holy Grail. Like the Holy Grail, it is unlikely to exist.

Have you quantified 'overfitting'?

Perhaps. One sense of overfitting is having a biased estimate of future performance based on the data used to construct the model. Perhaps this bound could be used to de-bias performance estimates, but it feels like there is an infinite regress.

Common Questions V

This Cramér-Rao bound feels very Frequentist.
I suppose it does.

How did you estimate γ ?
I used the 'KRS' estimator from the [SharpeR](#) package.