The Most Important Number in Finance

The Quest for the Market Risk Premium

Marc Zenner
marc.p.zenner@jpmorgan.com
(212) 834-4330

Scott Hill
scott.d.hill@jpmorgan.com
(415) 315-8842

John Clark
john.hs.clark@jpmorgan.com
(212) 834-2156

Nishant Mago
nishant.x.mago@jpmorgan.com
(212) 834-2172

JPMorgan

1. The most important number in finance

You will not find it in section C of *The Wall Street Journal*. *CNBC* will not mention it in its morning market recap. *The Economist* will not provide it in its back pages with other financial data. Yet it is one of the most critical metrics in finance, a figure implicit in the evaluation of financing and investment opportunities: the market risk premium. What is it? How and where should it be used? What is the right number to use? Does it change over time?

In this report, we (1) estimate a current range of risk premiums; (2) explain how the risk premium has increased since the beginning of the subprime crisis; (3) discuss how, thanks to Federal Reserve intervention, a higher risk premium does not necessarily lead to a higher cost of capital; and (4) debate how possible divergence between equity and credit markets since last summer may affect strategic and financial decision-making. In addition, we review some common methods used to estimate the market risk premium.

What is the market risk premium?

The *market risk premium* (MRP) reflects the incremental premium required by investors, relative to a risk-free asset like U.S. Treasury bonds, to invest in a globally diversified market portfolio. Below is a simple and generally accepted equation:

Expected return on the market portfolio = Risk-free rate of return + market risk premium

Should the market risk premium be higher for some assets and lower for others? Most likely yes, but how should the adjustment be made? The Capital Asset Pricing Model (CAPM) proposes one such adjustment. CAPM states that the expected return on an asset is the risk-free rate plus an MRP that is adjusted, through beta, to reflect the market risk of the asset:

*Expected return on an asset = Risk-free rate of return + beta * market risk premium*

The beta is a calibration factor that is higher (lower) than one if the asset has a systematic, or non-diversifiable, risk that is higher (lower) than the market's risk. In the CAPM framework, the MRP should apply to all assets, including bonds, real estate, art, etc. In practice, however, the risk premium is mostly used to estimate the expected return on equity (also referred to as the cost of equity). Bond markets rely on their own risk premium concept, the credit spread, which is the difference between the yield on a bond and the maturity-matched Treasury rate.

From a macroeconomic perspective, the MRP reflects the broader outlook on the whole economy. Factors influencing investors' views on market risk include outlooks for economic growth, consumer demand, inflation, interest rates, and geopolitical risks. As such, the MRP is a single metric that reflects these inputs in the expected returns of various asset classes.

## Why is the market risk premium so important?

While many finance professionals and executives actively manage their debt and debate the incremental basis points their firm may have to pay on new bonds, they do not tend to focus much on the cost of equity. Is it that debt financing is so much more prevalent than equity financing? Not really. Even with a tax system that favors debt financing, equity financing constitutes over 80% of the total market capitalization for a typical non-financial S&P 500 firm today.

Why then is there less focus on the cost of equity? Maybe because most firms manage debt actively and equity only passively; or because an economic cost of equity of 12% does not translate into an actual cash outlay of 12%; or perhaps because there is no consensus on how to estimate the market risk premium.

> **Practical Application:** Understanding and quantifying the MRP is critical to the value-creation process. With most of their capitalization in the form of equity, decision-makers require an estimate of the MRP to determine their cost of capital, identify projects that create shareholder value, decide how much to pay for acquisition targets, evaluate their capital structure, and compare the costs of various sources of financing. Not adjusting the cost of equity to new market realities may lead firms to (1) over or under-invest or (2) forgo capital-structure opportunities.

What is the market risk premium today?

No single method to estimate the MRP is used universally. Our review of various methods (detailed in Section 2) suggests that they each have strengths and weaknesses. They also generate a wide range of results as summarized in the figure below. We therefore recommend thinking about the MRP in terms of a range rather than a unique number. Based on our results, the MRP probably falls within a range of 5% - 7% today.

Figure 1: Summary of risk premium estimates
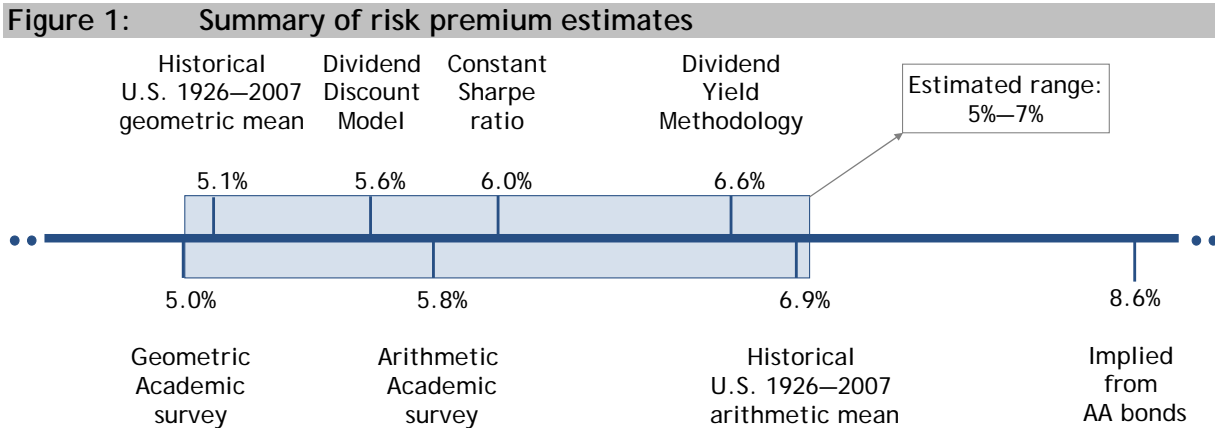
Source: JPMorgan

If I do not use CAPM, should I still focus on the market risk premium?

Most practitioners use CAPM as their method of choice to estimate the cost of capital. Interestingly, while academics often emphasize the limitations of CAPM, they still tend to focus on it when teaching about the cost of capital. Two of the risk premium estimation methods we used rely on CAPM (the Dividend Yield and the bond-based methods). The Dividend Discount and Sharpe ratio methods, as well as the historical analysis, do not rely on CAPM. Practitioners who do not use CAPM can still use the risk premium range we suggest by using the low (high) end of the range for projects they perceive to be at the low (high) end of their risk spectrum.

## My firm is global, so should I focus on a risk premium based on U.S. data?

The results we present are based on U.S. market data. Can you use these risk premium estimates for investments in other countries? We believe that the U.S.-based MRP is a reasonable estimate for developed economies for a couple of reasons. First, an unconstrained investor would not freely invest in a market in which he/she would earn a lower risk-adjusted return. Hence risk premiums should gravitate to each other across open developed markets, and the U.S.-based risk premium should serve as a good estimate for this. The situation may be different in emerging markets, however, where non-market risks may exist (e.g., political risk) or where investor segmentation and constraints limit the free flow of capital into and out of the country. Second, the U.S. market has some data advantages, namely very broad markets with long data histories. Many other markets tend to be over-weighted in some sectors (e.g., banking, shipping, energy, telecommunications) or have data series that have been interrupted by political events in the 20th century.

Has the risk premium changed since last summer?

Are we in a new risk premium environment? The figure below shows that the answer depends on the methodology. The historical method, as expected, suggests no change in the risk premium. On the other hand, methods that rely on current market information (which we discuss in detail later)

JPMorgan

signal that the risk premium has increased since the credit crisis began last summer, but that it has declined from its peak in February/March.

Figure 2:   Comparing risk premium estimates since last summer
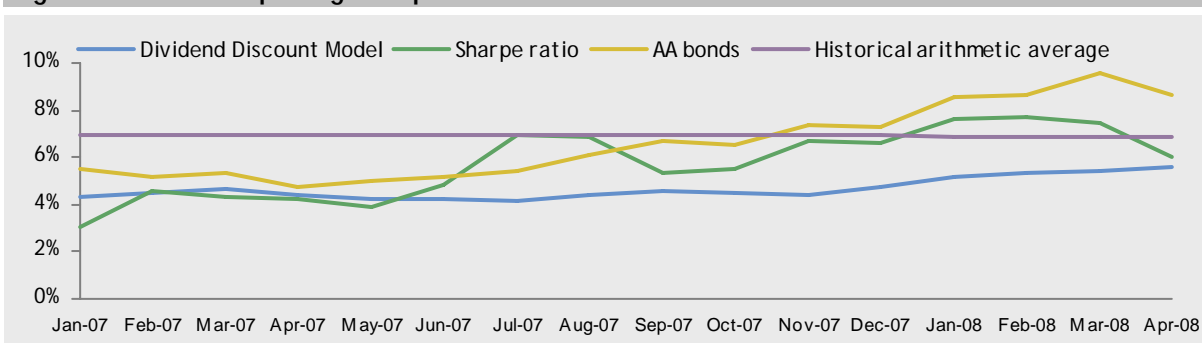
Source: JPMorgan, SBBI Market Report-Morningstar, Bloomberg

Should executives change their hurdle rates for capital allocation?

Boards of Directors and senior executives implicitly use the MRP when determining hurdle rates for new projects and acquisitions. There is a preference for hurdle rates that do not change often, possibly because stable hurdle rates facilitate communication with regional and divisional management. In some cases, however, it is critical to understand whether changing market conditions affect how the market prices risk. Financial decision-makers examine day-to-day data when they look at debt financing, so why not also for equity, often the biggest component of the capital structure? We believe that today's environment warrants re-estimating the cost of capital using new market information, in particular when considering large capital projects or acquisitions.

> **Practical Application:** The cost of capital for many S&P 500 firms has not increased since last summer. Why? While risk premiums increased in both credit and equity markets, the Fed's policy of lowering interest rates has succeeded in offsetting this increase for the largest firms in the economy. It is worth noting that, even in today's environment, many firms tend to use a hurdle rate that is a few percentage points higher than their true cost of capital, which may lead them to forgo valuable investment opportunities.

Which is right—equity or credit markets?

Many market observers have focused on how the equity and credit markets have behaved differently since last summer. While credit markets lost significant liquidity and experienced dramatic pricing changes, the non-financial component of equity markets remained relatively unaffected until the beginning of this year. Have credit markets overreacted, and should they revert to more normalized pricing? Have the equity markets failed to completely absorb the effects of the financial crisis, and should we expect a further decline in equity values, along with an increase in the MRP? Or do credit markets reflect a higher overall premium combining both a heightened risk premium and an increased liquidity premium? In many segments of the credit markets, liquidity diminished significantly over the last few months, but not so in the equity markets. We believe that both effects have taken place; i.e., the equity risk premium has increased, but the credit markets have been affected even more because they are also pricing in an additional premium for liquidity.

> **Practical Application:** Executives should consider this debt vs. equity market premium dynamic when making funding decisions. For example, the after-tax cost of hybrids should be compared to an updated after-tax cost of equity. Furthermore, as discussed above, given that low Treasury rates have offset rising risk premia for the largest firms, executives should consider locking in a low long-term cost of capital, especially if they have near-term refinancing, capital or liquidity needs, or if they expect rates to increase because of inflationary pressures.

2. Different methods to estimate the MRP

A. Historical average realized returns

A common way to estimate the MRP has been to compare realized annual equity returns to average returns of U.S. Treasury bonds over some historical time period.

MRP = average annual equity index return - average return on Treasury bonds

This method is widely used in practice but has a few weaknesses which diminish its usefulness.

Choice of averaging method: The choice of arithmetic vs. geometric averaging methods can lead to significant differences in MRP estimates. For example, if $100 grows to $110 in one year and then drops back to $100 the next, the arithmetic average annual return is [+10.0% - 9.1%]/2, or 0.5%. The arithmetic average represents the best estimate of annual expected return. The geometric mean, however, will be 0%, which is the compounded annual return the investor actually earned. Many academics prefer the arithmetic average because it represents an investor's expected return at any given point in time. But the geometric mean better reflects asset returns investors should expect over long horizons.

Time horizon: As evidenced in Figure 3 below, different time horizons also yield different MRP estimates. For example, an observer examining the U.S. data since 1978 using the geometric mean would determine that the MRP is 4.9%, whereas an observer viewing the data since 1946 would instead conclude it is 5.7%.

Figure 3:	Historical risk premium estimates across various time periods	
Large company stocks – Intermediate T bonds	**Arithmetic**	**Geometric**
1926-2007	6.9%	5.1%
1946-2007	6.8%	5.7%
1978-2007	5.7%	4.9%

Source: Morningstar, JPMorgan

Reaction to changing risk premium: In a changing risk-premium environment, this method can produce counterintuitive results. For example, if the risk premium increases and cash-flow projections remain unchanged, equity prices will drop. This drop in equity prices reflects investors' demand for higher *future* expected returns in the riskier environment. But the drop would cause lower realized returns, which in turn would lower the average *historical* returns, thereby suggesting a lower instead of higher risk premium. Though this backward-looking method may not capture the direction of the change in risk premium well, it may still be a viable long-term estimate of the risk premium investors expect to earn by investing in equity.

Figure 4:	Pros and cons of using the historical method
Pros	**Cons**
■ Easy to compute	■ Estimate depends on historical window
■ Has been a standard in business schools	■ Estimate depends on averaging method
■ Does not change often and rapidly	■ Does not change often and rapidly; i.e., does not incorporate new market realities
■ Can be sourced by a third-party provider such as Ibbotson Associates	■ Responds in a counterintuitive way to changes in actual risk premium

Source: JPMorgan

JPMorgan

B. Dividend Discount Model

Another means of estimating the MRP is through the Dividend Discount Model (DDM), which can be used to calculate the current market cost of equity. The model solves for an internal rate of return (cost of equity) based on the price level and expected dividend stream of an index (often the S&P 500 as a proxy for the broad market). Dividends are projected by applying an expected payout ratio to forecasted earnings. Earnings are forecasted, in turn, by combining near-term (i.e., 5 years) market estimates with a perpetuity growth rate equivalent to long-term nominal GDP growth. The dividend payout ratio is initially assumed to be the average of recent historical payout ratios, but increases over the long-term towards 80% in the terminal period as reinvestment opportunities are assumed to subside. Simplistically, the formula for the market cost of equity is:

$$Price_0 = \sum_{t=1}^{\infty} \frac{Dividend_t}{(1 + Cost\ of\ Equity)^t}$$

where *t* is time from now to infinity. Subtracting the 10-year government bond yield from the market cost of equity then provides the market risk premium. Thus, the MRP formula is as follows:

MRP = Cost of equity implied by DDM – 10-year government bond yield

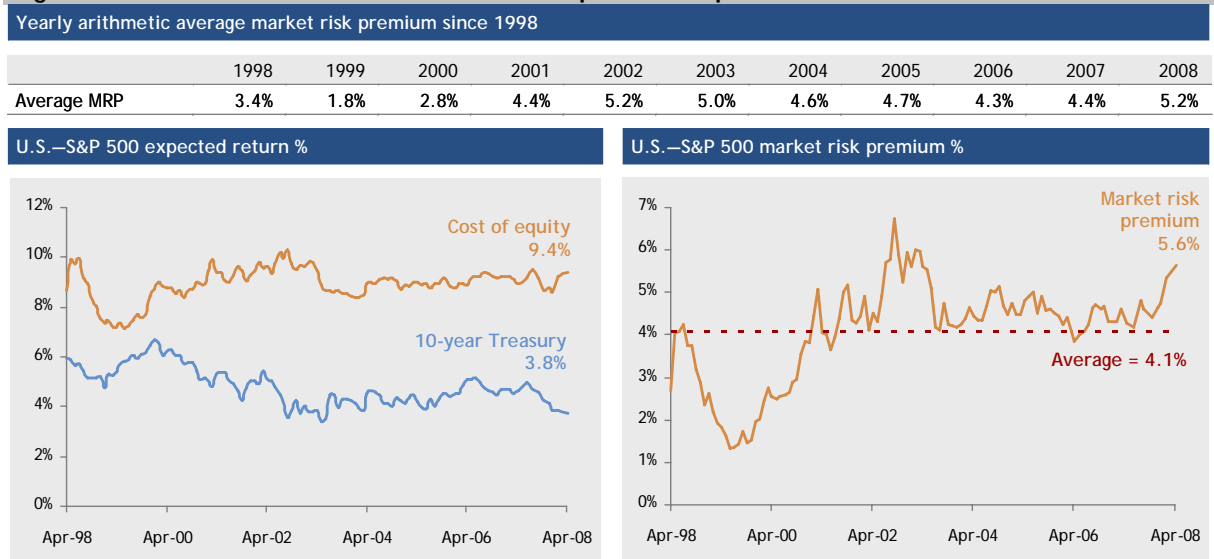| Figure 5:	Pros and cons of risk premium implied from Dividend Discount Model
Pros	**Cons**
■ Implied from equity market values	■ Price variable changes daily
■ Changes and responds to current market environment	■ Highly dependent on future dividend/cash flow estimates
■ Forward-looking; not heavily reliant on historical data	■ Dividend forecasts not updated frequently; may not take market cycles into account

Source: JPMorgan

Changes over time: The market cost of equity varies primarily with movements in the level of the index, but also with changes in expectations for future dividends. The chart below shows the market cost of equity based on the S&P 500, as well as the 10-year Treasury yield, over the last 10 years. The resulting MRP, shown to the right, varies from a low of 1.3% at the peak of the market to a high of over 6% in the post-9/11 era. After 2003, the MRP stabilized in the 4% range until the recent credit crisis, which has led to a re-pricing of risk and a higher MRP.

Figure 6:   Dividend Discount Model implied risk premium over time

Yearly arithmetic average market risk premium since 1998

	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
Average MRP	3.4%	1.8%	2.8%	4.4%	5.2%	5.0%	4.6%	4.7%	4.3%	4.4%	5.2%

U.S.—S&P 500 expected return %

U.S.—S&P 500 market risk premium %

Source: JPMorgan, Bloomberg

JPMorgan

Capital Structure Advisory & Solutions

C. Constant Sharpe ratio method

Another useful metric to estimate the risk premium, the Sharpe ratio, has been inherited from portfolio management theory. The Sharpe ratio measures a portfolio's excess return per unit of risk and can be used to estimate the MRP:

$$\textit{Market Sharpe ratio} = S_M = \frac{Portfolio\ MRP}{Volatility\ of\ MRP}$$

We estimate that, over the last 50 years, the Sharpe ratio for the broad market (using the S&P 500 index as a proxy) has been about 0.3, which is consistent with academic research. Assuming that this ratio is constant going forward, we can then solve for the forward-looking MRP by multiplying the S&P 500 Sharpe ratio by a measure of future market volatility. We estimate future market volatility via the VIX index, which measures the volatility implied from options on the S&P 500 index. Thus, the Sharpe ratio-implied MRP is:

*MRP = Market (S&P 500) Sharpe ratio * Market (S&P 500) implied volatility*

Figure 7:   Pros and cons of the Sharpe ratio method

| Pros	Cons
■ Estimate of Sharpe ratio based on more than 50 years of historical data; consistent with academic research	■ Some evidence that Sharpe ratio may change over time instead of remaining constant
■ VIX component is forward-looking; captures shifts in investor sentiment very quickly	■ VIX measures short-term volatility (<1 year), whereas risk premium is generally viewed long-term (10+ years)

Source: JPMorgan

Changes over time: Figure 8 displays the Sharpe ratio-implied MRP over the last 10 years. By definition, the Sharpe ratio-implied MRP moves proportionally with the VIX volatility index. At times of greater uncertainty and market panic, including the Long-Term Capital Management fallout in 1998, the 2000-2002 recession/tech bubble burst, and the current credit crisis, investors have fled to safer securities and demanded a greater MRP to keep their investments in riskier assets. Such shifts in risk preferences have been accompanied by spikes in volatility.

Figure 8:   Historical risk premiums computed from the Sharpe ratio method

Yearly arithmetic average market risk premium since 1998

	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
Average MRP	8.2%	7.8%	7.4%	7.9%	8.0%	6.6%	4.6%	3.9%	3.7%	5.2%	7.2%

U.S.—S&P 500 market risk premium %
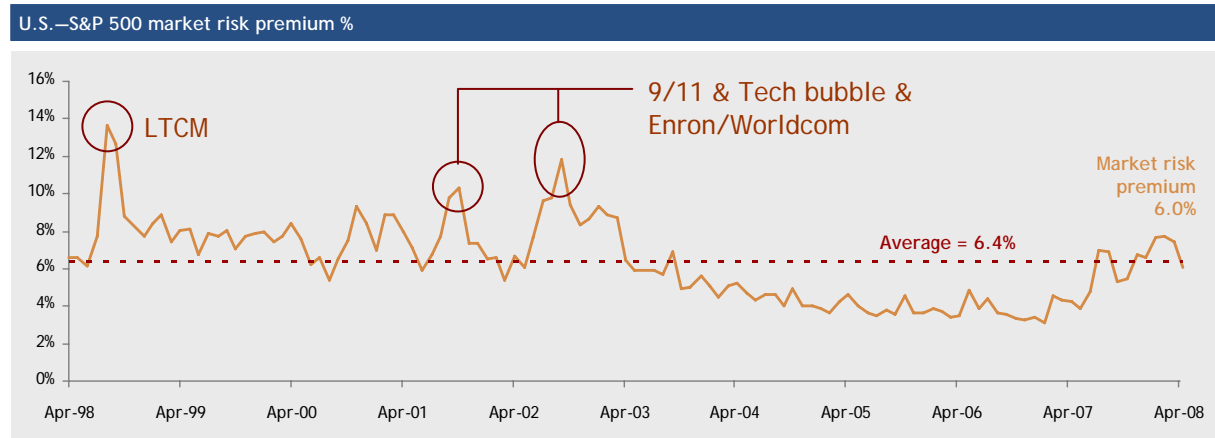
Source: Bloomberg, Federal Reserve Data

JPMorgan

D. Bond-market implied risk premium

Most of us think of the MRP in the context of cost of equity. Risk premiums do, however, also exist for corporate bonds. The expected return of a bond can therefore be expressed using the Capital Asset Pricing Model, as:

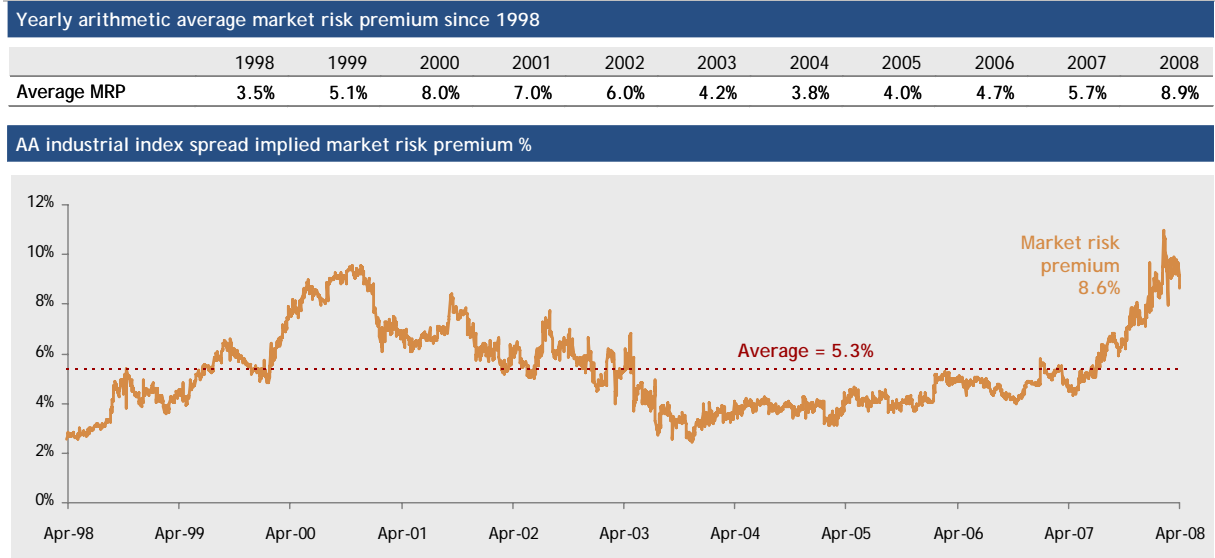*AA yield ≈ AA expected return = risk-free rate + beta * market risk premium*

Therefore, if we know the expected return on the bond and its beta, we can estimate the implied MRP. For high-yield bonds, we know the yield, but the expected return is likely to be significantly lower than the promised yield. For AA rated corporate bonds, on the other hand, the default probabilities are very low and we can use the yield as a proxy for expected returns. Hence, we use the price series of AA corporate bonds to estimate the MRP. The beta of AA bonds is between 0.15 and 0.20, depending on the estimation period. Using a beta of 0.15, we estimate that the bond-implied MRP was below 4% in 1998 and 2004-2005 but recently rose to about 8.6%.

Figure 9: Pros and cons of the bond-market data methodology

Pros	Cons
■ Based on daily market feedback regarding risk premium on bonds	■ Possibility that expected default rates change at the same rating
■ Assumes no capital-structure arbitrage; i.e., when bonds demand a higher risk premium, other assets such as equity also demand a higher return	■ Depends on CAPM and an assumption about bond betas
	■ Implied risk premium captures both a liquidity and risk premium

Source: JPMorgan

Figure 10: Recent changes in the bond-market implied risk premium

Yearly arithmetic average market risk premium since 1998

	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008
Average MRP	3.5%	5.1%	8.0%	7.0%	6.0%	4.2%	3.8%	4.0%	4.7%	5.7%	8.9%

AA industrial index spread implied market risk premium %

Source: JPMorgan, Bloomberg

E. Dividend Yield Method

A methodology that is closely related to the Dividend Discount Model method uses the dividend yield as a starting point. The price of a dividend-paying stock can be estimated using the constant-growth valuation model. This model assumes that the dividend will grow at a constant rate forever. We rewrite this model as a function of the cost of equity, stating that the cost of equity is the dividend yield plus the long-term growth rates. The formulas are:

$$Price = \frac{Div_1}{Cost\ of\ Equity - GrowthRate}\text{, and therefore } Cost\ of\ Equity = \frac{Div_1}{Price} + GrowthRate$$

This approach works well in sectors with large and steadily growing dividends. We applied the methodology to three industries known for their focus on dividend yields: Real Estate Investment Trusts (REITs), Master Limited Partnerships (MLPs), and regulated utilities. In the regulated utilities industry, regulators accept this method as a way to estimate the cost of equity. Another useful feature of the model is its closeness to the cash cost of the equity. In fact, some practitioners look at the dividend yield only and ignore the growth component of the equation.

MRP = (Cost of equity implied by Dividend Yield Method – 10-year government bond yield) / beta

Figure 11:   MRP implied by dividend yields in dividend-heavy sectors

	Dividend yield	IBES 5-yr EPS growth	Overall growth[1]	Cost of equity	Equity beta	Implied MRP
Regulated utilities	4.1%	6.1%	4.5%	8.6%	0.78	6.6%
MLPs	6.4%	6.5%	5.1%	11.4%	0.61	13.0%
REITs	5.4%	6.5%	4.7%	10.5%	1.13	6.1%
Mean	5.3%	6.4%	4.8%	10.2%	0.84	8.6%
Median	5.4%	6.5%	4.7%	10.5%	0.78	6.6%

Source: JPMorgan, FactSet

[1] Overall growth is weighted combination of 5-yr EPS growth and 4% perpetuity growth assumptions

We use EPS estimates and an assumption of constant payout ratios to forecast the dividend growth over the next five years, and an assumption that dividends will grow at 4% thereafter (long-term real growth plus inflation). Our results suggest that the cost of equity for these sectors is in the 9% to 12% range. The figures also display two clear weaknesses: (1) the need for assumptions to estimate overall or long-term growth, estimated in this case as a weighted-average of the 5-year EPS growth projection followed thereafter by a 4% perpetuity growth rate; and (2) the need to rely on CAPM and a beta estimate to extract the MRP implied by our cost-of-equity estimates. Today, this approach yields an MRP in the 6% range for REITs and utilities, and a higher number for MLPs.

Figure 12:   Pros and cons of MRP implied from Dividend Yield Method

Pros	Cons
■ Intuitive: cost of equity equals dividend yield plus a growth rate	■ Only applicable in a few dividend-heavy sectors
■ Widely accepted in dividend-heavy sectors	■ Capital structures of these sectors may not represent those of the market at large
■ Close to the actual cash cost on equity	■ Relies on perpetuity growth rate assumption
■ Dividend yield changes daily	■ Depends on CAPM and assumption about industry or firm beta

Source: JPMorgan

JPMorgan

F. Survey evidence

One relatively basic method for determining the MRP is to survey market participants for their views on required returns. Such surveys have polled academics, investors, and other corporate-finance practitioners such as CFOs.

An academic survey by Ivo Welch from Brown University provides useful insights on MRP estimates.[1] The typical finance professor responding to Welch's survey estimates that the long-term market risk premium is 5% on a geometric basis and 5.8% on an arithmetic basis. Interestingly, these numbers are very close to the MRP estimates of the historical realized returns methodology, suggesting that finance professors still primarily rely on that approach.

A similar survey conducted quarterly from 2000 to 2007 by John Graham and Campbell Harvey of Duke University compiled the views of U.S. CFOs regarding the current risk premium.[1] Their average risk premium in 1Q07 was 3.2%, and the range from 2000 to 2007 was 2.4% to 4.7%.

Relying on these survey results has some advantages. First and foremost, in the case of finance professors, participants may be biased in their preferred methodology, but they are typically unbiased in their MRP estimates—that is, they do not have any specific incentive to make low or high estimates. Secondly, academics tend to spend a lot of time on the subject and have significant influence on how regulators, practitioners, and even investors look at the MRP.

On the other hand, survey respondents can provide wide differences of opinion and express views that may be extreme (such as a negative MRP). Surveys can also reflect the collective views of the constituent base. As an example, academics' reliance on the historical-data approach suggests that their estimates will not change very often. This may be an advantage for executives looking for a MRP estimate that is robust through time, but it may not capture the realities of a new market environment (such as structural shifts, tax changes, etc.). Conversely, the CFO-based survey is different in that its results are quite volatile and might represent current market conditions and concerns.

Figure 13: Pros and cons of surveys

Pros	Cons
■ Significant time researching this topic	■ Wide differences in opinion
■ Academic estimate unbiased (no reasons for it being high or low)	■ Does not change often and rapidly; i.e., does not incorporate new market realities (e.g., tax rate changes)
■ Does not change often and rapidly	

Source: JPMorgan

> As stated above, none of these six estimation methods are used universally. Taken together, however, they provide an understanding of the drivers of the market risk premium, and allow decision-makers to consider using a method that reflects today's volatile market environment.

---

[1] Ivo Welch, "The Consensus Estimate for the Equity Premium by Academic Financial Economists in December 2007", January 18, 2008.
[2] John Graham and Campbell Harvey, "The Equity Risk Premium in January 2007: Evidence from the Global CFO Outlook Survey", January 18, 2007.