

**Benchmarks as Limits to Arbitrage:
Understanding the Low Volatility Anomaly***

Malcolm Baker
Harvard Business School and NBER

Brendan Bradley
Acadian Asset Management

Jeffrey Wurgler
NYU Stern School of Business and NBER

First Draft: December 2009

This Draft: March 22, 2010

Abstract

Over the past 41 years, high volatility and high beta stocks have substantially underperformed low volatility and low beta stocks in U.S. markets. We propose an explanation that combines the average investor's preference for risk and the typical institutional investor's mandate to maximize the ratio of excess returns and tracking error relative to a fixed benchmark (the information ratio) without resorting to leverage. Models of delegated asset management show that such mandates discourage arbitrage activity in both high alpha, low beta stocks and low alpha, high beta stocks. This explanation is consistent with several aspects of the low volatility anomaly including why it has strengthened in recent years even as institutional investors have become more dominant.

* We thank Gary Bergstrom, John Campbell, Edwin Elton, Ronald Frashure, Xavier Gabaix, Laurent De Greef, Robin Greenwood, Martin Gruber, Anthony Lynch, Ted Noon, Jay Ritter, Ryan Stever, Rui Tang, Luis Viceira, and Rob Whitelaw for helpful comments, and Rui Tang and Ferris Zhang for excellent research assistance. Baker gratefully acknowledges the Division of Research of the Harvard Business School for financial support.

While there are many candidates for the greatest anomaly in finance, perhaps the most worthy is the long-term success of low volatility and low beta stock portfolios. Over the 41 years between 1968 and 2008, low volatility and low beta portfolios have offered an enviable combination of high average returns and small drawdowns. This runs counter to the fundamental economic principle that risk is compensated with higher expected return. We apply principles of behavioral finance to shed light on the drivers of this anomalous performance and to assess the likelihood that it will persist.

Behavioral models of security prices combine two ingredients. The first is that some market participants are irrational in a particular way. In the context of the low risk anomaly, we believe that a preference for lotteries and well established biases of representativeness and overconfidence lead to demand for higher volatility stocks that is not warranted by fundamentals.

The second ingredient is limits to arbitrage—an explanation for why the “smart money” does not offset the price impact of any irrational demand. With respect to the low risk anomaly, we believe that the underappreciated limit on arbitrage is benchmarking. Many institutional investors who are in a position to offset the irrational demand for risk have fixed benchmark mandates, typically capitalization weighted, which by their nature discourage investments in low volatility stocks. We draw out the implications of Brennan’s [1993] model of agency and asset prices and show specifically that traditional fixed benchmark mandates with a leverage constraint cause institutional investors to pass up the superior risk-return tradeoff of low volatility portfolios. Rather than being a stabilizing force on prices, the typical institutional contract for delegated portfolio management is destabilizing, causing an increase in the demand for higher beta investments.

Ours is not the only explanation of the low-risk anomaly with behavioral elements. Karceski [2002] points out that mutual fund investors tend to chase returns over time and across funds, possibly as a result of an extrapolation bias, and that their fund flows are also sticky. These forces make fund managers care more about outperforming during bull markets than underperforming during bear markets, increasing their demand for high beta stocks and reducing their required returns. His model’s predictions

complement our own and the mechanisms can certainly work simultaneously. Our mechanism places the irrationality elsewhere and we focus on distortions introduced by benchmarking.

In this paper, we review the long-term performance of low risk portfolios, present our behavioral explanation, and discuss the practical implications for investors and investment managers. Perhaps the most important practical implication is that unless individual investors' preference for volatile stocks is somehow reversed, the ever-increasing importance of institutional investors with fixed benchmarks suggests that the low risk anomaly will persist.

THE LOW RISK ANOMALY

In an efficient market, investors realize above average returns only by taking above average risk. Risky stocks have high returns on average, and safe stocks do not. This simple empirical proposition has been hard to support based on the history of U.S. stock returns. The most widely used measures of risk actually point in the wrong direction, and rather strongly.

We take 41 years of data from January 1968 through December 2008 from the Center for Research on Security Prices (CRSP). We sort stocks into five groups each month according to trailing total volatility or trailing beta and track the returns on these portfolios. We also restrict the investing universe to the top 1000 stocks by market capitalization. **Exhibit 1** shows the results.

Regardless of whether we define risk as volatility or beta, or whether we consider all stocks or only large caps, low risk consistently outperforms high risk over this period. Panel A shows that a dollar portfolio invested in the lowest volatility portfolio in January 1968 would have increased to \$59.55. Over this period, inflation eroded the real value of a dollar to about \$0.17, meaning that the low-risk portfolio produced a \$10.12 gain in real terms. Contrast this with the performance of the highest volatility portfolio. A dollar invested here is worth 58 cents at the end of December 2008, assuming no transaction costs. Given the declining value of the dollar, the real value of the high volatility portfolio declines to less than 10 cents—a 90% decline in real terms! It is remarkable that over the last four decades the investor who aggressively pursued high volatility stocks would have borne almost a total loss in real terms.

Panel C considers beta as the measure of risk. Here, a dollar in the lowest beta portfolio grows to \$60.46 (\$10.28 in real terms) and a dollar in the highest beta portfolio grows to \$3.77 (64 cents in real terms). Like the high volatility investor, the high beta investor also fails to recover his dollar in real terms and underperforms his “conservative” beta neighbor by 964%.

Almost all mispricings are stronger for smaller firms than larger firms, but even within large firms the low risk anomaly has been dramatic. A dollar in low volatility large caps grew to \$55.53 over 41 years, while a dollar in high volatility large caps grew to \$24.14. For beta, the numbers are \$90.07 and \$10.14, respectively.

Finally, as if this puzzle was not bad enough, several facts only compound it.

- The low risk portfolios' paths to their higher dollar values have been much smoother. They are, as advertised, genuinely lower risk.
- The transaction costs of monthly rebalancing are greater, probably substantially, for the high volatility portfolio. This means that the relative performance in Exhibit 1 is understated.
- With the exception of the technology bubble, the return gap accelerated after 1983—a period during which institutional investment managers have become progressively more numerous, better capitalized, and more quantitatively sophisticated. Karceski [2002] also notes this trend.

These results are not entirely new, but they have not been sufficiently emphasized, explained, or exploited. In the 1970s, Black [1972], Black, Jensen, and Scholes [1972], and Haugen and Heins [1975] noted that the relationship between risk and return was much flatter than predicted by the CAPM. Haugen and Heins pointed out that it was not merely flat in their sample period, but actually inverted. Fama and French [1992] extended this analysis through 1990 and found that the relationship was flat, prompting a conclusion that “beta is dead.” More recently, Ang, Hodrick, Ying, and Zhang [2006, 2009] have drawn new attention to these results, finding that high volatility stocks have had “abysmally low returns” in longer U.S. samples and in international markets. Blitz and van Vliet [2007] provide a detailed analysis of the volatility anomaly and demonstrate its robustness across regions and to controls for size, value, and momentum effects. Bali, Cakici, and Whitelaw [2009] investigate a measure of lottery-like return distributions, which is highly correlated with other risk measures, and find that it is also associated with

poor performance. All told, the evidence for a risk-return tradeoff along the lines of the CAPM has, if anything, only deteriorated in the last few decades.

These patterns are hard to explain with traditional, rational theories of asset prices. In principle, beta might simply be the wrong measure of risk, too. The CAPM is just one equilibrium model of risk and return, with clearly unrealistic assumptions. For the past few decades, finance academics have devoted considerable energy to developing rational models, searching for the “right” measure of risk. Most of these newer models make the mathematics of the CAPM look quaint.

But despite superior computational firepower, the new models face an uphill battle. After all, the task is to prove that high volatility and high beta stocks are *less* risky. Granted a less risky stock might not be less volatile (although volatility and beta are positively correlated in the cross-section), but it must at least provide insurance against bad events. Even this notion of risk fails to resolve the anomaly. The high volatility quintile portfolio provided a relatively low return in precisely those periods when an insurance payment would have been most welcome, such as the downturns of 1973-1974 and 2000-2002, the crash of 1987, and the financial crisis beginning in the fall of 2008. Investors appear to be paying an insurance premium in the average month, only to lose even more when the equity market (and, often, the economy as a whole) is melting down.

We believe the long-term outperformance of low risk portfolios is perhaps the greatest anomaly in finance. It is large in economic magnitude and practical relevance and it challenges the basic notion of a risk-return tradeoff.

A BEHAVIORAL EXPLANATION

What is going on? We think two things drive these results: less than fully rational investor behavior and underappreciated limits to arbitrage.

The combination of irrational behavior in the presence of limited arbitrage is the core framework of behavioral finance, laid out in surveys such as Shleifer [2000], Barberis and Thaler [2003], and Baker and Wurgler [2007]. But what is the actual investor psychology that leads to a preference for volatile

stocks? What forces prevent smart institutions from taking advantage and in the process restoring the risk-return tradeoff? We address these questions in turn.

THE IRRATIONAL PREFERENCE FOR HIGH VOLATILITY

The preference for high volatility stocks derives from the biases that afflict the individual investor. We emphasize the following: the preference for lotteries, the representativeness bias, and the overconfidence bias. Each is well-known in the behavioral economics and finance literature. Each has its roots in the work of the Nobel-prize-winning psychologists Kahneman and Tversky. Each leads to a demand for volatile stocks that is not grounded in fundamentals.

Preference for lotteries. This one is the simplest. Would you take a gamble with a 50% chance of losing \$100 versus a 50% chance of winning \$110? Most people say no. Despite the positive expected payoff, the possibility of losing \$100 is enough to deter participation, even when \$100 is trivial in the context of wealth or income.

Kahneman and Tversky [1979] called this “loss aversion.” Taken on its own, loss aversion suggests investors would shy away from volatility for fear of realizing a loss. But something strange happens as the probabilities shift. This time, suppose you are presented a gamble with a near certain chance of losing \$1 and a small 0.12% chance of winning \$5,000. Like the first example, this has a positive expected payoff of around \$5. But in this case, most people take the gamble. The amounts spent on lotteries and roulette wheels, which have negative expected payoffs, are a clear manifestation of this tendency.

To be more precise, this is about positive skewness, where large positive payoffs are more likely than large negative ones, than it is about volatility. But Mitton and Vorkink [2007] remind us that volatile individual stocks, with limited liability, are also positively skewed. Buying a low priced, volatile stock is like a lottery ticket: There is some small chance of doubling or tripling in value in a short period, and a much larger chance of declining in value. Kumar [2009] finds that some individual investors do show a clear preference for stocks with lottery-like payoffs, measured as idiosyncratic volatility or skewness, and Barberis and Huang [2008] model this preference with the cumulative prospect theory approach of

Tversky and Kahneman [1992]. Blitz and van Vliet [2007] connect the preference for lottery tickets to the two-layer behavioral portfolio theory of Shefrin and Statman [2000].

Barberis and Xiong [2010] offer a preference-based explanation of the volatility effect that does not depend on a preference for skewness. They suggest that investors derive utility from *realizing* gains and losses on risky assets, rather than from paper gains and losses. These investors find volatile stocks very attractive: a volatile stock could go up a lot in the short-term, allowing investors to realize a large gain. Of course, it could also drop a lot in value. But in that case, investors simply delay selling it, thereby postponing the pain of realizing a loss into the distant future.

Representativeness. The classic way to explain representativeness is with a Tversky and Kahneman [1974] experiment. They described a fictional woman named Linda as “single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.” Then they asked subjects: Which is more probable? A: Linda is a bank teller. Or B: Linda is a bank teller who is active in the women’s movement. The fact that many subjects choose B suggests that probability theory and Bayes rule are not ingrained skills. If Linda is a bank teller who is active in the women’s movement (B is true), she *must* also be a bank teller (B is a proper subset of A). The mistake arises because the second job fits the description or is more “representative” of Linda. If we think of a bank teller who is active in the women’s movement, a person like Linda comes to mind more readily.

What does this have to do with stocks and volatility? Consider defining the characteristics of “great investments.” The layman and the quant address this question with two different approaches.

The layman tries to think of great investments – maybe buying Microsoft or Genzyme at their IPOs in 1986 – and concludes that the road to riches is paved with speculative investments in new technologies. The problem with this logic is similar to the Linda question. The layman largely ignores the high base rate at which small, speculative investments fail, and as a result is inclined to overpay for volatile stocks.

The quant, on the other hand, examines the full sample of stocks like Microsoft and Genzyme in an analysis like **Exhibit 1**. She concludes that without a way to separate the Microsofts from the losers, high risk stocks are generally to be avoided.

Overconfidence. Another pervasive bias underlying the preference for high volatility stocks is overconfidence (Fischhoff, Slovic, and Lichtenstein [1977] and Alpert and Raiffa [1982]). Experimenters ask subjects to estimate, for example, the population of Massachusetts and to provide a 90% confidence interval around this answer. The confidence interval should be wide enough so that for every 10 questions of this type, on average nine of them should contain the right answer. The experimental evidence shows that most people form confidence intervals that are far too narrow. Most people are simply overconfident in the accuracy of their information or judgment. Moreover, the more obscure the question—if it is the population of Bhutan instead of Massachusetts—the more this calibration deteriorates.

Valuing stocks involves this same sort of forecasting. What will revenues be five years hence? Overconfident investors are likely to disagree. Being overconfident, they will also agree to disagree, each sticking with the false precision of his or her estimate. The extent of disagreement is higher for more uncertain outcomes. For example, stocks that are either growing quickly or distressed – volatile stocks – will elicit a wider range of opinions.

The careful theorist will note that we require one extra assumption to connect overconfidence, or more generally differences of opinion, to the demand for volatile stocks. Pessimists must act less aggressively in markets than optimists. We must have a general reluctance or inability to short stocks relative to buying them. Empirically, the relative scarcity of short sales among individual investors and even institutional investors is evident, so this assumption is clearly valid. It means that prices are generally set by optimists, as pointed out by Miller [1977]. Stocks with a wider range of opinions will have more optimists among their shareholders and sell for higher prices, leading to lower future returns.

BENCHMARKING AS A LIMIT TO ARBITRAGE

Assuming that average investors have a psychological demand for high volatility stocks, the remaining and deeper economic question is why sophisticated institutions do not capitalize on the low

risk-high return anomaly. Indeed, as mentioned before, this anomaly has actually gained force over a period when institutional management in the U.S. has doubled from 30% to over 60% according to the Flow of Funds data presented in **Exhibit 2**.

One issue is why institutional investors do not short the very poor performing top volatility quintile. In the full CRSP sample, this has a simple answer: because the top volatility quintile tends to be small stocks, and they are costly to trade in large quantity both long and especially short, and because the volume of shares available to borrow is limited and borrowing costs are often high. In the large cap sample, the same frictions are present, albeit in considerably smaller measure. But the second and more interesting issue is why institutional investors do not at least overweight the low volatility quintile. We believe the answer involves benchmarking.

A typical contract for institutional equity management contains an implicit or explicit mandate to maximize the “information ratio” relative to a specific, fixed, capitalization-weighted benchmark, without using leverage. For example, if the benchmark is the S&P 500, the numerator of the information ratio is the expected difference between the return earned by the investment manager and the return on the S&P 500. The denominator is the volatility of this return difference, also called the tracking error. The investment manager is expected to maximize this information ratio through stock selection and without using leverage.

This contract is widely used because it has several appealing features. Although the ultimate investor cares more about total risk, not tracking error, it is arguably easier to understand the skill of an investment manager, and the risks taken, by comparing returns to a well-known benchmark. Knowing that each manager will stick at least roughly to a benchmark also helps the ultimate investor keep track of his or her overall risk across many asset classes and mandates.

But these advantages come at a cost. Roll [1992] analyzes the distortions that arise from a fixed benchmark mandate, and Brennan [1993] considers the effect on stock prices. In particular, *a benchmark makes institutional investment managers less likely to exploit the low volatility anomaly*. We lay this out formally in an appendix for the mathematically inclined, but the logic is simple.

In the Sharpe-Lintner CAPM, investors with common beliefs aim to maximize the expected return on their portfolios and minimize volatility. This leads to a simple relationship between risk and return. A stock's expected return equals the risk-free rate plus its beta times the market risk premium:

$$E(R) = R_f + \beta E(R_m - R_f) \quad (1)$$

Now imagine that there is some extra and somewhat irrational demand for high volatility stocks that comes from the preferences for lotteries and the representativeness and overconfidence biases. This extra demand will push up the price of higher risk stocks and thus drive down their expected returns, and vice versa for lower risk stocks.

PROBLEM CASES: LOW ALPHA / HIGH BETA AND HIGH ALPHA / LOW BETA

Will an institutional investor with a fixed benchmark exploit such mispricings? The short answer, which we illustrate with a simple example, is likely no. In fact, in empirically relevant cases, the manager's incentive is to exacerbate mispricings!

Low beta, high alpha. Consider an institutional manager who is benchmarked against the market portfolio. Suppose the expected return on the market is 10% over the risk-free rate and the volatility of the market is 20%. Take a stock with a β of 0.75, and imagine that it is undervalued, with an expected return greater than the CAPM benchmark in Equation (1) by an amount α . Overweighting the stock by a small amount, say 0.1%, will increase the expected active, or benchmark-adjusted, return by approximately $0.1\% \cdot \{\alpha - (1 - \beta) \cdot E(R_m - R_f)\} = 0.1\% \cdot \{\alpha - 2.5\%\}$. The extra tracking error of the portfolio is at least $\sqrt{0.1\%^2 \cdot \{\sigma_m^2 \cdot (1 - \beta)^2\}} = \sqrt{0.1\%^2 \cdot 0.0025}$, the component that comes from having a portfolio β that is not equal to 1.0.

This investment manager would not start overweighting such an undervalued low-beta stock until its α exceeds 2.5% per year. An undervalued stock with an alpha less extreme but still substantial, say 2%, is actually a better candidate for *underweighting*.

A key assumption here is that the institutional manager cannot use leverage. For example, by borrowing 33% of each dollar invested in the low beta stock, the problem of portfolio tracking error is solved, at least from the β component alone. Black [1972] also noted the relevance of a leverage

constraint to a flat return-beta relationship. Similarly, a balanced fund mandate without a fixed leverage constraint could solve this problem. For example, if a balanced fund mandate dictated a beta of 0.5 rather than a fixed 50% of the portfolio in stocks, the manager could choose low risk stocks in place of a greater percentage of the portfolio in low beta fixed income securities. There are also more elaborate solutions to the problem of delegated investment management, e.g., van Binsbergen, Brandt, and Kojien [2008].

In practice, our assumption of a leverage constraint seems reasonable. Very few mutual funds for example allow leverage, and the holdings of mutual funds with ‘balanced’ in their title had an average equity beta of 1.02 in December 2008 using data on U.S. holdings from Lionshares and CRSP betas. This is slightly lower than the average beta of other mutual funds at 1.10, but still above 1.0. (Karceski [2002] reports an average beta of 1.05 in older data from 1984 to 1996.) Moreover, the assets under management of balanced funds are just 2% of the total. While flexible mandates could in principle solve the low risk anomaly, they are rarely used in practice, and even when the explicit contract allows flexibility, investment managers do not overweight low risk stocks. One interpretation is that balanced funds, for example, are implicitly evaluated according to their allocation to equities, not their beta.

Low alpha, high beta. Now consider the case of overvalued high-beta stocks. By the same logic, the manager will not underweight a stock with a β of 1.25, for example, until its α is below -2.5%. And again, the manager becomes part of the problem unless the alpha is very negative—an α of -2%, for example, is still a candidate for *overweighting*.

The logic illustrates that an investment manager with a fixed benchmark and no leverage is best suited to exploit mispricings among stocks with close to market risk, i.e. β near 1. In those cases, managers will have a robust desire to overweight positive α and underweight negative α stocks, enhancing market efficiency. As beta decreases (increases), alpha must increase (decrease) to induce bets in this direction. All of this relates directly to the low risk anomaly. Its essence is that low risk is undervalued relative to high risk. Our logic suggests that this is not surprising in a benchmarked world.

Exhibit 3 gives a feel for just what these anomalies look like to the benchmarked manager. Let us focus on the case of large caps only, a universe of special practical relevance to benchmarked investors, and a perfectly dramatic illustration of the problem. We assume that the benchmark is the CRSP value-

weighted market return over the three major U.S. exchanges. For low volatility portfolios, the Sharpe ratio is reasonably high at 0.39. But the information ratio, the ratio of the excess return over the fixed benchmark to the tracking error, is much less impressive at 0.11. This is even somewhat lower than the information ratio of the high volatility portfolio, even though the cumulative return plot appears considerably more attractive. Results for the other three low risk portfolios give the same message.

Beta and volatility are highly correlated. It is of practical interest to determine which notion of risk is more fundamental to the anomaly. Is beta a noisy measure of idiosyncratic volatility or is volatility a proxy for the underlying beta? It is also of theoretical interest, because our mechanism centers on beta, with total volatility entering the picture only to the extent that portfolios are not sufficiently diversified to prevent idiosyncratic risk from affecting tracking error. In unreported results, we sort on volatility orthogonalized to beta (very roughly, idiosyncratic risk), and on beta orthogonalized to volatility. The results suggest that beta is closer to the heart of the anomaly. In large capitalization stocks, high orthogonalized beta portfolios have the lowest returns just as do high raw beta portfolios. However, large stocks with high orthogonalized volatility actually show higher returns. In other words, beta drives the anomaly in large stocks, while both measures of risk play a role in small stocks. Perhaps this pattern is consistent with the fact that benchmarked managers focus disproportionately on large stocks.

The bottom line of **Exhibit 3** is that a benchmarked institutional fund manager is likely to devote little if any long capital or risk-bearing capacity to exploiting these risk anomalies. Nor is aggressively shorting high risk stocks a particularly appealing strategy. Other anomalies generate far better information ratios over this period. This one will tend to be ignored, and thus the mispricings generated by risk-loving irrational investors survive.

We can also think of all of this in familiar CAPM terms. In a simple equilibrium described in the appendix along the lines of Brennan [1993], and with no irrational investors at all, the presence of delegated investment management with a fixed benchmark will cause the CAPM relationship to fail. In particular, it will be too flat as in **Exhibit 4**:

$$E(R) = (R_f + c) + \beta E(R_m - R_f - c) \quad (2)$$

The constant $c > 0$ depends intuitively on the tracking error mandate of the investment manager (a looser mandate leads to more distortion), and the fraction of asset management that is delegated (more assets lead to more distortion). The pathological regions are the areas between the CAPM and Delegated Management security market lines. For stocks in these regions, the manager will not enforce the CAPM. He or she will be reluctant to overweight low beta, high alpha stocks and will be reluctant to underweight high beta, low alpha stocks. This is consistent with the average mutual fund beta of 1.10 over the last ten years, as mentioned above. Of course, the presence of volatility-preferring, irrational investors only serves to further diminish the risk-return tradeoff.

PUTTING THE PIECES TOGETHER

The combination of irrational investor demand for high volatility and delegated investment management with fixed benchmarks and no leverage flattens the relationship between risk and return. High risk stocks, whether measured by β or σ , have not earned a commensurate return. Low risk stocks have outperformed. Indeed, the empirical results indicate that, over the long haul, the risk-return relationship has not just flattened, but inverted. Yet sophisticated investors are to a large extent sidelined by their mandates of maximizing active return subject to benchmark tracking error.

There is a solid investment thesis going forward for low volatility (and low beta) strategies, as long as fixed benchmark contracts are the dominant form of implicit or explicit contracts between investors and investment management firms. As long as the share of the market held by investment managers remains high or even increases, this anomaly is likelier than not to persist.

WITHIN VERSUS ACROSS MANDATES

There is an additional, more subtle, but interesting prediction from this analysis. Investment managers with fixed benchmarks will be unlikely to exploit mispricings where stocks of different risks have similar returns *within a particular mandate*. But risk and return are more likely to line up *across mandates* if the ultimate investors are thoughtful about asset allocation: for example, between

intermediate and long-term bonds; between government and corporate bonds; between stocks and bonds; and, between large and small company stocks.

Exhibit 5 shows that the CAPM indeed works to some extent across asset classes, in contrast to its long-term performance within the stock market. In other words, as the β rises from 0.05 for intermediate-term government bonds to 1.07 for small company stocks, average returns rise from 7.9% to 13.0%. There is still a small CAPM anomaly in bonds, whereby lower risk asset classes appear to outperform their risk-adjusted benchmarks, also suggesting a possible impact from fixed benchmarks in tactical asset allocation. The returns on small company stocks appear to be an exception, but it is worth noting that this comes from lower β small stocks. Higher β small stocks have underperformed.

SEARCHING FOR LOWER VOLATILITY

One last notable feature of **Exhibits 1 and 3** is compounding. The advantage of a low risk portfolio versus a high volatility portfolio is greater when displayed in compound returns than in average returns. The difference comes from the benefits of compounding a lower volatility monthly series.

Given the power of compounding low volatility returns and the outperformance of low volatility stocks, a natural question is whether we can do even better than the low risk quintile portfolios by taking further advantage of the benefits of diversification. Leaving returns aside, we can do better, if we have useful estimates of not only individual firm volatility, but also the correlations among stocks. A portfolio of two *uncorrelated* but slightly more individually volatile stocks can be even less volatile than a portfolio of two *correlated* stocks each with low volatility.

With this in mind, we construct two minimum variance portfolios that take advantage of finer detail in the covariance matrix. Following the method of Clarke, de Silva and Thorley [2006] in **Exhibit 6**, we use only large caps and a simple five-factor risk model—a realistic and implementable strategy—and compare the returns on two optimized low volatility portfolios to the performance of the lowest quintile sorted by volatility. The second column of **Exhibit 6** uses individual firm estimates of volatility, rather than a simple sort, to form a low volatility portfolio, but sets the correlations among stocks to zero. The third column also uses the covariance terms from the risk model. We are able to reduce the total

volatility of the portfolios from 12.8% with a simple sort to 11.5% in the optimized portfolio. It turns out that this volatility reduction comes entirely from the estimation of correlations, as the diagonal covariance model produces a higher risk portfolio than the simple sort. Moreover, the reduction in volatility comes at little expense in terms of average returns, so the Sharpe ratios are best in the optimized portfolio, as is visually apparent in **Exhibit 7**. These patterns are stable across both halves of our 41-year sample period.

The final column of **Exhibit 6** uses leverage. By relaxing the leverage constraint, the high Sharpe ratio of the low volatility portfolio in the third column can be converted into a respectable information ratio of 0.48. By using leverage to neutralize the portfolio beta, the extra tracking error that comes from focusing on lower beta stocks is reduced to the idiosyncratic component of stock selection. This portfolio produces higher than market returns at market levels of average risk.

THE BEST OF BOTH WORLDS

The majority of stock market anomalies can be thought of as “different returns, similar risks.” Value and momentum strategies, for example, are of this sort; cross-sectional return differences are emphasized, not risk differences. Institutional investment managers are well positioned to take advantage of such anomalies because they can generate high excess return while maintaining average risks, thereby matching their benchmark’s risk and controlling tracking error.

But the low risk anomaly is of a quite different character. Exploiting it involves holding stocks with more or less similar long-term returns, which does not help a typical investment manager’s excess returns, but with different risks, which only increases tracking error. So, even though irrational investors happily overpay for high risk and shun low risk, investment managers are generally not incentivized to exploit this mispricing. And thus, the anomaly persists, and even grown stronger over time.

Our behavioral finance diagnosis also implies a practical prescription. Investors who want to maximize return subject to total risk must incentivize their managers to do just that, by focusing on the benchmark-free Sharpe ratio, not the commonly employed information ratio. For them, our behavioral finance insights are good news, because they suggest that, as long as most of the investing world sticks with standard benchmarks, the advantage will be theirs.

REFERENCES

- Alpert, Marc, H. Raiffa. "A progress report on the training of probability assessors." In D. Kahneman, P. Slovic, A. Tversky, eds., *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press, 1982, pp. 294-305.
- Ang, Andrew, R. Hodrick, Y. Xing, X. Zhang. "The Cross-Section of Volatility and Expected Returns." *Journal of Finance*, 61 (2006), pp. 259-299.
- . "High Idiosyncratic Volatility and Low Returns: International and Further U.S. Evidence." *Journal of Financial Economics*, 91 (2009), pp. 1-23.
- Baker, Malcolm, J. Wurgler. "Investor Sentiment in the Stock Market." *Journal of Economic Perspectives*, 21 (2007), pp. 129-152.
- Barberis, Nicholas, M. Huang. "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices." *American Economic Review* 98 (2008), pp. 2066-2100.
- Barberis, Nicholas, R. Thaler. "A Survey of Behavioral Finance." In G. Constantinides, M. Harris, R. Stulz, eds., *Handbook of the Economics of Finance*, North-Holland.
- Barberis, Nicholas, W. Xiong. "Realization Utility." Working paper, Yale University, 2010.
- van Binsbergen, Jules H., M. Brandt, and R. Koijen. "Optimal Decentralized Investment Management." *Journal of Finance* 63 (2008), pp. 1849-1895.
- Black, Fischer. "Capital Market Equilibrium with Restricted Borrowing." *Journal of Business* 45 (1972), pp. 444-455.
- Black, Fischer, M. C. Jensen, M. Scholes. "The Capital Asset Pricing Model: Some Empirical Tests." In M. C. Jensen, ed., *Studies in the Theory of Capital Markets*. New York: Praeger, 1972, pp. 79-121.
- Blitz, David, P. V. Vliet. "The Volatility Effect: Lower Risk Without Lower Return." *Journal of Portfolio Management* 34 (2007), pp. 102-113.
- Brennan, Michael J. "Agency and Asset Pricing." Working paper, University of California, Los Angeles, 1993.

Clarke, Roger G., H. de Silva, S. Thorley. "Minimum-Variance Portfolios in the U.S. Equity Market." *The Journal of Portfolio Management* 33 (2006), pp. 10-24.

Fama, Eugene F., K. R. French. "The Cross-Section of Expected Stock Returns." *Journal of Finance* 47 (1992), pp. 427-465.

Fischhoff, Baruch, P. Slovic, S. Lichtenstein. "Knowing with Certainty: The Appropriateness of Extreme Confidence." *Journal of Experimental Psychology: Human Perception and Performance* 3 (1977), pp. 552-564.

Haugen, Robert A., A. J. Heins, "Risk and the Rate of Return on Financial Assets: Some Old Wine in New Bottles." *Journal of Financial and Quantitative Analysis* 10 (1975), pp. 775-784.

Kahneman, Daniel, A. Tversky. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47 (1979), pp. 263-291.

Karceski, Jason. "Returns-Chasing Behavior, Mutual Funds, and Beta's Death." *Journal of Financial and Quantitative Analysis* 37 (2002), pp. 559-594.

Kumar, Alok. "Who Gambles in the Stock Market?" *Journal of Finance* 64 (2009), pp. 1889-1933.

Miller, Edward M. "Risk, Uncertainty, and Divergence of Opinion." *Journal of Finance* 32 (1977), pp. 1151-1168.

Mitton, Todd, K. Vorkink. "Equilibrium Underdiversification and the Preference for Skewness." *Review of Financial Studies* 20 (2007), pp. 1255-1288.

Roll, Richard. "A Mean/Variance Analysis of Tracking Error." *Journal of Portfolio Management* 18 (1992), pp. 13-22.

Shefrin, Hersh, M. Statman. "Behavioral Portfolio Theory." *Journal of Financial and Quantitative Analysis* 35 (2000), pp. 127-151.

Shleifer, Andrei. *Inefficient Markets: An Introduction to Behavioral Finance*. Oxford, UK: Oxford UP, 2000.

Tversky, Amos, D. Kahneman. "Judgment Under Uncertainty: Heuristics and Biases." *Science* 185 (1974), pp. 1124-1131.

APPENDIX: DELEGATED PORTFOLIO MANAGEMENT AND THE CAPM

This is a short derivation that follows Brennan [1993]. It shows that delegated portfolio management with a fixed, market benchmark and no leverage will tend to flatten the CAPM relationship, even with no irrational investors, and make low volatility and low beta stocks and portfolios components of an attractive investment strategy. We start with assumptions that are sufficient to derive the CAPM.

1. **Stocks and bonds.** There are stocks $i = 1$ to N with expected returns \mathbf{R} and covariance $\mathbf{\Sigma}$. There is a risk-free bond that returns R_f .
2. **Investors.** There are two representative investors $j = 1, 2$, who are mean-variance utility maximizers over returns with a risk aversion parameter of ν .
3. **Investment strategies.** Each representative investor makes a scalar asset allocation decision a_j between stocks and the risk free asset, and a vector portfolio choice decision \mathbf{w}_j .
 - a. Investor 1 delegates his portfolio choice. Investor 1 allocates a fraction a_1 of his capital to an intermediary who chooses a portfolio \mathbf{w}_1 on investor 1's behalf.
 - b. Investor 2 chooses his own portfolio. Investor 2 allocates a fraction a_2 of his capital to stocks and he chooses a portfolio \mathbf{w}_2 . This can be collapsed without loss of generality to a single choice variable \mathbf{w}_2 . Mean-variance utility maximization means he chooses \mathbf{w}_2 to maximize $E\left(\mathbf{w}_2' \mathbf{R} - R_f \mathbf{1}\right) - \frac{\nu}{2} \mathbf{w}_2' \mathbf{\Sigma} \mathbf{w}_2$.

If there are only investors of type 2, then the CAPM holds in equilibrium:

$$E(R_i - R_f) = \beta_i E(R_m - R_f) \quad (3)$$

If we add investors of type 1 to the model, we need an extra assumption about what intermediaries do. It would be natural to assume, for example, that they have an information advantage. To keep the derivation simple, intermediation here simply involves selecting stocks on behalf of investors of type 1, with the objective of maximizing the information ratio of the portfolio, or maximizing returns subject to a tracking error constraint, which is governed by a parameter γ .

4. **Intermediation.** There is a single intermediary chooses a portfolio \mathbf{w}_1 to maximize

$$E(\mathbf{w}_1 - \mathbf{w}_b)' \mathbf{R} - \gamma (\mathbf{w}_1 - \mathbf{w}_b)' \mathbf{\Sigma} (\mathbf{w}_1 - \mathbf{w}_b), \text{ where } \mathbf{w}_b \text{ are the weights in the market portfolio and } (\mathbf{w}_1 - \mathbf{w}_b)' \mathbf{1} = \mathbf{0}.$$

Investor 1 allocates a fraction a_1 of his capital to an intermediary. The problem now is that the intermediary no longer cares about maximizing the Sharpe ratio for investors of type 1. The intermediary chooses \mathbf{w}_1 to maximize information ratio; the investors of type 2 choose \mathbf{w}_2 to maximize Sharpe ratio; and the two compete to set prices. Note that the budget constraint in the intermediary's objective means that the information ratio must be maximized through stock selection, i.e. without resorting to borrowing or investing in a risk-free asset. We make no claim that this contract is optimal in the sense of van Binsbergen, Brandt, and Kojien [2008], only that it is commonly used in practice.

The market must clear, so that $a_1\mathbf{w}_1 + a_2\mathbf{w}_2 = \mathbf{w}_b$. Substituting the optimal choices of \mathbf{w} into the market clearing condition delivers a flattened version of the CAPM.

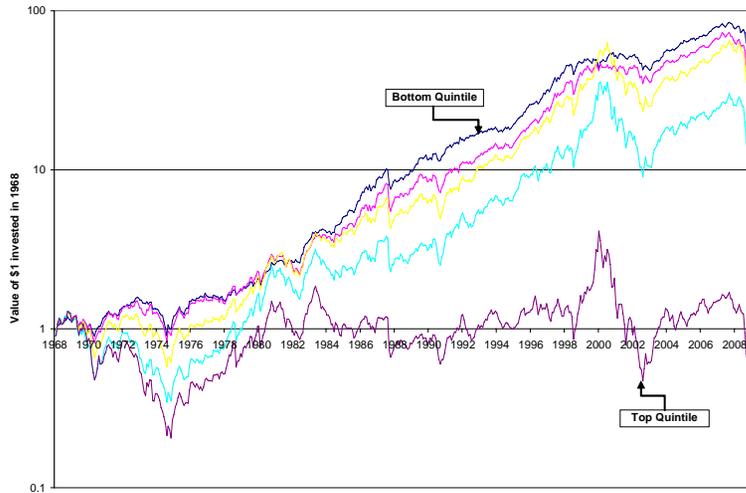
$$E(R_i - R_f) = \beta_i E(R_m - R_f) + c(1 - \beta_i), \text{ where } c = A \frac{a_1 v}{a_2 \gamma + a_1 v} > 0 \quad (4)$$

A is a constant that depends on the equilibrium distribution of risk and return and is positive provided the Sharpe ratio of the minimum variance portfolio is positive. The amount of capital delegated a_1 can easily be endogenized and determined as a function of the risk aversion of investors of type 1, the tracking error mandate γ , and the investment opportunity set, but this does not add much to the intuition of the model. The effects of changes in the other parameters are intuitive. The CAPM relation is especially flat: when γ is small so that there is a loose tracking error mandate; when investors of type 1 delegate a large amount of capital a_1 to the intermediary; when investors of type 2 are risk averse, or when v is large, leading them to stay out of stocks to a greater extent.

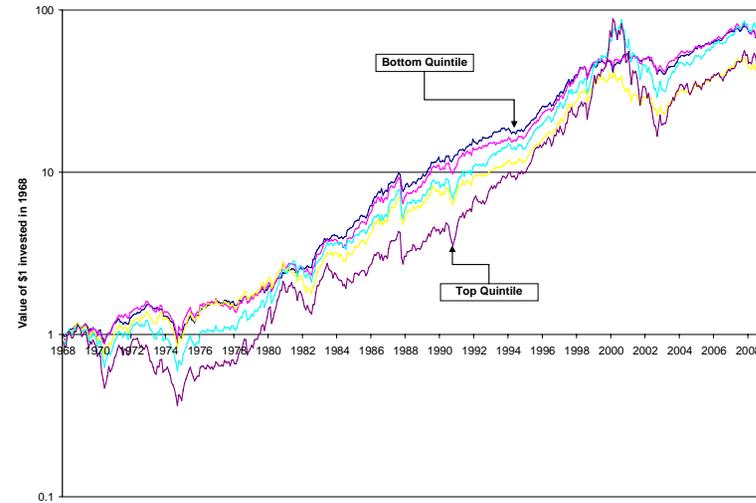
As Brennan [1993] shows, investors of type 2 will specialize in lower volatility stocks. In this example, they are rational, mean-variance utility maximizers and they partially offset the effects of an intermediary who tries to capture improvements in information ratio by holding higher volatility stocks. Introducing a set of irrational individual investors with a preference for high volatility will only exacerbate the flattening of the CAPM. Intermediaries will only start to act as arbitrageurs when the relationship between risk and return is inverted.

Exhibit 1. Returns by volatility quintile. In each month, we sort all publicly traded stocks (Panels A and C) or the top 1000 stocks by market capitalization (Panels B and D) tracked by the Center for Research on Security Prices (CRSP) with at least 24 months of return history into five equal quintiles according to trailing volatility (standard deviation) or beta. Volatility and beta are estimated using up to 60 months of trailing returns. In January 1968, \$1 is invested, according to capitalization weights. At the end of each month, each portfolio is rebalanced, with no transaction costs included. Source: Acadian calculation using data from the Center for Research on Security Prices (CRSP).

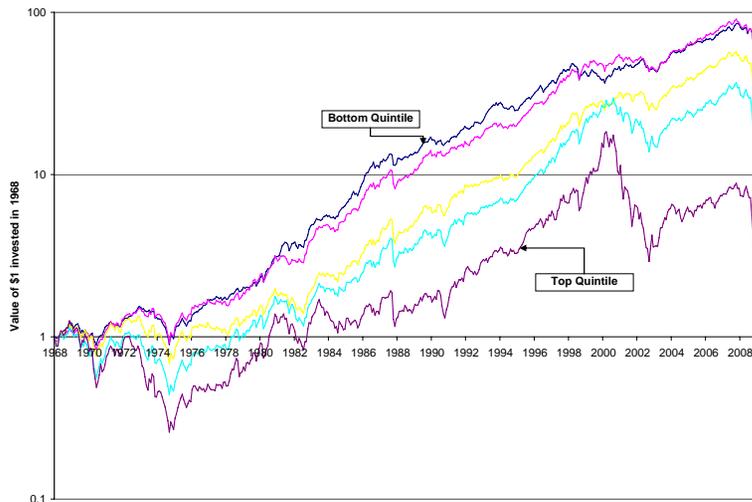
Panel A. All Stocks, Volatility Quintiles



Panel B. Top 1000 Stocks, Volatility Quintiles



Panel C. All Stocks, Beta Quintiles



Panel D. Top 1000 Stocks, Beta Quintiles

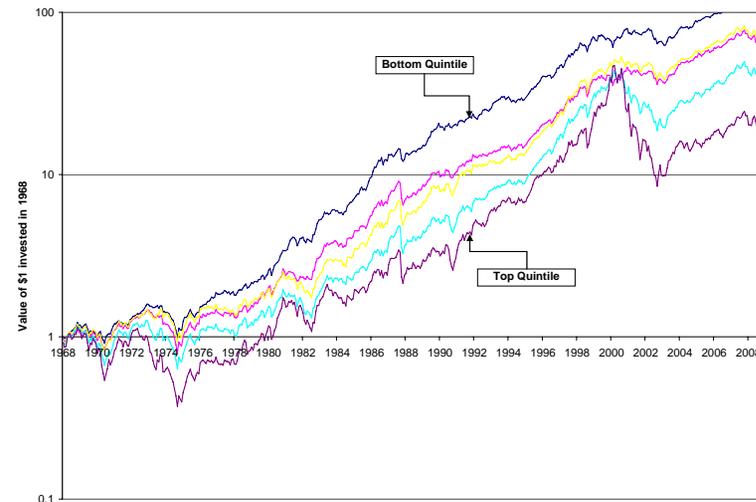


Exhibit 2. Institutional ownership, 1968-2008. Institutional ownership from the Flow of Funds Table L.213. Assets managed by insurance companies (lines 12 and 13), public and private pension funds (lines 14, 15, and 16), open- and closed-end mutual funds (lines 17 and 18), and broker dealers (line 20). Assets under management are scaled by the market value of domestic corporations (line 21).

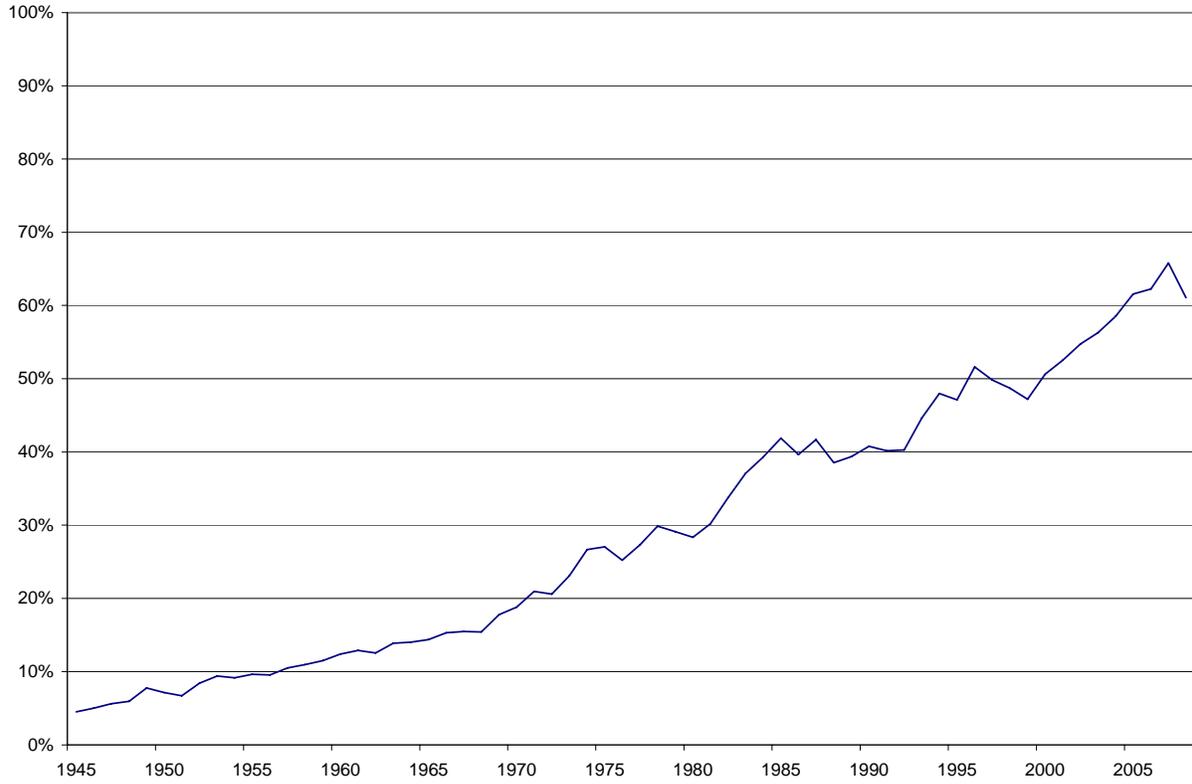


Exhibit 3. Returns by volatility quintile, 1968.01-2008.12. We form portfolios by dividing all publicly traded stocks (first five columns) or the top 1000 stocks by market capitalization (second five columns) tracked by the Center for Research on Security Prices (CRSP) into five equal sized quintiles according to trailing volatility (standard deviation) in Panel A and trailing beta in Panel B. Beta and volatility are measured using up to five years of monthly returns. The return on the market R_m and the risk free rate R_f are taken from Ken French's website. Average returns are monthly averages multiplied by 12. Standard deviation and tracking error are monthly standard deviations multiplied by the square root of 12.

	<i>All Stocks</i>					<i>Top 1000 Stocks</i>				
	Low	2	3	4	High	Low	2	3	4	High
Panel A. Volatility Sorts										
Geometric Average $R_p - R_f$	4.38%	3.37%	2.72%	0.46%	-6.78%	4.20%	4.17%	2.38%	3.64%	2.10%
Average $R_p - R_f$	5.15%	4.75%	5.04%	4.18%	-1.73%	4.94%	5.25%	3.91%	5.82%	5.81%
Standard Deviation	13.10%	16.72%	21.38%	26.98%	32.00%	12.75%	15.15%	17.47%	20.90%	27.02%
Sharpe Ratio	0.39	0.28	0.24	0.16	-0.05	0.39	0.35	0.22	0.28	0.21
Average $R_p - R_m$	1.05%	0.65%	0.94%	0.08%	-5.84%	0.83%	1.15%	-0.20%	1.72%	1.71%
Tracking Error	6.76%	4.59%	7.88%	14.23%	20.33%	7.92%	5.87%	4.59%	7.82%	14.39%
Information Ratio	0.16	0.14	0.12	0.01	-0.29	0.11	0.20	-0.04	0.22	0.12
Beta	0.75	1.01	1.28	1.54	1.71	0.70	0.88	1.06	1.24	1.53
Alpha	2.08%	0.61%	-0.21%	-2.12%	-8.73%	2.08%	1.62%	-0.44%	0.73%	-0.48%
t(Alpha)	2.44	0.85	-0.21	-1.19	-3.28	2.11	1.86	-0.63	0.69	-0.26
Panel B. Beta Sorts										
Geometric Average $R_p - R_f$	4.42%	4.49%	2.99%	1.27%	-2.42%	5.44%	4.02%	3.84%	2.12%	-0.04%
Average $R_p - R_f$	5.07%	5.30%	4.30%	3.36%	1.53%	6.08%	4.95%	5.11%	3.99%	3.40%
Standard Deviation	12.13%	13.39%	16.31%	20.24%	27.77%	12.41%	14.07%	16.22%	19.25%	25.85%
Sharpe Ratio	0.42	0.40	0.26	0.17	0.05	0.49	0.35	0.31	0.21	0.13
Average $R_p - R_m$	0.97%	1.20%	0.20%	-0.74%	-2.58%	1.97%	0.85%	1.01%	-0.11%	-0.70%
Tracking Error	9.74%	7.06%	5.15%	6.25%	14.52%	9.35%	6.51%	4.56%	5.41%	12.45%
Information Ratio	0.10	0.17	0.04	-0.12	-0.18	0.21	0.13	0.22	-0.02	-0.06
Beta	0.60	0.76	0.97	1.23	1.61	0.63	0.81	0.98	1.17	1.51
Alpha	2.60%	2.20%	0.31%	-1.69%	-5.06%	3.49%	1.64%	1.10%	-0.83%	-2.80%
t(Alpha)	2.23	2.39	0.39	-2.13	-2.97	3.05	1.82	1.54	-1.13	-1.90

Exhibit 4. Delegated investment management with a fixed benchmark or R_m flattens the CAPM relationship.

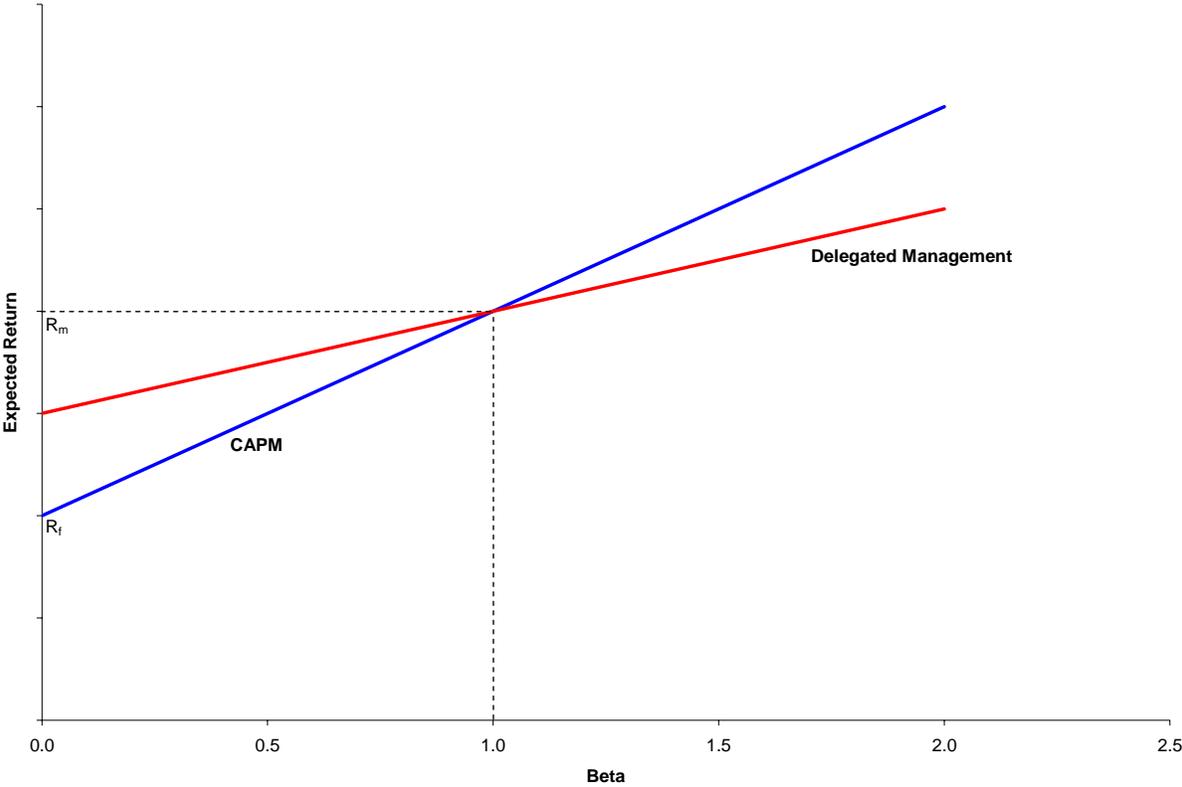


Exhibit 5. Risk and return across asset classes, 1968.01-2008.12. We compute the average return and beta by asset class, using data from Ibbotson Associates. The return on the market R_m (Large Company Stocks) and the risk free rate R_f are taken from Ibbotson Associates. Average returns are monthly averages multiplied by 12.

	Average Return	Sharpe Ratio			CAPM Performance		
		Excess Return	SD	Sharpe	Beta	Alpha	t(Alpha)
Short Government Bonds	5.70%	0.00%	0.02%				
Intermediate Government Bonds	7.88%	2.18%	5.58%	0.39	0.05	1.98%	2.28
Long Government Bonds	8.90%	3.20%	10.54%	0.30	0.14	2.61%	1.61
Corporate Bonds	8.48%	2.77%	9.58%	0.29	0.18	2.01%	1.40
Large Company Stocks	9.86%	4.16%	15.33%	0.27	1.00		
Small Stocks	13.04%	7.34%	21.79%	0.34	1.07	2.88%	1.28

Exhibit 6. A low volatility portfolio versus a portfolio of low volatility stocks, 1968.01-2008.12. We form a minimum variance portfolio of the top 1000 stocks by market capitalization in the CRSP universe using two methods and compare performance to a low volatility sort. The covariance matrix is estimated as in Clarke, de Silva and Thorley [2006]. We limit the individual stock weights to fall between zero and three percent. The third column uses a simple five-factor covariance matrix with a Bayesian shrinkage parameter applied to the correlations, and the second column uses only the diagonal of the covariance matrix. The fourth column levers the third column portfolio to produce an average beta of 1.0. The return on the market R_m and the risk free rate R_f are taken from Ken French's website. Average returns are monthly averages multiplied by 12. Standard deviation and tracking error are monthly standard deviations multiplied by the square root of 12.

	Low Volatility Quintile	Diagonal Only	Full Risk Model	Levered, Full Risk Model
Geometric Average $R_p - R_f$	4.20%	5.26%	4.85%	7.26%
Average $R_p - R_f$	4.94%	6.42%	5.41%	8.82%
Standard Deviation	12.75%	15.93%	11.50%	18.80%
Sharpe Ratio	0.39	0.40	0.47	0.47
Average $R_p - R_m$	0.83%	2.32%	1.31%	4.71%
Tracking Error	7.92%	5.08%	8.73%	9.90%
Information Ratio	0.11	0.46	0.15	0.48
Beta	0.70	0.95	0.61	1.00
Alpha	0.02	0.03	0.03	0.05
t(Alpha)	2.11	3.21	3.01	3.03

Exhibit 7. A low volatility portfolio versus a portfolio of low volatility stocks. The heavy solid line is a minimum variance portfolio of the top 1000 stocks by market capitalization in the CRSP universe using the covariance matrix estimate methodology of Clarke, de Silva and Thorley [2006]. We limit the individual stock weights to fall between zero and three percent. The dashed line is the portfolio of the lowest quintile by trailing volatility. Volatility is measured as the standard deviation of up to 60 months of trailing returns. The think solid line levers the heavy solid line portfolio to produce an average beta of 1.0.

