# Black Ice: Low-Volatility Investing in Theory and Practice

October 22, 2015

by Feifei Li, Engin Kose
of Research Affiliates

**Key Points**

1. Implementing a low-volatility strategy entails mitigating out-of-sample estimation errors; over-concentration in sectors, regions, and names; and high transaction costs.

2. Representative constraints succeed in making simulated minimum-variance portfolios more investable but push them in the direction of the cap-weighted benchmark.

3. Constraints that are similarly designed to improve the investability of heuristically constructed low-volatility portfolios tend to preserve the intended portfolio characteristics.

Equity investors have endured two extreme market downturns since the turn of the century. The broad U.S. market, represented by the S&P 500 Index, fell by 44% in the aftermath of the dot-com bubble and 51% in the great recession. These devastating experiences reawakened institutional and individual investors to the downside of market volatility and, for a while, prompted great interest in low-volatility investing. Over the last six years, however, the market has been climbing; at the end of July 2015, the price level of the S&P 500 was over 200% higher than its trough in March 2009.[1] Low-volatility strategies have languished, and many investors appear to be sleepwalking again—possibly toward a cliff.

While human nature conditions us to chase whatever has been working best—a strategy that we know will backfire badly for the long-term investor—we also know that inertia generally doesn't pay off. Given the immense gains of this bull market, it may be timely to take some profits off the table, and to dampen our overall portfolio risk through exposure to the well-documented low-volatility effect.[2] But, like most things that sound inviting, not all low-volatility portfolio strategies are equally attractive. It pays to understand the differences. Let's focus first on issues surrounding the implementation of minimum-variance strategies. The same challenges arise for heuristic low-volatility portfolio construction; we consider their impact below.

**The Need for Constraints**

There are essentially two approaches to low-volatility investing. One of them, called minimum-variance investing, is based on quantitative optimization techniques,[3] while the other employs heuristic portfolio construction rules. Some products use combinations of the two approaches, but for this purpose, we will focus on the two primary approaches.

- The minimum-variance portfolio approach uses a numerical optimizer to select a set of non-negative stock weights such that the resulting predicted portfolio volatility is minimized.

- A heuristic approach to low-volatility investing typically uses a common risk measure (e.g., beta or volatility) to screen out volatile companies, and assigns weights to the remaining securities by their market capitalizations or the inverse of the company-specific risk measure.

Solidly grounded in finance theory, the minimum-variance method is clearly a sound approach to constructing a low-volatility portfolio. Nonetheless, implementing this method may be more problematic than many investors realize, and the chosen solutions unavoidably affect investment results.[4] The challenges relate to "implementation shortfall," including disappointing out-of-sample performance due to estimation errors,[5] extreme and unstable portfolio characteristics, and high transaction costs.[6]

In addition to applying advanced statistical techniques,[7] asset managers and index providers often mitigate estimation errors—and address other minimum-variance implementation issues—by imposing constraints on the optimization process. They typically apply minimum and maximum weight constraints to avoid over-concentration in individual stocks; sector and regional weight constraints to forestall excessive allocations to any one industry group or geographical area; and turnover constraints to control trading costs.

These restrictions are successful in fixing the identified problems, and as a result, they make minimum-variance portfolios more investable. But the improvements come at a price. The constraints progressively nudge the portfolio closer to the market-cap-weighted index and, more importantly, introduce a link between the price of a stock and its weight in our portfolio. As we (and others) have demonstrated, the link between stock price and the portfolio weight has a cost; indeed, severing that link is the main source of alpha for fundamentally weighted and other non-cap-weighted strategies. As a practical matter, it appears that optimization-based minimum-variance strategies cannot be implemented without meaningful slippage.

**Empirical Study**

To evaluate the impact of typical constraints, we constructed three hypothetical long-only minimum variance portfolios[8] from the 1,000 stocks with the highest market capitalization in our universe: a U.S. portfolio, a developed markets portfolio, and an emerging markets portfolio. The baseline minimum-variance portfolios, which were rebalanced annually over the simulation periods, incorporated minimum and maximum weight constraints on individual stock positions. Then we serially applied a capacity constraint related to the stocks' weights in the market-cap-weighted benchmark; sector and regional concentration constraints; and a ceiling on one-way turnover. (See the Appendix for details on the constraints and regional makeup.)

In **Table 1**, we see that the stepwise imposition of constraints decreases turnover, increases weighted-average market capitalization (WAMC), increases the effective number of stocks,[9] and decreases the aggregate weight of the top 10 names. Just as intended, the constraints limit trading and give the minimum-variance portfolios greater liquidity, higher capacity, and lower concentration.

In Panel A of **Table 2**, we see how performance drops, risk rises, and the Sharpe ratio falters, as we apply more constraints to the simulated U.S. portfolio. Interestingly, the capacity constraint helps performance in the hypothetical developed markets (Panel B) and emerging markets (Panel C) portfolios. In all markets, tracking error against the cap-weighted benchmark decreases monotonically with each new constraint. By partially reversing the optimization, the added constraints move the portfolios away from the theoretical minimum-variance baseline toward the cap-weighted benchmark.

The effect of constraints on the ratios of excess return to volatility and value added to tracking error can be seen in **Figure 1**. Taken together, the constraints push the U.S. minimum-variance portfolio in the direction of the cap-weighted benchmark.

We also observe that the U.S. minimum-variance portfolio's sector allocation more closely resembles that of the cap-weighted benchmark when all constraints are in effect. **Figures 2a–2c** display simulated three-month smoothed sector weights using Kenneth French's 12-industry classification. In the baseline case, shown in Figure 2a, the utilities sector has a very large allocation over most of the measurement period. The fully constrained portfolio (Figure 2b) has a more balanced allocation to economic sectors, much like the cap-weighted benchmark (Figure 2c).

So far, we have studied the optimization-based approach to low-volatility investing. We confirm that the optimization process must be constrained to assure the minimum-variance portfolio is implementable. These constraints are also necessary to obtain reasonable portfolio characteristics such as diversification and capacity. But they have a cost. The portfolio becomes more like the market, and the risk increases, with mixed effects on risk-adjusted performance over the simulation periods. Let's now turn to the heuristic approach to low-volatility investing.

**The Heuristic Approach**

We conducted a similar analysis of a heuristic approach to low-volatility portfolio construction. To construct the simulated baseline heuristic portfolios, we selected the 200 stocks with the lowest volatility from fundamentally weighted indices for the U.S., developed, and emerging markets. To construct region- and sector-constrained portfolios, we selected from the fundamentally weighted indices' constituents the 20% of stocks with the lowest volatility within each region and sector, thereby conserving the original allocations. Finally, to incorporate a turnover constraint, we limited trading to removing stocks whose volatility moves outside a pre-established band and adding previously ineligible stocks whose volatility now falls within the band. This approach to turnover control suits heuristically constructed portfolios better than the explicit turnover constraints used in minimum-variance portfolios. Performance statistics for the baseline and constrained low-volatility portfolios are presented in **Table 3**. (We showed the same measures for the simulated minimum-variance portfolios in Table 2.) In the United States, the minimum-variance and heuristic low-volatility portfolios have roughly comparable absolute and risk-adjusted returns. In the developed markets, the heuristic strategy has higher absolute returns and higher Sharpe ratios; in the emerging markets, the minimum-variance approach has lower absolute returns but higher Sharpe ratios. Neither approach prevails in all regions.

The heuristic approach is, however, significantly superior in terms of transaction costs and valuation features. In**Table 4**, we see that, across regions, the baseline and constrained heuristic portfolios have substantially higher weighted-average market cap, lower price multiples, and higher dividend yields. (Table 1 displayed the same measures for the minimum-variance portfolios.) In addition, the heuristically constructed portfolios have lower turnover in the U.S. and developed markets. These characteristics make the heuristic approach cheaper in terms of fundamental valuations and, outside the emerging markets, more efficient in terms of trading activity.

**In Closing**

As the study summarized here demonstrates, constraints like those that index providers typically introduce in the optimization and portfolio construction process succeed in making minimum-variance portfolios more investable by improving liquidity, avoiding extreme allocations, and controlling transaction costs. All the same, there are side effects. In general, the constraints tend to make minimum-variance portfolios look a little more like cap-weighted indices. In so doing, the constraints increase portfolio volatility, compromising a key feature (and rendering the term "minimum variance" technically inaccurate). In comparison, constraints similarly designed to improve the investability of heuristically constructed low-volatility portfolios tend to preserve the intended portfolio characteristics. When evaluating smart beta alternatives, it clearly pays to understand the trade-offs that come into play in the transition from theory to practice.

**Endnotes**

[1] The S&P 500 Index closing price level was 676.53 on March 9, 2009, and 2103.84 on July 31, 2015, a change of 211%.

[2] See Chow, Hsu, Kuo, and Li (2014); Soe (2012); Blitz, Pang, and van Vliet (2012).

[3] The minimum-variance method is offered by several influential market providers, such as MSCI.

[4] See Behr, Guettler, and Miebs (2008).

[5] See Jagannathan and Ma (2003); Kempf and Memmel (2003); AGIC Systematic Investment Team (2012).

[6] See Chow, Hsu, Kuo, and Li (2014), and Arnott (2006).

[7] Methods available to mitigate the estimation errors inherent in sample covariance matrices include the Sharpe (1964) factor-based approach, the Elton and Gruber (1973) constant correlation approach, and the Ledoit and Wolf (2004) statistical shrinkage approach.

[8] In brief, we employed an optimization routine to find a numerical solution of portfolio weights that minimizes portfolio variance under constraints. To ensure that the covariance structure inputs were positive definite, we applied principal component analysis to the covariance matrix, which was estimated using up to five years of monthly excess returns.

[9] See the Appendix for the mathematical definition of effective N (here, the effective number of stocks).

**References**

AGIC Systematic Investment Team. 2012 “Specification of Constraints in Managed Volatility Strategies.”*Allianz Global Investors Capital*. (September) Available at http://www.allianzgic.com/en/Documents/Constraints-In-Managed-Volatility_FINAL2.pdf.

Arnott, Robert D. 2006. “Implementation Shortfall.” Editor’s Corner,*Financial Analysts Journal*, vol. 62, no. 3 (May/June):6–8.

Behr, Patrick, Andre Guettler, and Felix Miebs. 2008. “Is Minimum-Variance Investing Really Worth the While?An Analysis with Robust Performance Inference.”

Blitz, David, Juan Pang, and Pim van Vliet. 2012. “The Volatility Effect in Emerging Markets.” Robeco Research Paper (March).

Chow, Tzee-Man, Jason C. Hsu, Li-Lan Kuo, and Feifei Li. 2014. “A Study of Low Volatility Portfolio Construction Methods.”*Journal of Portfolio Management*, vol. 40, no. 4 (Summer):89–105.

Elton, Edwin J., and Martin J. Gruber. 1973. “Estimating the Dependence Structure of Share Prices—Implications for Portfolio Selection.” *Journal of Finance*, vol. 8, no. 5 (December):1203–1232.

Jagannathan, Ravi, and Tongshu Ma. 2003. “Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps.” *Journal of Finance*, vol. 58, no. 4 (August):1651–1684.

Kempf, Alexander, and Christoph Memmel. 2003. “On the Estimation of the Global Minimum Variance Portfolio”

Ledoit, Olivier, and Michael Wolf. 2004. “Honey, I Shrunk the Sample Covariance Matrix.”*Journal of Portfolio Management*, vol. 30, no. 4 (Summer):110–119.

Sharpe, William F. 1964. “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk.”*Journal of Finance*, vol. 19, no. 3 (September):425–442.

Soe, Aye M. 2012. “The Low Volatility Effect: A Comprehensive Look.” (August 1).

**Appendix**

A. PORTFOLIO CONSTRAINTS

1. Minimum weight constraint. Weights smaller than 0.05% are forced to zero.

2. Maximum weight constraint. Individual stock weights are capped at 5%.

3. Capacity constraint. The weight of a stock is capped at the lower of 1.5% or 20 times its weight in the corresponding cap-weighted portfolio. Note that this constraint dominates the maximum weight constraint.

4. Sector concentration constraint. Sector weights are not allowed to deviate more than ±5% from the corresponding cap-weighted sector weights.

5. Region concentration constraint. If the cap-weighted region weights are less than 2.5%, the minimum-variance region weights are capped at three times their weight in the cap-weighted portfolio. Otherwise, they are not allowed to deviate more than ±5% from the corresponding cap-weighted region weights.

 

6. Turnover constraint. The maximum allowable one-way index turnover is 20%.

B. MARKET AND REGION DEFINITIONS

*Developed Markets*

Region 1 = DevEME, which includes Austria, Belgium, Denmark, Finland, Greece, Ireland, Israel, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, and Switzerland

Region 2 = DevAPAC, which includes Australia, Hong Kong, New Zealand, and Singapore

Region 3 = France

Region 4 = Germany

Region 5 = United Kingdom

Region 6 = Japan

Region 7 = Canada

Region 8 = United States

*Emerging Markets*

Region 1 = EMEMEA, which includes Czech Republic, Egypt, Hungary, Morocco, Poland, and Turkey

Region 2 = EMAPAC, which includes Indonesia, Malaysia, Philippines, and Thailand

Region 3 = EMAME, which includes Chile, Colombia, Mexico, and Peru

Region 4 = South Africa

Region 5 = Russian Federation

Region 6 = India

Region 7 = China

Region 8 = Taiwan

Region 9 = South Korea

Region 10 = Brazil

C. EFFECTIVE NUMBER OF STOCKS

This is the reciprocal of the Herfindahl ratio, which was developed to gauge monopoly concentration in industry, repurposed for investment management. Hypothetically a portfolio of 100% weight in 1 stock has an Effective N of 1; a portfolio of equal weight to 1,000 stocks has an Effective N of 1,000. In another words, these minimum variance portfolios are as diversified as equally weighting only 30–40 stocks.