www.joim.com

Journal Of Investment Management, Vol. 11, No. 1, (2013), pp. 1–5

PRACTITIONER'S DIGEST

The "Practitioner's Digest" emphasizes the practical significance of manuscripts featured in the "Insights" and "Articles" sections of the journal. Readers who are interested in extracting the practical value of an article, or who are simply looking for a summary, may look to this section.

DECONSTRUCTING BLACK–LITTERMAN: HOW TO GET THE PORTFOLIO YOU ALREADY KNEW YOU WANTED

PAGE 6

Richard O. Michaud, David N. Esch and Robert O. Michaud

Since the publication of their original article in 1992, Black-Litterman (BL) has become a popular method in practical finance for creating superficially stable portfolios, adjusted to investor views. A popular perception is that BL can solve the instability problems of portfolios on Markowitz efficient frontiers. In fact, the instability issues of Markowitz portfolios are caused by estimation error (Michaud 1998, 2008), which BL does nothing to explicitly handle. The BL method assumes a perfectly known market portfolio in a state of undisturbed equilibrium, a perfectly known covariance matrix, and correct investor views numerically calibrated to perfectly quantify the exogenous knowledge of the investor. On top of these heroic assumptions, the BL formula itself is built on faulty statistical theory and is not optimal in any mathematical sense. Besides, since it is equivalent to a maximum Sharpe ratio Markowitz optimization with specific inputs, it inherits all of the instability of Markowitz optimization, especially when the frontier is extended beyond the BL portfolio.

Black and Litterman (1992) give a tuning parameter τ to adjust the strength of the views. This parameter may be fixed or adjusted, and is in practice often used to guarantee investable portfolios. Adjusting τ for investability amounts to either adjusting the data to fit the desired solution or adjusting one's "exogenous" views, and is a violation of fundamental principles of statistical analysis. Like the unadjusted BL portfolio, the τ-adjusted portfolio can also be found on a Markowitz frontier with particular inputs and inherits the properties and shortcomings of that method.

In our article, we provide a simple but detailed example of a realistic Black-Litterman analysis and show the corresponding Markowitz inputs and frontiers which contain the BL portfolios. Moving away from the BL portfolios at their maximum Sharpe ratio points, these frontiers veer quickly into uninvestable portfolios with short and/or leveraged positions in some assets and are not useful to managers who require access to multiple risk profiles tailored to investors' risk preferences. The BL portfolios and frontiers in our example are compared with better solutions created with methods that explicitly account for estimation error. Michaud efficient portfolios are better diversified and more intuitive, have superior out-of-sample performance by design, and do not rely on false assumptions or dial in a preordained result.

Users of Black-Litterman or its implied returns should be mindful of these methods' limitations. BL does not solve but rather conceals the instability and estimation error problems of Markowitz mean-variance optimization. Because it is not a proper optimization method and tends to assign too much confidence to personal views it may often miss useful information while exposing investors to unnecessary risk. The simplicity and apparent adequacy of the procedure comes at the peril of ignoring better statistically-based methods that merge all of the available information into a more effective portfolio creation process.

The first practical similarity test for mean-variance optimality is the Michaud rebalancing rule. However, the original procedure ignored an often important consideration in that much of the information used to construct the current portfolio may be implicitly or explicitly included in the new optimal. This partial input match results in an overly conservative Michaud rebalance signal. We develop new algorithms that address overlapping data in the Michaud test. The new distribution defines a critical range for the Michaud rule and extends its applicability and power. We describe two procedures and give examples. The method allows large-scale automatable non-calendar based portfolio monitoring. The procedure is generalizable as a statistical similarity rule for quadratic programming contexts with potential applications well beyond portfolio management.

INVESTING IN WHAT YOU KNOW: THE CASE OF INDIVIDUAL INVESTORS AND LOCAL STOCKS PAGE 21

Mark S. Seasholes and Ning Zhu

This paper empirically evaluates the returns to investing locally. Specifically, we study the investment decisions of a large number of individuals who invest through a discount broker and test whether individuals generate positive alpha when investing in local companies.

We evaluate individuals' performance in local vs. non-local stocks, by studying their both their portfolio holdings and their transactions. Using either portfolios or holdings generates consistent results. That is, investing in local stocks *does not* help individual investors earn excess returns. Investing locally can, at best, alleviate the negative excess returns that individual investors tend to earn (on average.)

Our study points out that some investment folklore may be unfounded and may mislead investors. Understanding the limitation of such folklore can help households make better finance-related decisions.

STOCK STRATEGIES WITH THE JANUARY BAROMETER AND THE YIELD CURVE PAGE 32

Licheng Sun, Chris Stivers and Ajay Kongera

We study annual stock-market strategies with signals that are based on the January Barometer and the slope of the yield curve. The January Barometer reflects that February-to-December stock returns have historically been higher following positive January stock returns. Cooper et al. (2010) show that the best way to use the January Barometer is to be long the stock market following positive Januarys and invest in T-bills following negative Januarys.

In this study, over 1954 to 2010, we similarly consider an annual signal based on the sign of the slope of the yield curve, which we believe has better underlying economic justification than the January Barometer. We find that the February-to-December average stock return following upward-sloping yield curves is significantly higher than the comparable 11-month average stock return following downward-sloping yield curves.

Motivated by this finding, we further consider strategies that combine the trading signals from the January Barometer and the yield curve. We find that a combined positive signal, with both a positive January return and an upward-sloping yield curve, has historically performed better than either stand-alone indicator, both in terms of higher average returns and lower risk. We show that the best strategy is to be long the market when both the January Barometer and the yield curve give us a buy signal, and stay invested in T-bills otherwise.

VARGAMMA: A UNIFIED MEASURE OF PORTFOLIO RISK PAGE 50

Kent Osband

Mean-variance analysis has long been the mainstay of portfolio optimization. Yet it copes poorly with unstable or fat-tailed risks. The usual responses supplement mean-variance analysis with stress tests or estimated Value-at-Risk percentiles.

These responses rarely hit the spot. Stress tests tend to be ad hoc: how much stress with what probability? Value-at-Risk encourages sloppiness in both estimation and risk mitigation, since the magnitude of shortfall is largely irrelevant. Even when forecasts of tail risks are accurate, it is unclear how to integrate them with mean-variance analysis.

VarGamma offers a unified measure. Formally, it estimates the expected risk-adjusted return when risk aversion is approximately constant and uncertainty is modeled as conditionally normal with gamma-distributed variances. VarGamma is much easier to understand than other risk measures and much less prone to being gamed. The distributional assumptions allow VarGamma to capture non-zero skewness, excess kurtosis, and slow-decaying tails. Stress tests are incorporated as independent VarGamma shocks, with readily quantified impact on risk-adjusted returns.

Despite the sophistication, VarGamma is tractable. Handy formulas relate VarGamma to portfolio weights. They are relatively easy to optimize, without requiring complicated integrations or Monte

Carlo sampling. Indeed, the formulas are as simple as they can be, given non-Gaussian risk. VarGamma also readily incorporates forward-looking estimates of risk from the options market.

VarGamma estimated under high risk aversion is a superior alternative to Value-at-Risk. It is more reliable and discourages wasteful regulatory arbitrage. Its predictions of near-exponential decay in outer tail risks (far slower than normality suggests) sits well with crisis theory and practice. Even in forecasting expected percentile losses, the VarGamma approach tends to outperform empirical Value-at-Risk.

PUT OPTION EXERCISE AND SHORT STOCK INTEREST ARBITRAGE PAGE 66

Kathryn Barraclough and Robert E. Whaley

U.S. exchange-traded stock options are exercisable before expiration. While put options should frequently be exercised early to earn interest, they are not. In this paper, we present the key results from our longer and more comprehensive study, Barraclough and Whaley (2012). We explain a decision rule for exercising a put option early and then examine actual early exercise behavior. Using a sample of put options on stocks over the period January 1996 through September 2008, we find that more than 3.96 million puts that should have been exercised early remain unexercised, representing over 3.7% of all outstanding puts. We also find that failure to exercise cost put option holders \$1.9 billion in forgone interest income over this period.

The failure of put option holders to exercise when they should has given rise to a trading game dubbed the short stock interest arbitrage strategy. We describe this trading strategy and show how it is used by market makers and proprietary firms to earn a dominant share of the interest income forfeit by put option holders.

ERRATUM

The Practitioner's Digest of the paper, "Portfolio Monitoring in Theory and Practice by Richard O. Michaud, David N. Esch and Robert O. Michaud", which was published in the Fourth Quarter 2012 of Journal of Investment Management (JOIM) should read as

PORTFOLIO MONITORING IN THEORY AND PRACTICE PAGE 5

Richard O. Michaud, David N. Esch and Robert O. Michaud

The when-to-trade decision is a critical yet neglected component of modern asset management. The need-to-trade decision is typically based on suboptimal heuristic rules. Until relatively recently no effective decision rule has existed for deciding whether a currently held portfolio has aged sufficiently to make trading desirable. The fundamental problem is that portfolio managers and financial theoreticians persist in ignoring the statistical nature of asset management and the impact of estimation error on effective decision making. The rebalancing decision is necessarily a statistical similarity test between the current drifted portfolio and a proposed new optimal. Many rebalancing rules in use have had little theoretical or statistical foundation and often lead to trading in noise or ineffectively using useful information. While statistical similarity tests are available in the financial literature none treat the real world portfolio management problem that requires inequality constraints, targeted risk portfolios, trading costs, and asset manager style customization. Inclusion of these necessary features requires compute-intensive resampling methods that overcome the limitations of familiar null distributions. The first practical similarity test for mean-variance optimality is the Michaud rebalancing rule. However, the original procedure ignored an often important consideration in that much of the information used to construct the current portfolio may be implicitly or explicitly included in the new optimal. This partial input match results in an overly conservative Michaud rebalance signal. We develop new algorithms that address overlapping data in the Michaud test. The new distribution defines a critical range for the Michaud rule and extends its applicability and power. We describe two procedures and give examples. The method allows large-scale automatable non-calendar based portfolio monitoring. The procedure is generalizable as a statistical similarity rule for quadratic programming contexts with potential applications well beyond portfolio management.