The Devil Is in the Details: *The Risks Often Ignored in Low-Volatility Investing*

NICHOLAS ALONSO AND OLEG NUSINZON

NICHOLAS ALONSO
is a director of multi-asset strategies at PanAgora Asset Management in Boston, MA.
nalonso@panagora.com

OLEG NUSINZON
is an independent consultant based in Boston, MA.
onusinzo@chicagobooth.edu

KEY FINDINGS

- Using an optimizer to gain exposure to the low-volatility anomaly/premiums subjects the portfolio construction process to input sensitivity, which can lead to wide variation in results.
- Input sensitivity affects not only the portfolio's sector and stock positioning but also its risk exposures. This variation in risk exposures can lead to very different outcomes for investors.
- There are more robust ways to gain exposure to the low-volatility anomaly/premiums, and the authors show that a risk-balanced approach not only captures the anomaly/premiums well but does so under different risk model assumptions.

ABSTRACT: *With increasing investor interest in low-volatility equity strategies comes a need for greater scrutiny of different methodologies used to achieve low-volatility exposure. In an earlier article, the authors investigated the analytical differences between a variety of approaches to constructing low-volatility portfolios. In this article, the authors turn their attention to the empirical differences between common approaches to low volatility. They find that a traditional optimizer-based approach to building low-volatility portfolios has large sensitivities to the risk inputs used in the process. In fact, using the same portfolio construction methodology but changing the risk inputs even slightly can lead to large differences. The magnitude of this sensitivity should give investors pause; even across risk inputs in which differences are valid, variations persist and can be harmful to portfolio performance. The authors show that there are other, more robust, ways of achieving low-volatility portfolios without this input sensitivity (e.g., risk balancing) and suggest that investors should consider this lack of input sensitivity as a valuable characteristic in low-volatility investing.*

TOPICS: ***Volatility measures, exchanges/markets/clearinghouses, risk management****

*All articles are now categorized by topics and subtopics. **View at PM-Research.com.**

Investors are always interested in getting more return for less risk, or, as stated in a more modern colloquialism, a bigger bang for their buck. As they seek to increase their reward/risk ratio, they consider many different types of strategies above and beyond their passive alternatives. Often, these investors look to the empirical record for guidance.

Theoretical arguments aside, one class of investment strategy shown to have higher reward/risk ratios than many other strategies is low volatility. Using this approach, high reward/risk ratios (higher than those of passive alternatives) are achieved by favoring stocks with lower realized risk or stocks that lower the portfolio's total volatility. The improved reward/risk ratios of this strategy class are largely attributable to the low-volatility anomaly written about in many academic and practitioner studies (Haugen and Baker 1991; Chan, Karceski, and Lakonishok 1999; Jagannathan and Ma 2003; Clarke, De Silva, and Thorley 2006; Baker, Bradley, and Wurgler 2011).[1] Relative to the capitalization-weighted passive alternative, the premium associated with the low-volatility anomaly is driven by an empirically observed asymmetry between this strategy's upside capture and downside protection, as discussed further in the following.[2] Generally, the more upside capture and downside protection a strategy has, the higher that strategy's reward/risk ratio is over time. When compared to a cap-weighted alternative, any persistence in higher reward/risk ratios suggests that investors can do better than investing in the capitalization-weighted passive benchmarks. This empirical observation has likely led to the popularity of low-volatility investing. However, simply seeking higher reward ratios is not the only driver of the increased popularity of this type of investing.

Another more tactical reason for low-volatility investing is an effort to de-risk overall portfolios. This can happen because of a timing decision related to the current market environment or the increased funded status of some liability-driven investment plans such as pension funds. Although improved performance and better liability matching are two good reasons to consider low-volatility investing, there is more to consider. With increased interest in these strategies, there comes a need for an increased understanding of what investors' expectations should be when investing in them.

There are many different approaches to low-volatility investing. Each method can lead to different outcomes in both performance and portfolio characteristics. Maybe less obvious is that different outcomes can occur within each strategy simply because the strategy is sensitive to the inputs used in portfolio construction. In this article, we do not address a detailed discussion of the merits of different approaches, as other research papers have done. Instead, we focus on the sensitivity each strategy has to its inputs. Input sensitivity is important not only because it drives many of the differences in outcomes but also because it is, in our opinion, the least addressed aspect of low-volatility investing, even though there is a long history of academic research on the subject (Michaud 1989). This general lack of attention to input sensitivity can blindside investors and expose them to potentially undesirable outcomes. It is important to note that variability in outcomes is not necessarily attributable to the superiority of one input over another. Instead, it is attributable to how those inputs are used in the portfolio construction process. In this context, we consider input sensitivity to be the devil in the details of low-volatility portfolio construction.

In this article, we discuss many aspects of low-volatility investing. First, we discuss how low-volatility investing differs from other types of factor investing. Second, we identify three popular portfolio construction approaches to achieving low-volatility exposure: inverse volatility, minimum variance, and risk balanced. We also discuss a low-beta-strategy, maximum diversification. Next, we generate backtests for each portfolio construction methodology using three risk models. Differences in performance and portfolio characteristics across the three risk models for each of the four strategies are then analyzed. Given that the portfolio construction methodology for a specific strategy is kept unchanged, the only driver of differences must be attributable to those across the three risk models. Finally, we offer some explanations for the differences seen and conclude with our comments.

WHAT MAKES LOW VOLATILITY DIFFERENT?

Unlike other traditional factor-based strategies, low volatility is typically measured as a portfolio characteristic, rather than simply a stock characteristic. For example, in a typical factor strategy such as value,

[1] As early as 1972, Robert Haugen and A. James Heins produced a working paper titled "On the Evidence Supporting the Existence of Risk Premiums in the Capital Market." Studying the period from 1926 to 1971, they concluded that "over the long run stock portfolios with lesser variance in monthly returns have experienced greater average returns than their 'riskier' counterparts" (Haugen and Heins 1972).

[2] Upside capture is the percentage of a benchmark's return a strategy captures when that benchmark has a positive return. Conversely, the downside capture (the inverse of downside protection) is defined as the percentage of a benchmark's return a strategy captures when that benchmark has a negative return.

in which portfolio managers favor cheap stocks over expensive ones (all else equal), the portfolio manager will calculate a value score independently for each stock. This results in a portfolio value score that is a simple weighted average of each stock's value score. In mathematical jargon, the portfolio value score is a simple linear transformation of the stock value scores. This means that only a stock's value score and its weight in the portfolio contribute to the portfolio's value score; any interaction among the stocks has no effect.

For low-volatility strategies, the interaction among stocks in the portfolio is an important consideration because the portfolio's expected volatility is a combination of both the individual stock volatility and the impact of the interaction among stocks in the portfolio. This is because stocks with low correlations, or interaction effects, also lower a portfolio's volatility independent of the impact of the stock volatility. Because the targeted portfolio characteristic is dependent on both the stock's volatility and the interaction effect among stocks, the portfolio's volatility score is not a simple transformation of the stocks' volatilities. This added complexity has two effects. First, greater complexity leads to more potential solutions, which is why there are many different ways to achieve a portfolio with low volatility. Second, with greater complexity comes more sensitivity to the estimates of the stock risks and interaction effects used as inputs to create low-volatility portfolios. For the former, having a variety of strategies to achieve low volatility can be useful for investors assessing the options and building multimanager strategies. However, for the latter, increased sensitivity to inputs can be dangerous because even identical strategies can lead to very different outcomes given differences in the inputs they use as a starting point.

STRATEGY DEFINITIONS AND COMPARISONS

In an earlier article (Qian, Alonso, and Barnes 2015), the authors addressed the analytical differences among many of the approaches we discuss here (i.e., minimum variance, maximum diversification, and risk balanced). In this article, we add a fourth strategy called *inverse volatility*. Instead of delving into the mathematics, we focus on the empirical results of generating backtests from these strategies using different risk models: RM1, RM2, and RM3. Although the risk models differ in how they estimate risk, in our view they all represent robust factor-based risk models that are widely used and commonly accepted within our industry. Therefore, we assume that any differences among the risk models are not necessarily errors, but rather the result of fundamentally different approaches to risk estimation.

We give concise descriptions of each strategy and how we calculate backtests for each:

- Inverse volatility (IV): The IV strategy does not have a specific volatility target and does not use an optimizer in its construction. The IV strategy simply picks the 100 lowest volatility stocks in the universe (we use the S&P 500 in this article) and weights them in proportion to the inverse of their standard deviations (a construction process similar to that of the S&P 500 Low Volatility Index). The standard deviations are derived from the risk models we use. In the first iteration of the IV strategy, we use the individual stock volatility estimates that come directly from RM1. The second iteration uses individual volatilities estimated by the RM2 risk model, and the third iteration uses the RM3 risk model. Thus, the only differences among the three iterations of the IV strategy come from the differences in stock volatility measurements in the three risk models. The greater the differences in the risk models, the greater the differences in the names selected and the weights given.
- Minimum variance (MV): The MV strategy is calculated using an optimizer in which the objective is to find a set of stocks and weights, given the risk inputs that have the lowest possible portfolio volatility. Again, we run three versions of these optimizations in which the only difference is in the risk and covariance estimations found in the three risk models we consider. Other than a long-only constraint, no constraints are added to the optimization. Although we acknowledge that constraints are often used for these strategies, our goal here is to show the unconstrained or unadulterated results for minimum variance portfolios given only changes to the risk models used in their calculation.
- Maximum diversification (MD): The MD strategy (Choueifaty and Coignard 2008) is also calculated using an optimizer. In this case, the objective is to find a set of stocks and weights that maximizes the

ratio of the portfolio-weighted stock volatilities over the portfolio volatility. This ratio is referred to as the *diversification ratio*, and the optimization aims to maximize its value. Other than a long-only constraint, this optimization is also run without constraints and again our objective is to understand how different the solutions to the same objective are, given only changes in the risk inputs. MD is a lower beta strategy. As we show later, low beta can be much different from low volatility.

- Risk balanced (RB): Rather than maximize a diversification ratio or minimize portfolio volatility, RB investing aims to reduce risk concentration in a portfolio by balancing the contributions to portfolio risk from every bet in a given portfolio (Qian 2005). In our definition, a *bet* is an independent source of risk coming from either a stock's sector or the stock itself. As such, in RB portfolios, one seeks to balance risk across the sectors in the portfolio and across the stocks within the sectors represented in the portfolio. Although an optimizer is not used to achieve this balance, there is an element of stock selection. The approach in this study starts with the 75 stocks representing the lowest risk estimates from our risk models and then balances the risk from these stocks, first according to their sector representation and then their individual risks. It is important to note that a stock's risk in this context is the stock's total contribution to portfolio risk—a combination of a stock's own volatility and its co-movement with other stocks in the portfolio. Again, we run three versions of this approach in which the only difference is the risk models used as inputs to the process.

DATA

To compare the four portfolio construction approaches (IV, MV, MD, and RB) across the three risk models (RM1, RM2, and RM3), we derive a total of 12 backtests. For brevity, we constrain our study to the US market constituents of the S&P 500 Index. The constituent and return data used are from January 1995 through July 2019. We use monthly total returns to compare returns and risks across these backtests. Finally, we use a relatively standard list of risk factors to compare exposures across construction methodologies. We also examine the impact that changing risk models has on these exposures within each construction methodology.

RISK MODELS

Before we describe the differences among the risk models we use in this study, an obvious question may be why we use risk models in the first place and how, in general, they are constructed. The use of risk models is common to quantitative investing as the estimation of a stock's volatility, and its covariance with other stocks, coupled with the relatively short stock price history we have, leads to issues of dimensionality. With 500 stocks in the S&P 500, the task of estimating all of the pairwise correlations among stocks requires the calculation of 124,750 ($[500^2 - 500]/2$)[3] covariances. Given that we are working with monthly data from January 1995 through July 2019, there are only 295 return observations for any given stock. (This is assuming that the stock has existed throughout the entire time period.) In short, data are insufficient to ensure that the calculation of each covariance is not overwhelmed with estimation error.

The presence of estimation error then makes it difficult to ensure that the resultant variance–covariance matrix retains the statistical properties required to subject it to an optimization (e.g., invertibility). To solve this problem, risk models are used because they significantly reduce the dimensionality issue. Essentially, the approach used to build a risk model requires the identification of common risk factors—in the risk models used for the purposes of this article, fundamental factors (e.g., valuation ratios) or behavioral factors (e.g., momentum), because they have been shown to explain a large percentage of common stock risk (Cochrane 2011). Once the common risk factors are identified, the problem becomes tenable. All that is required is to calculate the variance–covariance matrix of the common risk factors and then the sensitivity of each stock to those common risk factors. In this scenario, the dimensionality of the problem is reduced from the number of stocks (i.e., 500) to the number of risk factors (usually around 10–20) and sectors or industries. After a relatively simple

[3] In a 500×500 covariance matrix, there are $500^2 = 250{,}000$ elements, of which 500 are on the diagonal. Therefore, 250,000 − 500 = 249,500 are off the diagonal. Covariance matrixes are symmetric, and so there are 249,500/2 = 124,750 unique pairwise correlation estimates.

matrix operation, a full variance covariance matrix can be generated (500 × 500 in the case of the S&P 500) that not only represents each stock's volatility and covariation with other stocks but also retains the statistical properties needed for optimization.

Finally, the process of constructing risk models can also allow for the inclusion of other statistical manipulations that account for potentially pesky issues in higher-order moments such as skewness and kurtosis or even include techniques to forecast risk, such as the generalized autoregressive conditional heteroskedasticity model (GARCH). The variation in risk models then largely comes down to the common factors used to describe stock risk and the length of time over which the stock sensitivities to these common risk factors is estimated. In each of the risk models we use, fundamental risk factors are the dominant descriptors of stock-level risk. Risk models can also use macroeconomic and statistical common risk factors. However, for the sake of comparability, we limit ourselves to three fundamentally based long-term risk models.

PERFORMANCE DATA ANALYSIS

In this section, we show a variety of results for the 12 backtests we generated. Our backtests can be broken into the following groupings; these will be important as we make comments on the results of our analysis.

Nonoptimized Portfolio Construction Approaches:

- IV
  - Backtest 1 (IV_RM1): Inverse volatility portfolio construction using the RM1 risk model
  - Backtest 2 (IV_RM2): IV portfolio construction using the RM2 risk model
  - Backtest 3 (IV RM3): IV portfolio construction using the RM3 risk model
- RB
  - Backtest 4 (RB_RM1): RB portfolio construction using the RM1 risk model
  - Backtest 5 (RB_RM2): RB portfolio construction using the RM2 risk model
  - Backtest 6 (RB_RM3): RB portfolio construction using the RM3 risk model

EXHIBIT 1

Total and Relative Performance of Four Low-Volatility Strategies from January 1995 through July 2019

	IV	MV	MD	RB	S&P 500
Panel A: Performance across Risk Models					
Avg. Total Return	11.57	10.62	9.46	11.96	10.00
Avg. Excess Return	8.96	8.03	6.89	9.34	7.42
Avg. Volatility	11.03	11.11	13.86	11.73	14.64
Avg. Sharpe Ratio	0.81	0.72	0.50	0.80	0.51
Panel B: Relative Performance across Risk Models					
Avg. Value Added	0.93	0.06	–0.65	1.36	
Avg. Tracking Error	9.65	11.42	10.04	8.99	
Avg. Information Ratio	0.10	0.01	–0.06	0.15	

Optimized Portfolio Construction Approaches:

- MV
  - Backtest 7 (MV_RM1): MV portfolio construction using the RM1 risk model
  - Backtest 8 (MV_RM2): MV portfolio construction using the RM2 risk model
  - Backtest 9 (MV_RM3): MV portfolio construction using the RM3 risk model
- MD
  - Backtest 10 (MD_RM1): MD portfolio construction using the RM1 risk model
  - Backtest 11 (MD_RM2): MD portfolio construction using the RM2 risk model
  - Backtest 12 (MD_RM3): MD portfolio construction using the RM3 risk model

In Exhibit 1, we show the average performance characteristics of the S&P 500 Index and different portfolio construction techniques across all three risk models. This exhibit illustrates that the strategies considered in this study, on average, are lower-volatility strategies and generally beat the S&P 500 Index in either total return or Sharpe ratio. IV has the highest average Sharpe ratio because it is the strategy with the lowest realized volatility, but RB has a higher return. Although both MV and MD have lower volatilities than the S&P 500, they also have lower returns than IV or RB.

Exhibit 2

Total Performance of All Variations of Different Low-Volatility Strategies from January 1995 through July 2019

Panel A: Total Returns in Excess of the Risk-Free Rate

	IV	MV	MD	RD	S&P 500 Return – Risk-Free Rate
RM1	9.12	9.43	7.22	9.41	7.42
RM2	8.73	9.09	8.34	9.50	
RM3	9.02	5.55	5.12	9.10	
Range	0.38	3.88	3.22	0.40	

Panel B: Portfolio Volatility

	IV	MV	MD	RB	S&P 500 Volatility
RM1	11.52	11.76	14.43	11.66	14.64
RM2	10.94	10.65	13.31	11.75	
RM3	10.63	10.93	13.83	11.76	
Range	0.89	1.11	1.12	0.10	

Panel C: Sharpe Ratio

	IV	MV	MD	RB	S&P 500 Sharpe Ratio
RM1	0.79	0.80	0.50	0.81	0.51
RM2	0.80	0.85	0.63	0.81	
RM3	0.85	0.51	0.37	0.77	
Range	0.06	0.35	0.26	0.03	

The lower returns for the MD strategy may also be explained by a distinction between low volatility and low beta. These terms are not synonymous; the objective function in the MD portfolio calculation favors stocks with higher volatility while still achieving lower beta portfolios. Hence, the lower average returns for this strategy may be explained by the MD strategy's lower loading on low-volatility names and reduction in overall exposure to the low-volatility premiums that these stocks empirically carry.

Now that we have shown that (on average) all four of the strategies we studied exhibit lower realized volatility than the S&P 500 Index, we will examine the differences in performance characteristics within each strategy and across the three risk models. In Exhibit 2, we first show the differences in return realizations for each portfolio construction technique.

Exhibit 2 also illustrates the absolute performance characteristic of each iteration of our backtests. We first observe that the variability or range in each statistic (return, risk, and Sharpe ratio) is an order of magnitude greater for the optimized strategies (MV and MD) than for the nonoptimized strategies (IV and RB). For both return and Sharpe ratio, the range of values for the MV portfolio construction approach is higher than that of MD, although both ranges are similar for portfolio risk. This result is not surprising (although the magnitude may be) when taking into consideration earlier analytical findings (Qian, Alonso, and Barnes 2015) on the optimized approaches discussed here with regard to higher dependency on different components of risk compared with nonoptimized approaches. Namely, MV is highly sensitive to volatility estimates, whereas MD is highly sensitive to covariance estimates.

Because the goal of a risk model is to estimate risk, a certain level of model-specific estimation error is always present. Any portfolio that disproportionately relies on the model estimation of volatility (MV) or covariance (MD) is prone to maximizing model estimation error. More generally, other papers have been written on the sensitivity of optimizers, and we believe that these sensitivities are principally responsible for the range of outputs we see in Exhibit 2 (Michaud 1989; Jorion 1992; Broadie 1993; Ledoit and Wolf 2004).

Thus far, we have limited our analysis to the absolute returns of the portfolio. To get a better idea of what might be driving these results, we also looked at results in excess of the S&P 500 return. Using excess returns allows for a different and important perspective on how the portfolio construction approaches diverge but also on how (even within the same portfolio construction methodology) sensitivities to risk models can drive different results. Exhibits 3 and 4 show different relative characteristics of the performance of our backtests.

In Exhibit 3, we again show that the range of outcomes is an order of magnitude higher for the optimized approaches than for the nonoptimized approaches. This result may not be too surprising given the absolute performance results we saw in Exhibit 2. However, one additional observation is that only for the optimized strategies do we see backtest iterations that lose value versus the benchmark. This, in turn, results in negative information ratios in some cases. The potential for positive or negative value-added results depends only on which risk model was used and not on the expectations inherent to the strategy's objectives. This is concerning, considering the variability in outcomes does not point to the superiority of one risk model over another. Simply

Exhibit 3

Performance Relative to the S&P 500 from January 1995 through July 2019

	IV	MV	MD	RB
Panel A: Value-Added Return				
RM1	1.13	1.45	–0.26	1.41
RM2	0.71	1.01	0.62	1.51
RM3	0.94	–2.27	–2.33	1.14
Range	0.42	3.72	2.95	0.37
Panel B: Tracking Error				
RM1	9.53	10.41	8.48	9.28
RM2	9.86	11.69	11.05	8.83
RM3	9.56	12.17	10.60	8.84
Range	0.34	1.76	2.57	0.45
Panel C: Information Ratio				
RM1	0.12	0.14	–0.03	0.15
RM2	0.07	0.09	0.06	0.17
RM3	0.10	–0.19	–0.22	0.13
Range	0.05	0.33	0.28	0.04

put, any differences in risk models, even those that are valid, can lead to very different outcomes.

Exhibit 4 shows relative performance as we look at performance patterns conditional on the benchmark performance. In Panel A, we see that a new distinction can possibly be made among the strategies. Both IV and MV have the lowest upside participation ratios, whereas MD and RB have higher upside participation. The corollary to upside participation can be seen in Panel B, which shows downside participation (downside protection = 1 – downside participation); IV and MV have the lowest downside participations. This may suggest that the value add from IV and MV versus MD and RB is more dependent on downside protection and thus occurs during market downturns. However, when we look at Panel D, we see that the win percentage in down markets is nearly identical among IV, MV, and RB. This suggests that persistence in downside protection among the nonoptimized approaches is the same as for MV, meaning that any advantage MV has in average downside participation is likely the result of a few observations and not a superior robustness. One last observation is that for both IV and RB, the upside and downside participations are less extreme than those for either MV or MD. A conclusion we can draw from these observations is that optimized solutions lean toward more extreme behavior (i.e., extreme but relatively infrequent upside versus downside performance), which ultimately can lead to less consistent and potentially inferior results.

Exhibit 4

Conditional Performance Patterns from January 1995 through July 2019

	IV	MV	MD	RB
Panel A: Upside Participation				
RM1	0.73	0.73	0.88	0.76
RM2	0.70	0.64	0.77	0.76
RM3	0.70	0.52	0.71	0.76
Average	0.71	0.63	0.79	0.76
Range	0.04	0.21	0.17	0.00
Panel B: Downside Participation				
RM1	0.47	0.44	0.81	0.49
RM2	0.43	0.31	0.56	0.49
RM3	0.42	0.33	0.67	0.51
Average	0.44	0.36	0.68	0.49
Range	0.05	0.14	0.25	0.02
Panel C: Participation Ratio				
RM1	1.57	1.65	1.08	1.55
RM2	1.61	2.07	1.37	1.56
RM3	1.67	1.58	1.07	1.49
Average	1.62	1.77	1.18	1.53
Range	0.10	0.48	0.29	0.07
Panel D: Win Percentage in Down Markets				
RM1	0.80	0.74	0.57	0.81
RM2	0.78	0.86	0.69	0.81
RM3	0.80	0.82	0.61	0.80
Average	0.80	0.81	0.62	0.81
Range	0.02	0.11	0.12	0.01

In our last piece of analysis on performance, we look briefly at its variations across our backtests in different subperiods. To perform this analysis (and to try and avoid any biases in our selection of subperiods), we first calculate the cross-sectional variation in rolling one-year returns across all of the backtests. We then simply choose the top five non-overlapping periods as representative timeframes in which our backtests deviate from each other most. Exhibit 5 graphs the cross-sectional rolling one-year performance standard deviation of our backtests.

The top five standard deviations of non-overlapping periods that occurred in the subperiods are shown in Exhibit 6.

EXHIBIT 5

Cross-Sectional One-Year Performance Standard Deviation

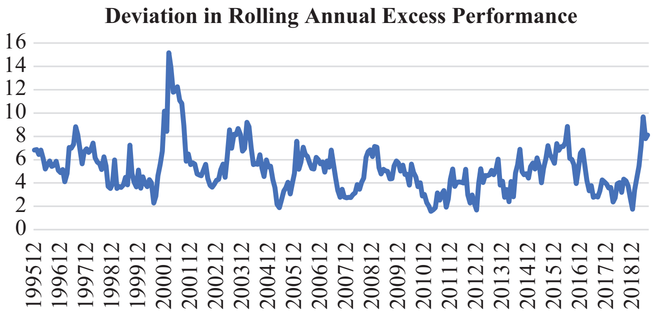

Exhibit 6 was generated using highlights to show above median performance within each row for the different portfolio construction approaches. First, as in our prior exhibits, the range of outcomes is almost always higher for the optimized solutions (with the one-year performance of MD through May 2009 being the exception). Second, in four of the five subperiods (with the one-year period of MD through February 2004 being the exception), the nonoptimized solutions, and especially the RB solution, produced the best results. Finally, four of the five subperiods had positive returns in the S&P 500. This last observation is important because these are all lower volatility strategies and one might expect less deviation among them in a down market. (Given they all tend to perform disproportionately better in down markets.) Our observation potentially suggests that the inconsistency we see in the optimized solutions tends to occur in up markets and not in strong down markets.

We address this in Exhibit 7, which shows two of the strongest non-overlapping down periods in this analysis.

In Exhibit 7, we now see that the MV strategy indeed provides strong downside protection, but again the ranges for the optimized strategies (MV and MD) remain significantly larger than those for the nonoptimized strategies (IV and RB).

PORTFOLIO CHARACTERISTIC ANALYSIS AND ROBUSTNESS

A strategy's performance is ultimately a byproduct of the exposures taken in the portfolio, be it sector exposures, factor exposures, or some combination of each.

EXHIBIT 6

Top Five Subperiods in which Cross-Sectional Standard Deviation of One-Year Performance across All Backtests Was Highest

	IV	MV	MD	RB	S&P 500
Panel A: One-Year Performance through January 1997					
RM1	40.27	33.96	38.64	42.13	52.14
RM2	36.88	29.84	20.78	43.95	
RM3	36.73	17.60	26.96	44.44	
Average	37.96	27.13	28.80	43.51	
Range	3.54	16.35	17.87	2.31	
Panel B: One-Year Performance through February 2001					
RM1	41.21	27.02	19.61	32.78	–8.20
RM2	43.26	24.60	8.27	36.56	
RM3	44.02	9.33	–2.49	34.69	
Average	42.83	20.31	8.46	34.68	
Range	2.81	17.68	22.11	3.77	
Panel C: One-Year Performance Through February2004					
RM1	33.30	44.90	52.73	33.93	38.52
RM2	32.97	36.92	59.62	31.12	
RM3	34.21	42.74	31.79	33.20	
Average	33.49	41.52	48.05	32.75	
Range	1.24	7.98	27.83	2.80	
Panel D: One-Year Performance through June 2016					
RM1	27.74	32.73	2.17	30.44	3.99
RM2	20.18	16.12	19.06	28.32	
RM3	17.46	15.83	13.19	28.40	
Average	21.79	21.56	11.47	29.05	
Range	10.28	16.90	16.89	2.12	
Panel E: One-Year Performance through May 2019					
RM1	17.83	13.85	–2.54	20.17	3.78
RM2	18.60	24.21	–2.37	19.68	
RM3	14.79	10.87	–2.43	18.58	
Average	17.07	16.31	–2.45	19.48	
Range	3.81	13.33	0.16	1.59	

Note: Highlights mark results above the median in each row.

The differences in performance shown in Exhibit 7 then must be the result of differences in the exposures each backtest has taken on over time. In this final analysis, we look at the variations of each strategy's exposure to risk factors and sectors over time and across risk models.

To gain some insight into how each of our four strategies differ in exposures, Exhibits 8 and 9 show the

Exhibit 7

Performance in the Strongest Non-overlapping One-Year Down Markets

	IV	MV	MD	RB	S&P 500
Panel A: One-Year Performance through September 2001					
RM1	8.82	7.23	–3.12	7.23	–26.62
RM2	7.71	3.53	–9.27	10.79	
RM3	11.70	0.47	–18.09	6.03	
Average	9.41	3.74	–10.16	8.01	
Range	3.99	6.76	14.97	4.77	
Panel B: One-Year Performance through February 2009					
RM1	–30.82	–28.42	–41.47	–30.16	–43.32
RM2	–31.51	–18.10	–25.62	–31.49	
RM3	–26.88	–21.98	–43.09	–31.91	
Average	–29.74	–22.83	–36.73	–31.18	
Range	4.64	10.32	17.47	1.75	

Exhibit 8

Average Time-Series Sector Exposures across Backtests

Average Sector Exposures	IV	MV	MD	RB	S&P 500
Consumer Discretionary	0.09	0.08	0.10	0.09	0.12
Consumer Staples	0.25	0.39	0.19	0.17	0.10
Energy	0.02	0.02	0.06	0.06	0.09
Financials	0.14	0.07	0.05	0.15	0.17
Healthcare	0.10	0.07	0.15	0.10	0.13
Industrials	0.10	0.08	0.06	0.11	0.11
Materials	0.04	0.04	0.12	0.08	0.03
Information Technology	0.03	0.02	0.08	0.03	0.18
Telecommunications	0.02	0.02	0.02	0.02	0.05
Utilities	0.21	0.19	0.15	0.19	0.03
Dispersion	0.08	0.11	0.05	0.06	0.05
Herfindahl Index	0.16	0.22	0.12	0.13	0.12

sector and risk factor exposures for IV, MV, MD, and RB, averaged across each strategy's backtests.

Exhibit 8 illustrates how all four strategies have relatively large weights to the utility and consumer staples sectors. (This is reasonable, given that these are historically low-volatility sectors.) MV tends to have a much larger concentration in consumer staples than the other strategies. The dispersion and Herfindahl metrics at the bottom of Exhibit 8 show that MD and RB tend to have the most balanced sector exposures.

Exhibit 9 shows the risk factor exposures of each strategy. All of the strategies have a low exposure to beta, which is to be expected. More interestingly, though, MD stands apart from the other strategies, with higher exposures to a variety of risk factors (i.e., book to price, earnings quality, earnings yield, and growth), meaning that some of the differences in performance for MD come from relative positive loadings to these factors. Also interesting is that both MV and MD have higher exposures to profitability and lower exposure to size (MV and MD are thus biased toward smaller companies), although neither of these exposures is explicitly targeted.

Another important perspective is how variable these exposures are within a given strategy, as the risk model used to build the strategy is changed. Exhibits 10 and 11 illustrate the range of exposures to sectors and risk factors experienced by each strategy, as a result of differences in the risk models.

In Exhibit 10, it is notable how much variation can be seen in the IV, MV, and MD strategies, whereas

Exhibit 9

Average Factor Exposures over Time and across Backtests

Average Risk Factor Exposure	IV	MV	MD	RB	S&P 500
Beta	–0.80	–1.02	–0.52	–0.73	0.01
Book to Price	0.05	–0.01	0.17	0.08	–0.05
Dividend Yield	0.56	0.60	0.17	0.60	0.10
Earnings Quality	0.01	0.09	0.24	0.02	0.00
Earnings Variability	–0.42	–0.42	0.10	–0.42	–0.10
Earnings Yield	0.10	0.01	–0.03	0.13	0.06
Growth	–0.43	–0.44	–0.16	–0.47	–0.09
Investment Quality	0.34	0.34	0.11	0.35	0.09
Leverage	0.25	0.21	0.12	0.17	0.00
Liquidity	–0.23	–0.23	0.22	–0.24	–0.09
Long-Term Reversal	0.04	0.07	0.07	0.04	–0.05
Mid Cap	0.37	0.44	0.37	0.35	–0.15
Momentum	0.00	–0.02	–0.05	–0.03	0.04
Profitability	0.01	0.19	0.11	–0.08	0.05
Residual Volatility	–0.57	–0.58	–0.06	–0.63	–0.13
Size	–0.31	–0.49	–0.40	–0.28	0.39

RB shows almost no variation. RB is a risk-balanced strategy, and sector risk is explicitly balanced, which leads to the low range of outcomes. However, risk balance is not weight balanced. Therefore, the stability of the sector weights for the RB strategy stems from both the explicit sector risk balance and the stability of each sector subportfolio that comes from balancing risk at the stock level. Exhibit 10 also suggests that many

EXHIBIT 10

Range in Time-Series Sector Exposures across Backtests

Average Range (maximum–minimum) Sector Exposures	Inverse Volatility	Minimum Variance	Maximum Diversification	Risk Balanced
Consumer Discretionary	7%	11%	14%	2%
Consumer Staples	14%	24%	19%	3%
Energy	4%	4%	5%	2%
Financials	14%	14%	7%	3%
Healthcare	7%	10%	14%	3%
Industrials	5%	13%	10%	3%
Materials	4%	7%	9%	3%
Information Technology	2%	5%	8%	2%
Telecommunications	2%	4%	5%	1%
Utilities	10%	21%	16%	3%
Average	7%	11%	11%	3%
Herfindahl Index	0.064	0.170	0.137	0.007

low-volatility strategies can take on very different sector exposures at different times in the market.

Exhibit 11 shows how sensitive the different exposures are to changes in the risk model by showing the range (maximum – minimum) of exposures to each risk factor across the three backtests for each strategy. For example, it shows that the range of exposures for the MD strategy to beta is 0.66, whereas the range of exposures to beta for the RB strategy is only 0.07. Interestingly, RB shows very little variability and thus minimal sensitivity to changes in the risk models. Both MV and MD, by contrast, show the most extreme sensitivity to risk model changes. This is an important observation because the characteristics of the portfolios can change drastically, even if the portfolio construction methodology does not.

EXPLANATION

Constructing portfolios with an optimizer, especially MV portfolios, can maximize the idiosyncrasies of the risk model used in the process (Chopra and Ziemba 1993). A more balanced, less risk model–dependent approach provides similar levels of ex post risk while creating more stability across time and different risk models. RB and IV exhibit less dependency on the correlation structure of the risk model, compared to MV and MD. Because MV and MD portfolios are optimized to a specific target, they are much more sensitive to the risk model used. Hence, they exhibit much larger fluctuations in their risk exposures than do RB and IV.

Cross-sectional and time-series ordinal correlation of stocks' total risk, as predicted by the three risk models, is fairly high. That being the case, there is much agreement across risk models regarding which stocks are more or less volatile. Therefore, the IV and RB approaches, which rely mostly on selecting the lowest-risk stocks, tend to be very stable across time. However, approaches that rely on optimizations have higher sensitivity to the correlation structure of the risk models. Stock correlation estimates tend to be less stable across time and across different risk models. There is less certainty regarding the relationships among stocks across risk models and time. Therefore, optimized MV and MD approaches that rely on that factor correlation structure are generally less stable across time and risk models.

Exhibit 12 shows two statistics that help empirically justify our explanation. In the first column, we show the average monthly correlation of stock volatility estimates across our three risk models. In the second column, we show the average monthly correlation of pairwise stock correlations across our three risk models. The average monthly correlations of stock volatilities are higher, indicating that there is more agreement among

EXHIBIT 11

Range in Time-Series Risk Factor Exposures across Backtests

Average Range (maximum–minimum) Risk Factor	Inverse Volatility	Minimum Variance	Maximum Diversification	Risk Balanced
Beta	0.20	0.49	0.66	0.07
Book to Price	0.21	0.38	0.40	0.07
Dividend Yield	0.24	0.37	0.49	0.08
Earnings Quality	0.13	0.29	0.27	0.08
Earnings Variability	0.06	0.26	0.42	0.04
Earnings Yield	0.10	0.22	0.27	0.05
Growth	0.12	0.26	0.35	0.05
Investment Quality	0.10	0.25	0.34	0.06
Leverage	0.13	0.40	0.40	0.08
Liquidity	0.06	0.43	0.44	0.06
Long Term Reversal	0.08	0.22	0.40	0.05
Mid Cap	0.09	0.34	0.40	0.08
Momentum	0.08	0.24	0.29	0.05
Profitability	0.37	0.51	0.36	0.09
Residual Volatility	0.11	0.30	0.45	0.06
Size	0.14	0.51	0.65	0.12
Average	0.14	0.34	0.41	0.07

EXHIBIT 12

Average Monthly Correlation across Risk Models (2003–2019)

	Stock Volatilities	Stock Correlations
Average Monthly Correlations	0.77	0.50

the risk models. On the other hand, the correlations of the pairwise stock correlations are lower, indicating that there is less agreement along this dimension.

CONCLUSION

There are many different ways of getting exposure to the low volatility premiums (or anomaly, as the case may be). In this article, we compared what we consider to be some of the most popular approaches. Each has a very different portfolio construction methodology and philosophy. By and large, however, they all focus on achieving low-volatility (or reduced-volatility low beta) portfolios. We point out that the risk inputs used in the portfolio construction processes represent a common sensitivity in each of these approaches. To understand how these sensitivities differ, and how they can lead to different results, we ran backtests within four implementation approaches—three low-volatility approaches (IV, MV, RB) and one low-beta approach (MD)—using three risk models and analyzed the results.

We find that the two optimized approaches—MV and MD—have much larger sensitivities to changes in their risk input, which can lead to very different results in their backtests and portfolio characteristics. We believe this sensitivity should be of concern to asset allocators and asset managers alike because the different outcomes do not necessarily point to the superiority of one risk model over another. Instead, these outcomes are the result of differences in the risk models that may be completely valid. Given these variations, we believe that a more stable and consistent exposure to low-volatility investing should be desirable to investors.

REFERENCES

Baker, M., B. Bradley, and J. Wurgler. 2011. "Benchmarks as Limits to Arbitrage: Understanding the Low-Volatility Anomaly." *Financial Analysts Journal* 67 (1): 40–54.

Broadie, M. 1993. "Computing Efficient Frontiers using Estimated Parameters." *Annals of Operations Research* 45 (1): 21–58.

Chan, L., J. Karceski, and J. Lakonishok. 1999. "Portfolio Optimisation: Forecasting Covariances and Choosing the Risk Model." *Review of Financial Studies* 12: 937–974.

Chopra, V. K., and W. T. Ziemba. 1993. "The Effect of Errors in Means, Variances and Covariances on Optimal Portfolio Choice." *The Journal of Portfolio Management* 19 (2): 6–11.

Choueifaty, Y., and Y. Coignard. 2008. "Toward Maximum Diversification." *The Journal of Portfolio Management* 35 (1): 40–51.

Clarke R., H. DeSilva, and S. Thorley. 2006. "Minimum-Variance Portfolios in the US Equity Market." *The Journal of Portfolio Management* 33 (1): 10–24.

Cochrane, J. H. 2011. "Presidential Address: Discount Rates." *Journal of Finance* 66 (4): 1047–1108.

Haugen, R. A., and N. L. Baker. 1991. "The Efficient Market Inefficiency of Capitalization-Weighted Stock Portfolios." *The Journal of Portfolio Management* 17 (3): 35–40.

Haugen, R. A., and A. J. Heins. "On the Evidence Supporting the Existence of Risk Premiums in the Capital Markets." Working paper, SSRN No. 1783797, 1972.

——. 1975. "Risk and the Rate of Return on Financial Assets: Some Old Wine in New Bottles." *Journal of Financial and Quantitative Analysis* 10 (5): 775–784.

Jagannathan, R., and T. Ma. 2003. "Risk Reduction in Large Portfolios: Why Imposing the Wrong Constraints Helps." *Journal of Finance* 58 (4): 1651–1684.

Jorion, P. 1992. "Portfolio Optimization in Practice." *Financial Analysts Journal* 48 (1): 68–74.

Ledoit, O., and M. Wolf. 2004. "Honey, I Shrunk the Sample Covariance Matrix." *The Journal of Portfolio Management Summer* 30 (4): 110–119.

Michaud, R. O. 1989. "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal* 45 (1): 31–42.

Qian, E. 2006. "On the Financial Interpretation of Risk Contribution: Risk Budgets Do Add Up." *Journal of Investment Management* 4 (4).

Qian, E., N. Alonso, and M. Barnes. 2015. "The Triumph of Mediocrity: A Case Study of Naïve Beta." *The Journal of Portfolio Management* 41: 19–34.

PanAgora is exempt from the requirement to hold an Australian financial services license under the Corporations Act 2001 in respect of the financial services.

PanAgora is regulated by the SEC under US laws, which differ from Australian laws.

Past performance is not a guarantee of future results.

ADDITIONAL READING

The Volatility Effect

David C. Blitz and Pim van Vliet

The Journal of Portfolio Management

https://jpm.pm-research.com/content/34/1/102

ABSTRACT: *There is empirical evidence that stocks with low historical volatility have high risk-adjusted returns, with annual alpha spreads of global low-versus high-volatility decile portfolios of 12 percentage points over 1986–2006. This volatility effect appears independently in US, European, and Japanese markets. It is similar in size to classic effects such as value, size, and momentum, and cannot be explained by implicit loadings on these well-known effects. These results indicate that equity investors overpay for risky stocks. Possible explanations include leverage restrictions, inefficient two-step investment processes, and behavioral biases of private investors. To exploit the volatility effect in practice, investors might include low-risk stocks as a separate asset class in the strategic asset allocation phase of the investment process.*

The Volatility Effect Revisited

David Blitz, Pim van Vliet, and Guido Baltussen

The Journal of Portfolio Management

https://jpm.pm-research.com/content/46/2/45

ABSTRACT: *High-risk stocks do not have higher returns than low-risk stocks in all major stock markets. This article provides a comprehensive overview of this low-risk effect, from the earliest asset pricing studies in the 1970s to the most recent empirical findings and interpretations. Volatility appears to be the main driver of the anomaly, which is highly persistent over time and across markets and which cannot be explained by other factors such as value, profitability, or exposure to interest rate changes. From a practical perspective, low-risk investing requires little turnover, volatilities are more important than correlations, low-risk indexes are suboptimal and vulnerable to overcrowding, and other factors can be efficiently integrated into a low-risk strategy. Finally, there is little evidence that the low-risk effect is being arbitraged away because many investors are either neutrally positioned or even on the other side of the low-risk trade.*