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

Malcolm Baker, Brendan Bradley, and Jeffrey Wurgler

Contrary to basic finance principles, high-beta and high-volatility stocks have long underperformed low-beta and low-volatility stocks. This anomaly may be partly explained by the fact that the typical institutional investor's mandate to beat a fixed benchmark discourages arbitrage activity in both high-alpha, low-beta stocks and low-alpha, high-beta stocks.

Among the many candidates for the greatest anomaly in finance, a particularly compelling one is the long-term success of low-volatility and low-beta stock portfolios. Over 1968–2008, low-volatility and low-beta portfolios offered an enviable combination of high average returns and small drawdowns. This outcome runs counter to the fundamental principle that risk is compensated with higher expected return. In our study, we applied 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 some particular way. In the context of the low-risk anomaly, we believe that a preference for lotteries and the well-established biases of representativeness and overconfidence lead to a demand for higher-volatility stocks that is not warranted by fundamentals.

The second ingredient is limits on arbitrage, which explain why the "smart money" does not offset the price impact of any irrational demand. With respect to the low-risk anomaly, we examined whether the underappreciated limit on arbitrage is benchmarking. Many institutional investors in a position to offset the irrational

Malcolm Baker is professor of finance at Harvard Business School, research associate at the National Bureau of Economic Research, and senior consultant at Acadian Asset Management, Boston. Brendan Bradley is director of managed volatility strategies at Acadian Asset Management, Boston. Jeffrey Wurgler is professor of finance at New York University Stern School of Business, research associate at the National Bureau of Economic Research, and senior consultant at Acadian Asset Management, Boston.

demand for risk have fixed-benchmark mandates (typically capitalization weighted), which, by their nature, discourage investments in low-volatility stocks. Drawing out the implications of Brennan's (1993) model of agency and asset prices, we looked at whether traditional fixed-benchmark mandates with a leverage constraint cause institutional investors to pass up the superior risk–return trade-off of low-volatility portfolios; we also examined the appropriateness of a leverage constraint assumption. Rather than being a stabilizing force on prices, the typical institutional contract for delegated portfolio management could increase the demand for higher-beta investments.

Other researchers have attempted to explain the low-risk anomaly on the basis of behavioral elements. For example, Karceski (2002) pointed out that mutual fund investors tend to chase returns over time and across funds, possibly because of an extrapolation bias. These forces make fund managers care more about outperforming during bull markets than underperforming during bear markets, thus increasing their demand for high-beta stocks and reducing their required returns. In our study, we placed the irrationality elsewhere and focused on distortions introduced by benchmarking. Nevertheless, his model's predictions appear to complement our own, and the mechanisms could certainly work simultaneously.
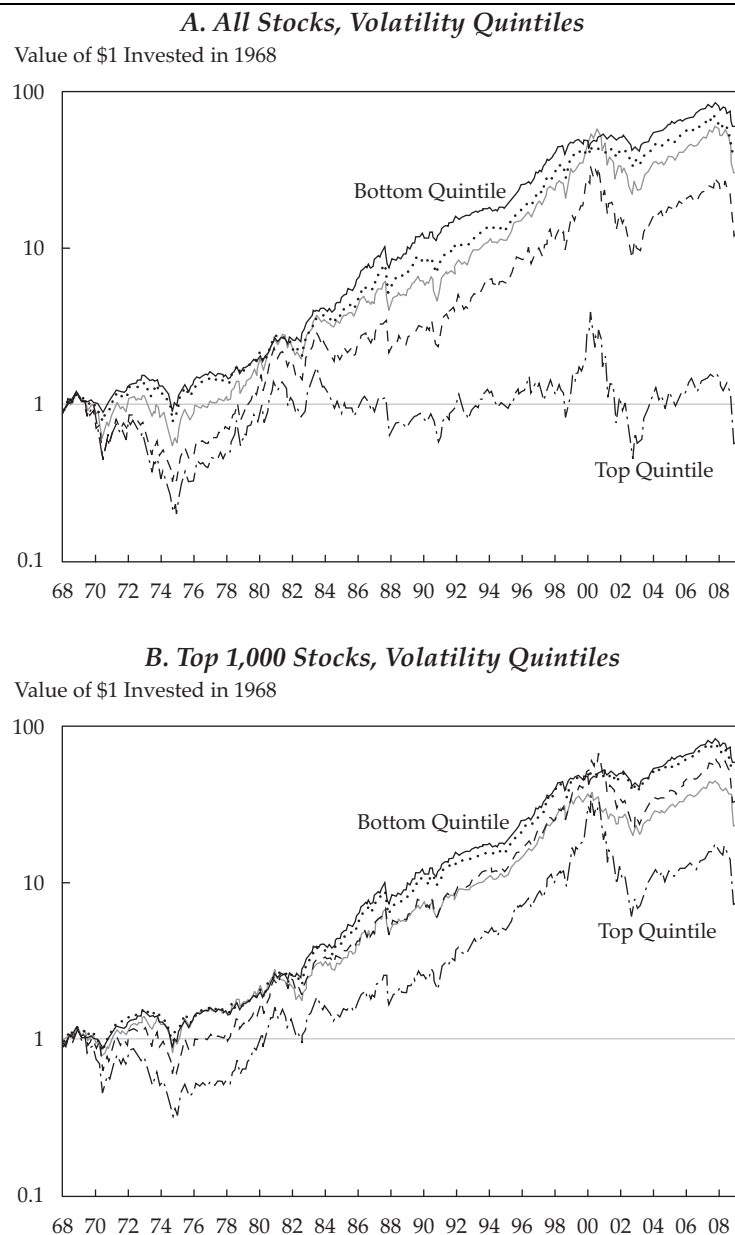
The Low-Risk Anomaly

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

We obtained data for January 1968–December 2008 (a span of 41 years) from CRSP. We sorted stocks into five groups for each month according to either five-year trailing total volatility or trailing beta—thus using data going back to January 1963—and tracked the returns on these portfolios. We also restricted the investing universe to the top 1,000 stocks by market capitalization. **Figure 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 outperformed high risk over the period. Panel A shows that a dollar invested in the lowest-volatility portfolio in January
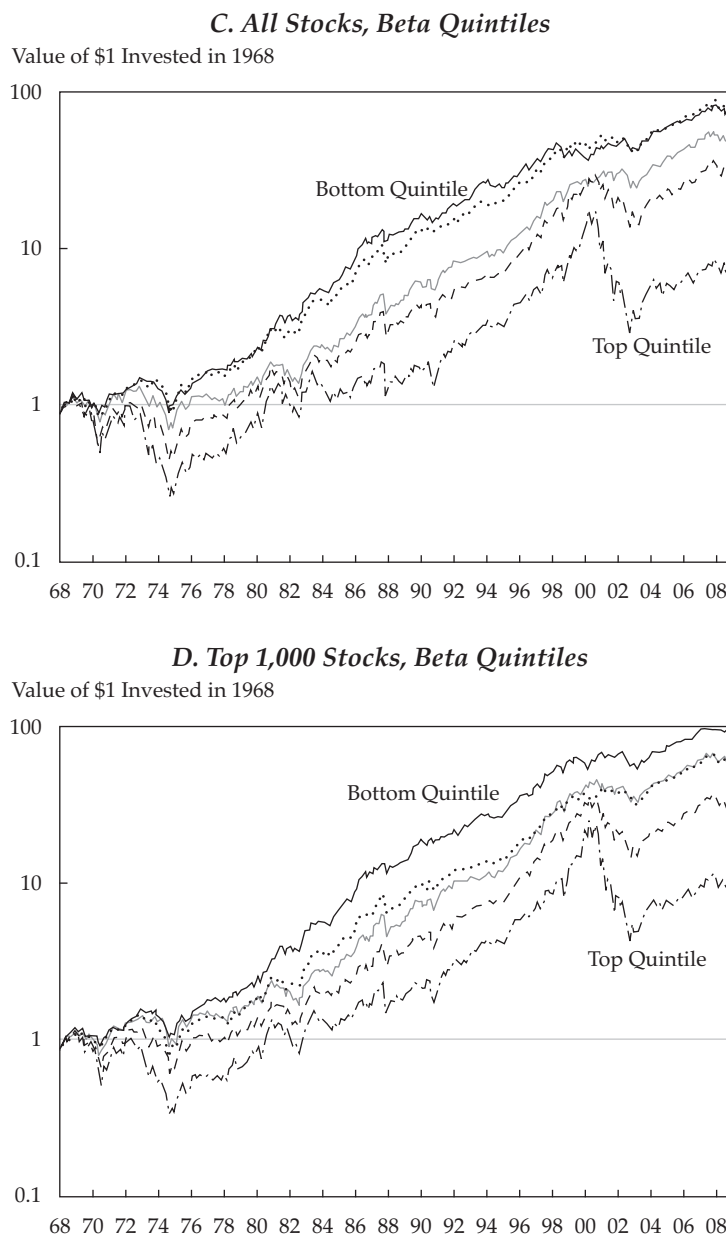
1968 increased to $59.55. Over this period, inflation eroded the real value of a dollar to about 17 cents, meaning that the low-risk portfolio produced a gain of $10.12 in real terms. Contrast this performance with that of the highest-volatility portfolio. A dollar invested there was worth 58 cents at the end of December 2008, assuming no transaction costs. Given the declining value of the dollar, the real value of the high-volatility portfolio declined to less than 10 cents—a 90 percent decline in real terms! Remarkably, an investor who aggressively pursued high-volatility stocks over the last four decades would have borne almost a total loss in real terms.

Figure 1.   Returns by Volatility and Beta Quintile, January 1968–December 2008

A. All Stocks, Volatility Quintiles

Value of $1 Invested in 1968

B. Top 1,000 Stocks, Volatility Quintiles

Value of $1 Invested in 1968

(continued)

**Figure 1.   Returns by Volatility and Beta Quintile, January 1968–
December 2008** (continued)

C. All Stocks, Beta Quintiles

Value of $1 Invested in 1968

D. Top 1,000 Stocks, Beta Quintiles

Value of $1 Invested in 1968

Notes: For each month, we sorted all publicly traded stocks (Panels A and C) and the top 1,000 stocks by market capitalization (Panels B and D) tracked by CRSP (with at least 24 months of return history) into five equal quintiles according to trailing volatility (standard deviation) and beta. In January 1968, $1 is invested according to capitalization weights. We estimated volatility and beta by using up to 60 months of trailing returns (i.e., return data starting as early as January 1963). At the end of each month, we rebalanced each portfolio, excluding all transaction costs.

Source: Acadian calculation with data from CRSP.

Panel C considers beta as the measure of risk. A dollar invested in the lowest-beta portfolio grew to $60.46 ($10.28 in real terms), and a dollar invested in the highest-beta portfolio grew to $3.77 (64 cents in real terms). Like the high-volatility investor, the high-beta investor failed to recover his dollar in real terms and underperformed his "conservative" beta neighbor by 964 percent.

Although almost all mispricings were stronger for small companies than for large companies, the low-risk anomaly was dramatic even for large companies. A dollar invested in low-volatility large caps grew to $53.81 over 41 years, whereas a dollar invested in high-volatility large caps grew to $7.35. For beta, the numbers are $78.66 and $4.70, respectively.

Finally, as if this puzzle were not bad enough, other facts only compound it.

- The low-risk portfolios' paths to their higher dollar values have been much smoother than those of the high-risk portfolios. They are as advertised: genuinely lower risk.

- Motivated by the analysis of Pettengill, Sundaram, and Mathur (1995), we repeated the analysis separately for months in which market returns were above or below their median. Consistent with Pettengill et al. (1995), we found that high-beta stocks earned higher (lower) total returns than did low-beta stocks in up (down) markets, but on a capital asset pricing model (CAPM) market-adjusted basis, the low-beta anomaly was present in both environments. That low beta is high alpha is a robust historical pattern.

- The transaction costs of monthly rebalancing were substantially higher for the high-volatility portfolio, which means that the relative performance in Figure 1 is understated. We found similar results under yearly rebalancing; constraining turnover did not have a material effect.

- With the exception of the technology bubble, the return gap has, if anything, widened a bit since 1983—a period in which institutional investment managers became progressively more numerous, better capitalized, and more quantitatively sophisticated. Karceski (2002) also noted this trend.

These results are not 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 the relationship was not merely flat in their sample period but was actually inverted. Extending this analysis through 1990, Fama and French (1992) also found that the relationship was flat, prompting many to conclude that beta was dead. More recently, Ang, Hodrick, Xing, and Zhang (2006, 2009) drew renewed attention to these results, finding that high-risk stocks have had "abysmally low average returns" (2006, p. 296) in longer U.S. samples and in international markets. Blitz and van Vliet (2007) provided a detailed analysis of the volatility anomaly and demonstrated its robustness across regions and to controls for size, value, and momentum effects. Frazzini and Pedersen (2010) documented that low-risk securities have high risk-adjusted returns in global stock, Treasury, credit, and futures markets. Bali, Cakici, and Whitelaw

(forthcoming) investigated a measure of lottery-like return distributions, which is highly correlated with other risk measures, and found that it too is associated with poor performance. All told, the evidence for a risk–return trade-off 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. 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.

Despite superior computational firepower, however, the new models face an uphill battle. After all, the task is to prove that high-volatility and high-beta stocks are *less* risky. 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 that began in the fall of 2008. Investors appeared to be paying an insurance premium only to lose even more whenever the equity market (and often the economy) underwent a meltdown.

We believe that the long-term outperformance of low-risk portfolios is perhaps the greatest anomaly in finance. Large in magnitude, it challenges the basic notion of a risk–return trade-off.

A Behavioral Explanation

In our study, we hypothesized two drivers of these results: (1) less than fully rational investor behavior and (2) underappreciated limits on arbitrage. The combination of these two forces is the basic framework of behavioral asset pricing, as laid out in such surveys as Shleifer (2000), Barberis and Thaler (2003), and Baker and Wurgler (2007). We explored a new combination that could explain the low-risk anomaly.

The Irrational Preference for High Volatility. The preference for high-volatility stocks derives from the biases that afflict the individual investor. We examined three such biases.

■ *Preference for lotteries.* Would you take a gamble with a 50 percent chance of losing $100 versus a 50 percent chance of winning $110? Most people would say no. Despite the positive expected payoff, the possibility of losing $100 is enough to deter participation, even when $100 is trivial compared with wealth or income.

Kahneman and Tversky (1979) called this behavior "loss aversion." Taken on its own, loss aversion suggests that investors would shy away from volatility for fear of realizing a loss. But something strange happens as the probabilities shift. Now suppose that you are offered a gamble with a near-certain chance of losing $1 and a small (0.12 percent) chance of winning $5,000. As in the first example, this gamble has a positive expected payoff of around $5. In this case, however, most people take the gamble. Gambling on lotteries and roulette wheels, which have negative expected payoffs, is a manifestation of this tendency.

To be precise, this behavior is more about positive skewness, whereby large positive payoffs are more likely than large negative ones, than it is about volatility. But Mitton and Vorkink (2007) reminded us that volatile individual stocks, with limited liability, are also positively skewed. Buying a low-priced, volatile stock is like buying a lottery ticket: There is a small chance of its doubling or tripling in value in a short period and a much larger chance of its declining in value. Boyer, Mitton, and Vorkink (2010) argued that volatility is a proxy for expected skewness. Kumar (2009) found that some individual investors do show a clear preference for stocks with lottery-like payoffs, measured as idiosyncratic volatility or skewness. Modeling this preference with the cumulative prospect theory approach in Tversky and Kahneman (1992), Barberis and Huang (2008) examined an array of circumstantial evidence that volatile stocks are overvalued because of a lottery preference. Blitz and van Vliet (2007) tied the preference for lottery tickets to the behavioral portfolio theory in Shefrin and Statman (2000). Barberis and Xiong (2010) offered a preference-based explanation of the volatility effect that is separate from a skewness preference in which investors obtain utility from *realizing* gains and losses on risky assets, not from paper gains and losses.

■ *Representativeness.* The classic way to explain representativeness is with an experiment from Tversky and Kahneman (1983). 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" (p. 297). They then asked subjects which was 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 chose B shows that probability theory and Bayes' rule are not ingrained skills. As a proper subset of A, B is less likely than A but seems more "representative" of Linda.

What does this experiment 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. On the one hand, the layman, trying to think of great investments—perhaps buying Microsoft Corporation and Genzyme Corporation at their IPOs in 1986—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. Largely ignoring the high base rate at which small, speculative investments fail, the layman is inclined to overpay for volatile stocks.

The quant, on the other hand, analyzes the full sample of such stocks as Microsoft and Genzyme, as shown in Figure 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 (see Fischhoff, Slovic, and Lichtenstein 1977; Alpert and Raiffa 1982). For example, experimenters ask subjects to estimate the population of Massachusetts and to provide a 90 percent confidence interval around their answer. Most people form confidence intervals that are too narrow. And the more obscure the question— Bhutan instead of Massachusetts—the more this calibration deteriorates.

Valuing stocks involves the 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, sticking with the false precision of their estimates. The extent of disagreement is likely higher for more uncertain outcomes—such as the returns on high-volatility stocks. Cornell (2009) viewed overconfidence as an important part of the demand for volatile stocks.

The careful theorist will note that one extra assumption is needed to connect overconfidence— or, more generally, differences of opinion—to the demand for volatile stocks. In markets, pessimists must act less aggressively than optimists. Investors 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 wide range of opinions will have more optimists among their shareholders and will sell for higher prices, leading to lower future returns. Diether, Malloy, and Scherbina (2002) provided empirical support for this idea.

Benchmarking as a Limit on 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 earlier, this anomaly gained force over a period when institutional management in the United States went from 30 percent to 60 percent (**Figure 2**).

Figure 2. Institutional Ownership, 1968–2008
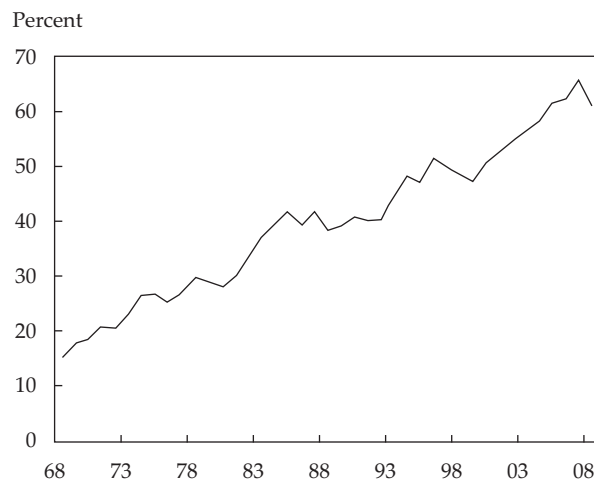
Note: This figure depicts data on institutional ownership from the Federal Reserve's Flow of Funds Accounts of the United States, Table L.213: assets managed by insurance companies (lines 12 and 13), public and private pension funds (lines 14, 15, and 16), open- and closed-end mutual funds (lines 17 and 18), and broker/dealers (line 20); assets under management are scaled by the market value of domestic corporations (line 21).

One issue is why institutional investors do not short the very poor-performing top volatility quintile. For the full CRSP sample, this question has a simple answer: The top volatility quintile tends to be small stocks, which are costly to trade in large quantities—both long and, especially, short—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, which begs the second and more interesting question: Why do institutional investors

not at least overweight the low-volatility quintile? We believe that 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 Index, the numerator of the information ratio (IR) 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 IR through stock selection and without using leverage. Sensoy (2009) reported that 61.3 percent of U.S. mutual fund assets are benchmarked to the S&P 500 and 94.6 percent are benchmarked to some popular U.S. index. Under current U.S. SEC rules, all mutual funds must select a benchmark and show fund returns versus the benchmark in their prospectuses. In this segment of the asset management industry, however, the use of the IR is less formalized.

This contract is widely used because it has several appealing features. Although the ultimate investor cares more about total risk than tracking error, it is arguably easier to understand the skill of an investment manager—and the risks taken—by comparing returns with those of a well-known benchmark. Knowing that each manager will at least roughly stick to a benchmark also helps the ultimate investor keep track of the overall risk across many asset classes and mandates.

But these advantages come at a cost. Roll (1992) analyzed the distortions that arise from a fixed-benchmark mandate, and Brennan (1993) considered the effect on stock prices. Cornell and Roll (2005) developed a similar model. In particular, *a benchmark makes institutional investment managers less likely to exploit the low-volatility anomaly*. We lay this model out formally in Appendix A, 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 goal leads to the famously simple relationship between beta risk and return. A stock's expected return equals the risk-free rate plus its beta times the market risk premium:

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

Now imagine some extra demand for high-volatility stocks. This demand will push up the price of higher-risk stocks and drive down their expected returns, and vice versa for lower-risk stocks.

Problem Cases: Low Beta/High Alpha and Low Alpha/High Beta

An institutional investor with a fixed benchmark is surprisingly unlikely to exploit such mispricings. In fact, in empirically relevant cases, the manager's incentive is to exacerbate them.

Low Beta/High Alpha. Consider an institutional manager who is benchmarked against the market portfolio. Suppose that the expected return on the market is 10 percent more than the risk-free rate and the volatility of the market is 20 percent. 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 percent—will increase the expected active, or benchmark-adjusted, return by approximately

$$0.1\%\left\{\left[\alpha-(1-\beta)\right]\left[E\left(R_m-R_f\right)\right]\right\}=0.1\%\left(\alpha-2.5\%\right).$$

The extra tracking error of the portfolio is at least

$$\sqrt{0.1\%^2\left[\sigma_m^2(1-\beta)^2\right]}=\sqrt{0.1\%^2(0.0025)},$$ the component that comes from having a portfolio β that is not equal to 1.

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

A key assumption here is that the manager cannot use leverage. By borrowing 33 percent of each dollar invested in the low-beta stock, the problem of portfolio tracking error is solved, at least with respect to the β component. 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 percent 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 Koijen 2008).

Therefore, our assumption of a leverage constraint deserves consideration. We believe that it is a reasonable assumption for a large portion of the asset management industry. Although precise statistics are hard to come by, conventional wisdom says that few mutual funds use leverage. We spot-checked and found that the five largest active domestic equity mutual funds did not use any leverage as of 1 July 2010. The Investment Company Act of 1940 prohibits mutual funds from using more than 33 percent leverage; our example assumes that funds use this statutory maximum.

Using data on U.S. holdings from LionShares and CRSP betas, we found that mutual funds with *balanced* in their names had an average equity beta of 1.02 in December 2008. This number is slightly lower than the average beta of other mutual funds (1.10) but is still above 1. (Using older data [1984–1996], Karceski [2002] reported an average beta of 1.05.) In addition, the assets under management of balanced funds were only 2 percent of the total.

Closed-end funds use leverage, but they are a small portion of total assets under management. Moreover, Anand (2009) reported that only 2 funds (out of the 18 he considered) used substantial leverage and had assets greater than $100 million. He also found that most of the funds used a 130/30 strategy, which typically involves leverage through borrowed stock, not bonds, and so the same benchmarking challenges arise in attempting to exploit the low-volatility anomaly. We are unaware of any comprehensive tabulation of institutional mandates—in either benchmarking or leverage.

Certainly, some strategies allow the flexibility to take advantage of the anomaly without running into the benchmark limit on arbitrage. These include maximum Sharpe ratio, managed volatility, balanced, and a variety of hedge fund strategies. Although data on the total assets managed under these strategies are unavailable, anecdotal evidence suggests that the numbers are relatively modest. Even when the explicit contract allows flexibility, investment managers do not overweight low-risk stocks. One possible interpretation is that balanced funds, for example, are implicitly evaluated according to their allocation to equities, not their beta. As mentioned earlier, Sensoy (2009) reported that almost all actively managed U.S. equity mutual funds are benchmarked to an S&P or Russell index.

In the end, our model is a substantial simplification of an enormously heterogeneous market, but the assumption of a leverage constraint seems likely to be a reasonable approximation. The documentable amount of assets under management that use leverage or that can use leverage is small relative to the market caps of the stocks involved in the anomaly—which, in some sense, is the total capitalization of the stock market—and is small relative to the amount of capital that would be required not only to *flatten* the relationship between risk and expected return but also to *reverse* it, as traditional finance theory would prefer.

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 percent. And again, the manager becomes part of the problem unless the alpha is very negative—an α of –2 percent, 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., a β near 1). In those cases, managers will have a robust desire to overweight positive-alpha stocks and underweight negative-alpha stocks, thus enhancing market efficiency. As beta decreases (increases), alpha must increase (decrease) to induce bets in that direction. All of this relates directly to the low-risk anomaly, whose essence is that low risk is undervalued relative to high risk. This finding is not surprising in a benchmarked world.

Table 1 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 for the three major U.S. exchanges. For low-volatility portfolios, the Sharpe ratio is reasonably high, at 0.38. But the IR—the ratio of the excess return over the fixed benchmark to the tracking error—is much less impressive, at 0.08. The results for the other three low-risk portfolios offer a similar message.

Beta and volatility are highly correlated. Determining which notion of risk is more fundamental to the anomaly is of practical interest. 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 sorted on volatility orthogonalized to beta (roughly, idiosyncratic risk) and on beta orthogonalized to volatility.

The results suggest that beta is closer than volatility to the heart of the anomaly. For large-cap stocks, high-orthogonalized-beta portfolios have the lowest returns, just as high-raw-beta portfolios do. But large stocks with high orthogonalized volatility actually show higher returns. In other words, beta drives the anomaly in large stocks, but both measures of risk play a role in small stocks. This pattern is consistent with the fact that benchmarked managers focus disproportionately on large stocks.

The main point of Table 1 is that a benchmarked institutional fund manager is likely to devote little long capital or risk-bearing capacity to exploiting these risk anomalies. Nor is aggressively shorting high-risk stocks a particularly appealing strategy. Other anomalies generated far better IRs over this period. Using data from Ken French's website,[1] we found (in unreported results) an IR of 0.51 for a simple, long-only, top quintile value strategy over the same period (1968–2008); the IR of a simple, long-only, top quintile momentum strategy over the same period was 0.64. These vanilla quantitative strategy IRs suggest that the IR of a long-only strategy in low-volatility stocks is unappealing at 0.08 to 0.17. It will not draw much risk capital, and thus, the mispricings are likely to survive.

We can also think of all this in familiar CAPM terms. In a simple equilibrium described in Appendix A along the lines of Brennan (1993), 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 shown in **Figure 3**:

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

The constant, $c > 0$, depends intuitively on the tracking error mandate of the investment manager (a looser mandate leads to more distortion) and on the fraction of asset management that is delegated (more assets increase distortion). The pathological regions are the areas between the CAPM and the delegated management security market lines. For stocks in these regions, the manager will not enforce the CAPM and will be reluctant to overweight low-beta, high-alpha stocks and to underweight high-beta, low-alpha stocks. This finding is consistent with the average mutual fund beta of 1.10 over the last 10 years.

The presence of volatility-preferring, irrational investors serves only to further diminish the risk–return trade-off. As a theoretical matter, behavioral biases are not needed for benchmarking to flatten the CAPM relationship. We included this element in the discussion because behavioral biases are a fact, and so including them allows for a more accurate description of the phenomenon. In addition, we would be the first to acknowledge that benchmarking per se is unlikely to generate the full magnitude of the low-risk anomaly, whereby the risk–return relationship is inverted. Time variation in the strength of behavioral biases, attached to bubbles and crashes, may also help explain some of the time variation in the returns to risky versus safe securities; Baker and Wurgler (2007) studied fluctuations in investor sentiment.

Table 1.   Returns by Volatility and Beta Quintile, January 1968–December 2008

	All Stocks					Top 1,000 Stocks				
	Low	2	3	4	High	Low	2	3	4	High
A. Volatility sorts										
Geometric average $R_p - R_f$	4.38%	3.37%	2.72%	0.46%	–6.78%	4.12%	4.03%	2.06%	2.81%	–0.82%
Average $R_p - R_f$	5.15%	4.75%	5.04%	4.18%	–1.73%	4.86%	5.12%	3.60%	5.02%	2.95%
Standard deviation	13.10%	16.72%	21.38%	26.98%	32.00%	12.74%	15.15%	17.48%	20.86%	27.13%
Sharpe ratio	0.39	0.28	0.24	0.16	–0.05	0.38	0.34	0.21	0.24	0.11
Average $R_p - R_m$	1.05%	0.65%	0.94%	0.08%	–5.84%	0.62%	0.88%	–0.64%	0.78%	–1.29%
Tracking error	6.76%	4.59%	7.88%	14.23%	20.33%	7.45%	5.54%	4.53%	7.91%	14.95%
Information ratio	0.16	0.14	0.12	0.01	–0.29	0.08	0.16	–0.14	0.10	–0.09
Beta	0.75	1.01	1.28	1.54	1.71	0.70	0.88	1.06	1.24	1.54
Alpha	2.08%	0.61%	–0.21%	–2.12%	–8.73%	2.00%	1.49%	–0.76%	–0.07%	–3.36%
t(Alpha)	2.44	0.85	–0.21	–1.19	–3.28	2.03	1.70	–1.07	–0.07	–1.84
B. Beta sorts										
Geometric average $R_p - R_f$	4.42%	4.49%	2.99%	1.27%	–2.42%	5.09%	3.75%	3.44%	1.46%	–1.89%
Average $R_p - R_f$	5.07%	5.30%	4.30%	3.36%	1.53%	5.74%	4.69%	4.72%	3.35%	1.56%
Standard deviation	12.13%	13.39%	16.31%	20.24%	27.77%	12.40%	14.07%	16.24%	19.27%	25.95%
Sharpe ratio	0.42	0.40	0.26	0.17	0.05	0.46	0.33	0.29	0.17	0.06
Average $R_p - R_m$	0.97%	1.20%	0.20%	–0.74%	–2.58%	1.50%	0.45%	0.48%	–0.89%	–2.68%
Tracking error	9.74%	7.06%	5.15%	6.25%	14.52%	8.83%	6.13%	4.31%	5.70%	13.02%
Information ratio	0.10	0.17	0.04	–0.12	–0.18	0.17	0.07	0.11	–0.16	–0.21
Beta	0.60	0.76	0.97	1.23	1.61	0.63	0.81	0.98	1.18	1.52
Alpha	2.60%	2.20%	0.31%	–1.69%	–5.06%	3.16%	1.38%	0.70%	–1.47%	–4.66%
t(Alpha)	2.23	2.39	0.39	–2.13	–2.97	2.77	1.53	0.99	–2.02	–3.15

Notes: For each month, we formed portfolios by sorting all publicly traded stocks (first five columns) and the top 1,000 stocks by market capitalization (second five columns) tracked by CRSP into five equal-sized quintiles according to trailing volatility (standard deviation) for Panel A and trailing beta for Panel B. We estimated volatility and beta by using up to 60 months of trailing returns (i.e., return data starting as early as January 1963). The return on the market, R_m, and the risk-free rate, R_f, are from Ken French's website (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html). The information ratio uses the market return for the relevant universe, all stocks in the first five columns and the top 1,000 stocks in the last five columns. Average returns are monthly averages multiplied by 12. Standard deviation and tracking error are monthly standard deviations multiplied by the square root of 12.

Figure 3. Delegated Investment Management with a Fixed Benchmark: R_m **Flattens the CAPM Relationship**
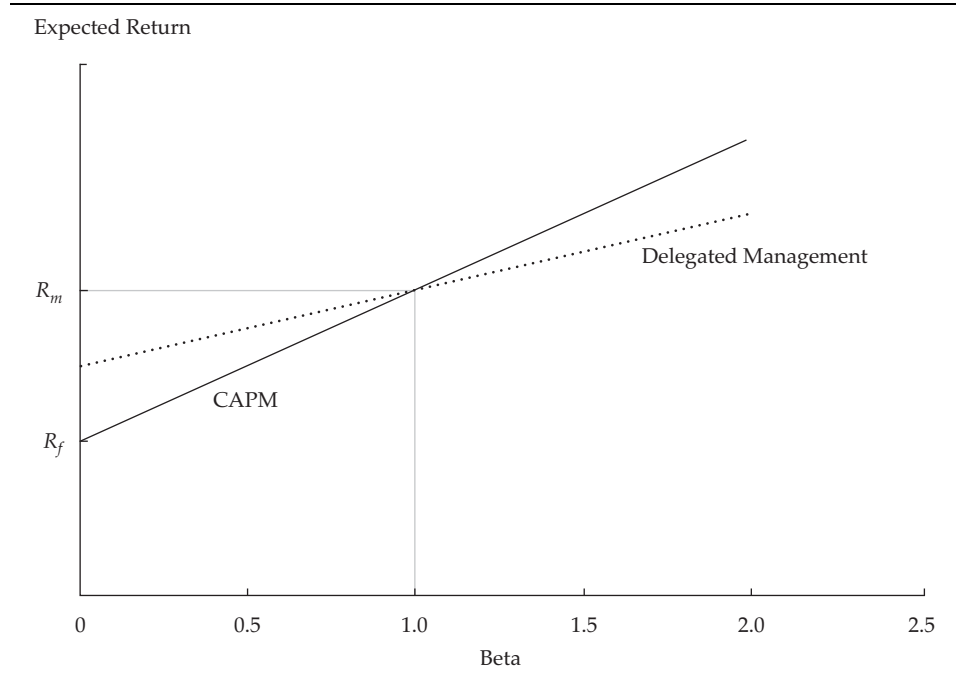
Putting the Pieces Together

To summarize, 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. Indeed, the empirical results suggest that, over the long haul, the risk–return relationship has not merely been flattened but inverted. Yet sophisticated investors are, to a large extent, sidelined by their mandates to maximize active returns subject to benchmark tracking error.

Unfortunately, conducting a direct test of our proposed mechanism is difficult. Instead, we have presented evidence that is consistent with it. We consider the process-of-elimination findings of Ang et al. (2009), who ruled out several potential explanations, to be supportive evidence. To our knowledge, the most direct evidence for our mechanism is provided in Brennan and Li (2008), following up on the framework of Brennan (1993), which we also used. Brennan and Li found evidence that beta with the idiosyncratic component of the S&P 500 should have a negative payoff, all else being equal, consistent with investment managers' attempting to minimize tracking error by holding such stocks. Brennan and Li did not connect their results to the low-risk anomaly.

Another testable prediction is that as benchmarking has increased, the low-risk anomaly should likewise have become more severe. In unreported results, we found that this prediction is directionally correct and, depending on the sample, marginally statistically significant. For example, for the top 500 capitalization stocks in the full CRSP sample (1931–2008), we found that the relative return of low-minus-high volatility may have increased by 1 or 2 bps a year. But this test is not a powerful one. Return data are quite noisy, partly because the preference for volatility has varied with time. For example, the internet bubble, which focused on high-volatility stocks, would have swamped the effect of benchmarks as a limit on arbitrage over that period. We hope that future research will develop more powerful tests of our proposed mechanism.

From a practitioner perspective, the takeaway is that there is a solid investment thesis for low-volatility (and low-beta) strategies. If our explanation is valid, this thesis will be the case so long as fixed-benchmark contracts remain pervasive and the share of the market held by investment managers remains high. There is no reason to expect that the anomaly will go away any time soon.

Within vs. Across Mandates

There is an additional, more subtle prediction that we can test empirically. Investment managers with fixed benchmarks may not exploit mispricings when stocks of different risks have similar returns *within a particular mandate*. But risk and return are 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-cap and small-cap stocks.

Table 2 shows that the CAPM does indeed work, 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 percent to 13.0 percent. A small CAPM anomaly in bonds still exists, whereby lower-risk asset classes appear to outperform their risk-adjusted benchmarks, which also suggests a possible impact from fixed benchmarks in tactical asset allocation. Although the returns on small-company stocks appear to be an exception, note that these returns are on lower-beta small stocks. Higher-beta small stocks have underperformed.

Searching for Lower Volatility

One last notable feature of both Figure 1 and Table 1 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. Returns aside, we can do better if we have useful estimates of not only individual company 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 with low volatility.

With that in mind, we constructed two minimum-variance portfolios that took advantage of finer detail in the covariance matrix. Following the method of Clarke, de Silva, and Thorley (2006), we used only large caps and a simple five-factor risk model—a realistic and implementable strategy—and compared the returns on two optimized low-volatility portfolios with the performance of the lowest quintile sorted by volatility (**Table 3**). In the second column of Table 3, we used individual company estimates of volatility, rather than a simple sort, to form a low-volatility portfolio but set the correlations among stocks to zero. In the third column, we also used the covariance terms from the risk model. We were able to reduce the total volatility of the portfolios from 12.7 percent with a simple sort to 11.5 percent in the optimized portfolio. This volatility reduction comes entirely from the estimation of correlations because the diagonal covariance model produces a higher-risk portfolio than does the simple sort. Moreover, because the reduction in volatility comes at no expense in terms of average returns, the Sharpe ratios are best in the optimized portfolio, as is visually apparent in **Figure 4**. These patterns are stable across both halves of our 41-year sample period.

The final column of Table 3 concerns leverage. With the leverage constraint relaxed, the high Sharpe ratio of the low-volatility portfolio in the third column can be converted into a respectable IR of 0.45. With 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.

Table 2. Risk and Return across Asset Classes, January 1968–December 2008

|		Sharpe Ratio			CAPM Performance		
	Average Return	Excess Return	Std. Dev.	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-company stocks	13.04	7.34	21.79	0.34	1.07	2.88	1.28

Notes: Using data from Ibbotson Associates, we computed the average return and beta by asset class. The return on the market, R_m (large-company stocks), and the risk-free rate, R_f, are from Ibbotson Associates. Average returns are monthly averages multiplied by 12.

Table 3. A Low-Volatility Portfolio vs. a Portfolio of Low-Volatility Stocks, January 1968–December 2008

	Low-Volatility Quintile	Diagonal Only	Full-Risk Model	Levered, Full-Risk Model
Geometric average $R_p - R_f$	4.12%	5.26%	4.85%	7.26%
Average $R_p - R_f$	4.86%	6.42%	5.41%	8.82%
Standard deviation	12.74%	15.93%	11.50%	18.80%
Sharpe ratio	0.38	0.40	0.47	0.47
Average $R_p - R_m$	0.62%	2.18%	1.17%	4.58%
Tracking error	7.45%	5.61%	8.67%	10.12%
Information ratio	0.08	0.39	0.13	0.45
Beta	0.70	0.95	0.61	1.00
Alpha	2.00%	2.52%	2.90%	4.70%
t(Alpha)	2.03	3.21	3.01	3.03

Notes: For each month, we formed a minimum-variance portfolio of the top 1,000 stocks by market capitalization in the CRSP universe by using two methods and compared performance with a low-volatility sort. We estimated volatility by using up to 60 months of trailing returns (i.e., return data starting as early as January 1963). We estimated the covariance matrix as in Clarke, de Silva, and Thorley (2006). We limited the individual stock weights to between 0 and 3 percent. The third column has a simple five-factor covariance matrix with a Bayesian shrinkage parameter applied to the correlations, and the second column has only the diagonal of the covariance matrix. The fourth column levers the third-column portfolio to produce an average beta of 1. Average returns are monthly averages multiplied by 12. Standard deviation and tracking error are monthly standard deviations multiplied by the square root of 12. The return on the market, R_m, and the risk-free rate, R_f, are from Ken French's website (http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html).

Figure 4. A Low-Volatility Portfolio vs. a Portfolio of Low-Volatility Stocks, January 1968–December 2008
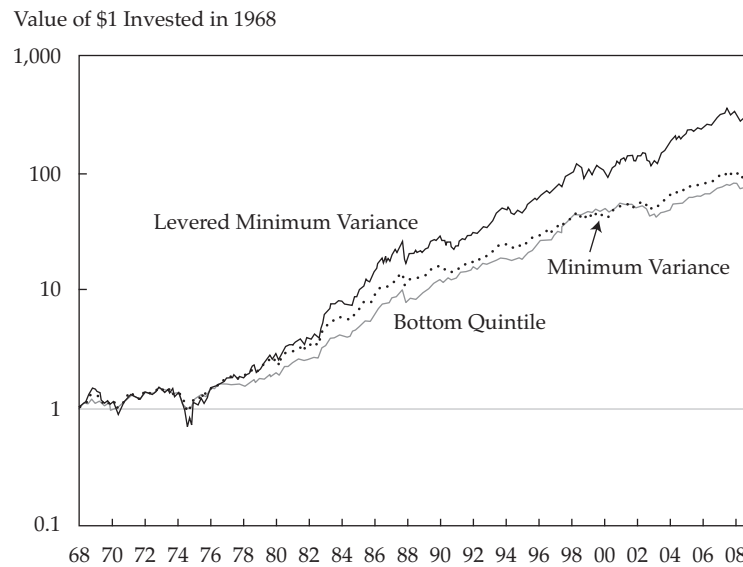
Notes: For each month, we formed portfolios by using the following three methods: (1) A minimum-variance portfolio of the top 1,000 stocks by market capitalization in the CRSP universe under the covariance matrix estimate methodology of Clarke, de Silva, and Thorley (2006). We limited the individual stock weights to between 0 and 3 percent. (2) The portfolio of the lowest quintile by trailing volatility. We measured volatility as the standard deviation of up to 60 months of trailing returns (i.e., return data starting in January 1963). (3) A levered minimum-variance portfolio to produce an average beta of 1.

Conclusion: 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, not risk differences, are emphasized. Institutional investment managers are well positioned to take advantage of such anomalies because they can generate high excess returns while maintaining average risks, thereby matching their benchmark's risk and controlling tracking error.

But the low-risk anomaly is of a very 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 such mispricing. We developed this argument and introduced some preliminary evidence.

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

This article qualifies for 1 CE credit.

Appendix A. Delegated Portfolio Management and the CAPM

This short derivation 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 $\mathbf{\Sigma}$. A risk-free bond 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 v.

3. *Investment strategies.* Each representative investor makes a scalar asset allocation decision a_j between stocks and the risk-free asset, as well as a vector portfolio choice decision $\mathbf{w}_j$.

   a. Investor 1 delegates his portfolio choice. He 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 her own portfolio. She allocates a fraction a_2 of her capital to stocks and chooses a portfolio $\mathbf{w}_2$, which can be collapsed without loss of generality to a single choice variable $\mathbf{w}_2$. Mean–variance utility maximization means that she chooses $\mathbf{w}_2$ to maximize

$$E\left(\mathbf{w}_2{'}\mathbf{R} - R_f \mathbf{1}\right) - \frac{v}{2}\mathbf{w}_2{'}\mathbf{\Sigma}\mathbf{w}_2.$$

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

$$E\left(R_i - R_f\right) = \beta_i E\left(R_m - R_f\right). \tag{A1}$$

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

4. *Intermediation.* A single intermediary chooses a portfolio $\mathbf{w}_1$ to maximize $E(\mathbf{w}_1 - \mathbf{w}_b)'\mathbf{R} - \gamma(\mathbf{w}_1 - \mathbf{w}_b)'\mathbf{\Sigma}(\mathbf{w}_1 - \mathbf{w}_b)$, where $\mathbf{w}_b$ 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 Type 1 investors. The intermediary chooses $\mathbf{w}_1$ to maximize IR; Investor 2 chooses $\mathbf{w}_2$ to maximize the Sharpe ratio; and the two compete to set prices. Note that the budget constraint in the intermediary's objective means that the IR 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 Koijen (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:

$$E\left(R_i - R_f\right) = \beta_i E\left(R_m - R_f\right) + c\left(1 - \beta_i\right), \qquad \text{(A2)}$$

where

$$c = A\frac{a_1 v}{a_2 \gamma + a_1 v} > 0$$

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

As Brennan (1993) showed, Type 2 investors will specialize in lower-volatility stocks. In this example, they are rational mean–variance utility maximizers who partially offset the effects of an intermediary who tries to capture improvements in the IR 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 start to act as arbitrageurs only when the relationship between risk and return is inverted.

Notes

1. http://mba.tuck.dartmouth.edu/pages/faculty/ken. french/data_library.html.

References

Alpert, Marc, and Howard Raiffa. 1982. "A Progress Report on the Training of Probability Assessors." In *Judgment under Uncertainty: Heuristics and Biases.* Edited by Daniel Kahneman, Paul Slovic, and Amos Tversky. New York: Cambridge University Press.

Anand, Shefali. 2009. "Leverage Shakes Up Mutual Funds, Which Discover a Strategy's Downside." *Wall Street Journal* (24 January):B2.

Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang. 2006. "The Cross-Section of Volatility and Expected Returns." *Journal of Finance*, vol. 61, no. 1 (February):259–299.

———. 2009. "High Idiosyncratic Volatility and Low Returns: International and Further U.S. Evidence." *Journal of Financial Economics*, vol. 91, no. 1 (January):1–23.

Baker, Malcolm, and Jeffrey Wurgler. 2007. "Investor Sentiment in the Stock Market." *Journal of Economic Perspectives*, vol. 21, no. 2 (Spring):129–152.

Bali, Turan, Nusret Cakici, and Robert F. Whitelaw. Forthcoming. "Maxing Out: Stocks as Lotteries and the Cross-Section of Expected Returns." *Journal of Financial Economics*.

Barberis, Nicholas, and Ming Huang. 2008. "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices." *American Economic Review*, vol. 98, no. 5 (December):2066–2100.

Barberis, Nicholas, and Richard Thaler. 2003. "A Survey of Behavioral Finance." In *Handbook of the Economics of Finance*, vol. 1B. Edited by George M. Constantinides, Milton Harris, and René Stulz. Amsterdam: Elsevier.

Barberis, Nicholas, and W. Xiong. 2010. "Realization Utility." Working paper, Yale University (July).

Black, Fischer. 1972. "Capital Market Equilibrium with Restricted Borrowing." *Journal of Business*, vol. 4, no. 3 (July):444–455.

Black, Fischer, Michael C. Jensen, and Myron Scholes. 1972. "The Capital Asset Pricing Model: Some Empirical Tests." In *Studies in the Theory of Capital Markets*. Edited by Michael C. Jensen. New York: Praeger.

Blitz, David C., and Pim van Vliet. 2007. "The Volatility Effect: Lower Risk without Lower Return." *Journal of Portfolio Management*, vol. 34, no. 1 (Fall):102–113.

Boyer, Brian, Todd Mitton, and Keith Vorkink. 2010. "Expected Idiosyncratic Skewness." *Review of Financial Studies*, vol. 23, no. 1 (January):169–202.

Brennan, Michael. 1993. "Agency and Asset Pricing." Working paper, University of California, Los Angeles (May).

Brennan, Michael, and Feifei Li. 2008. "Agency and Asset Pricing." Working paper, University of California, Los Angeles (11 March).

Clarke, Roger G., Harindra de Silva, and Steven Thorley. 2006. "Minimum-Variance Portfolios in the U.S. Equity Market." *Journal of Portfolio Management*, vol. 33, no. 1 (Fall):10–24.

Cornell, Bradford. 2009. "The Pricing of Volatility and Skewness: A New Interpretation." *Journal of Investing*, vol. 18, no. 3 (Fall):27–30.

Cornell, Bradford, and Richard Roll. 2005. "A Delegated-Agent Asset-Pricing Model." *Financial Analysts Journal*, vol. 61, no. 1 (January/February):57–69.

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina. 2002. "Differences of Opinion and the Cross Section of Stock Returns." *Journal of Finance*, vol. 57, no. 5 (October):2113–2141.

Fama, Eugene F., and Kenneth R. French. 1992. "The Cross-Section of Expected Stock Returns." *Journal of Finance*, vol. 47, no. 2 (June):427–465.

Fischhoff, Baruch, Paul Slovic, and Sarah Lichtenstein. 1977. "Knowing with Certainty: The Appropriateness of Extreme Confidence." *Journal of Experimental Psychology. Human Perception and Performance*, vol. 3, no. 4 (November):552–564.

Frazzini, Andrea, and Lasse H. Pedersen. 2010. "Betting against Beta." Working paper, New York University (October).

Haugen, Robert A., and A. James Heins. 1975. "Risk and the Rate of Return on Financial Assets: Some Old Wine in New Bottles." *Journal of Financial and Quantitative Analysis*, vol. 10, no. 5 (December):775–784.

Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, vol. 47, no. 2 (March):263–291.

Karceski, Jason. 2002. "Returns-Chasing Behavior, Mutual Funds, and Beta's Death." *Journal of Financial and Quantitative Analysis*, vol. 37, no. 4 (December):559–594.

Kumar, Alok. 2009. "Who Gambles in the Stock Market?" *Journal of Finance*, vol. 64, no. 4 (August):1889–1933.

Miller, Edward M. 1977. "Risk, Uncertainty, and Divergence of Opinion." *Journal of Finance*, vol. 32, no. 4 (September):1151–1168.

Mitton, Todd, and Keith Vorkink. 2007. "Equilibrium Underdiversification and the Preference for Skewness." *Review of Financial Studies*, vol. 20, no. 4 (July):1255–1288.

Pettengill, Glenn N., Sridhar Sundaram, and Ike Mathur. 1995. "The Conditional Relation between Beta and Return." *Journal of Financial and Quantitative Analysis*, vol. 30, no. 1 (March):101–116.

Roll, Richard. 1992. "A Mean/Variance Analysis of Tracking Error." *Journal of Portfolio Management*, vol. 18, no. 4 (Summer):13–22.

Sensoy, Berk. 2009. "Performance Evaluation and Self-Designated Benchmark Indexes in the Mutual Fund Management Industry." *Journal of Financial Economics*, vol. 92, no. 1 (April):25–39.

Shefrin, Hersh, and Meir Statman. 2000. "Behavioral Portfolio Theory." *Journal of Financial and Quantitative Analysis*, vol. 35, no. 2 (June):127–151.

Shleifer, Andrei. 2000. *Inefficient Markets: An Introduction to Behavioral Finance*. Oxford, U.K.: Oxford University Press.

Tversky, Amos, and Daniel Kahneman. 1983. "Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment." *Psychological Review*, vol. 90, no. 4 (October):293–315.

———. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty*, vol. 5, no. 4 (October):297–323.

van Binsbergen, Jules H., Michael W. Brandt, and Ralph S.J. Koijen. 2008. "Optimal Decentralized Investment Management." *Journal of Finance*, vol. 63, no. 4 (August):1849–1895.