

Measuring alliance performance

Large companies often have dozens of alliances—and little idea how they are performing. Here’s how to evaluate them.

Jim Bamford and David Ernst

As corporations have evolved from command-and-control structures with sharply defined boundaries into loosely-knit organizations, corporate alliances¹ have become central to many business models. Most large companies now have at least 30 alliances, and many have more than 100.

Yet despite the ubiquity of alliances—and the considerable assets and revenues they often involve—very few companies systematically track their performance, creating a substantial risk of negative surprises. We believe that every corporate manager, including CFOs, should be well tuned into the performance of their alliances. In our work with more than 500 companies around the world, we have found that fewer than one in four alliances have adequate performance metrics in place.² As a result, alliances tend to be run by intuition and with incomplete information. Partners may not agree about the progress of their ventures and senior management can’t intervene quickly enough to correct problems. In many companies, 30 to 60 percent of alliances are underperforming—and three to five major deals are in desperate need of restructuring. Unfortunately, management does not really know where the problems lie or how it should best invest its time. To get a better grip on performance, companies must develop a more structured approach to evaluating the health of their alliances. Doing so is not a straightforward task.


The challenge

Why is alliance performance so hard to measure? The answer can be traced to three characteristics. First, alliances are by definition arrangements between separate companies, each employing its own reporting processes and systems and each pursuing its own goals for the alliance. This makes it hard to agree on a single measure of performance and creates incentives to conceal benefits and inflate costs.

A second critical characteristic is an operational interdependence that makes benefits and costs difficult to track. Most alliances receive some inputs from their parent companies (including raw materials, customer data, and administrative services) and in turn provide outputs to them, creating complicated transfer-pricing issues. Before Airbus Industrie was revamped in 2001, for instance, the four consortium members made aircraft sections and “sold” them to the joint venture, which then assembled and marketed the airplanes. Setting transfer prices was a challenge because of the partners’ sensitivity to sharing detailed cost data. In another case, two global technology companies agreed to jointly market a new product. This alliance involved more than 30 working teams whose 300-odd members spent between 20 percent and 60 percent of their time on the alliance. One executive admitted that he had no real idea how much the company had spent on the

Exhibit 1. Get to know your alliance

Net present value to parent company (disguised example) \$ million



¹Includes terminal value

²Includes management costs

Source: McKinsey analysis

venture, so large were its hidden costs. Measuring benefits is no easier. Alliances often generate sales of related products for parent companies, which should also be taken into account in assessing performance and value. So should longer term benefits like opportunities for learning, access to new technologies and markets, and improved competitive positioning.

The third characteristic is the noncore position of alliances within the corporate portfolio. Because they are inside the corporation, not totally outside, they often do not receive the same level of management scrutiny as business units and internal initiatives; nor are they subject to the same level of market forces as standard customer or supplier relationships.

Measuring performance

To overcome these difficulties, companies must assess the performance of their alliances

on three levels, each focusing on different aspects of the problem and prompting distinct managerial responses.

At the first level, every alliance should be individually assessed to establish how it is performing and whether the parent company needs to intervene. At the second level, a company should periodically search for performance patterns across the portfolio—a process that often leads to adjustments in the types of deals a company pursues and sometimes to additional investments in a drive to build alliance-related skills. At the third level, once a company better understands how its portfolio is performing, it can conduct a top-down review of overall strategy to ensure not only that its alliance portfolio is configured for optimal performance, but also that it has ranked new opportunities in a clear order of priority. The following excerpt focuses on the performance of individual alliances.³

Developing a detailed view of the economics of an alliance is indispensable to measuring its performance. This measurement should go well beyond the usual cash flow metrics to include transfer-pricing benefits, benefits outside the scope of the deal (for instance, sales of related products), the value of options created by the alliance, as well as start-up and ongoing management costs (Exhibit 1). This information is vital for managers to evaluate deals up front and to monitor their continuing performance.

For example, one company in the power industry calculated the embedded option value of a potential alliance to commercialize a critical new technology, evaluating the odds of different possible outcomes and the associated payoffs for each. Its calculation showed that for the given alliance the firm had a

10 percent chance of creating \$1 billion in annual income within three years, a 20 percent chance of creating a modestly successful business producing \$10 million to \$30 million in annual income, a 60 percent chance of losing \$10 million to \$30 million, and a 10 percent chance of losing more than \$200 million. This profile of potential option value was extremely sensitive to assumptions regarding technology and construction costs, leading the company to closely monitor the alliance's early performance, while reserving the right to cut off funding in the event that technical progress slowed.

Likewise, one biotechnology company developed explicit values for the potential learning benefits from a planned development and marketing alliance with a large pharmaceutical company. This exercise provided the information that ultimately pushed the firm to choose one partner over another and led it to closely monitor the benefits of the alliance, eventually allowing the firm to “migrate” into new capabilities.

Having formed an alliance with a clear and integrated view of the economics in mind, a company must develop, within 30 days of the launch, a scorecard to track the venture's performance. Partners must decide whether to share a single scorecard, to run separate scorecards, or to use a combination of the two. For a joint venture with its own P&L, a single scorecard is often possible. For most other alliances, the combination approach works best. Each partner can supplement a shared scorecard with additional metrics that track progress against goals that aren't shared by the other partners. This approach also enables each partner to devise internal metrics that allow it to compare the performance of an alliance with the performance of business

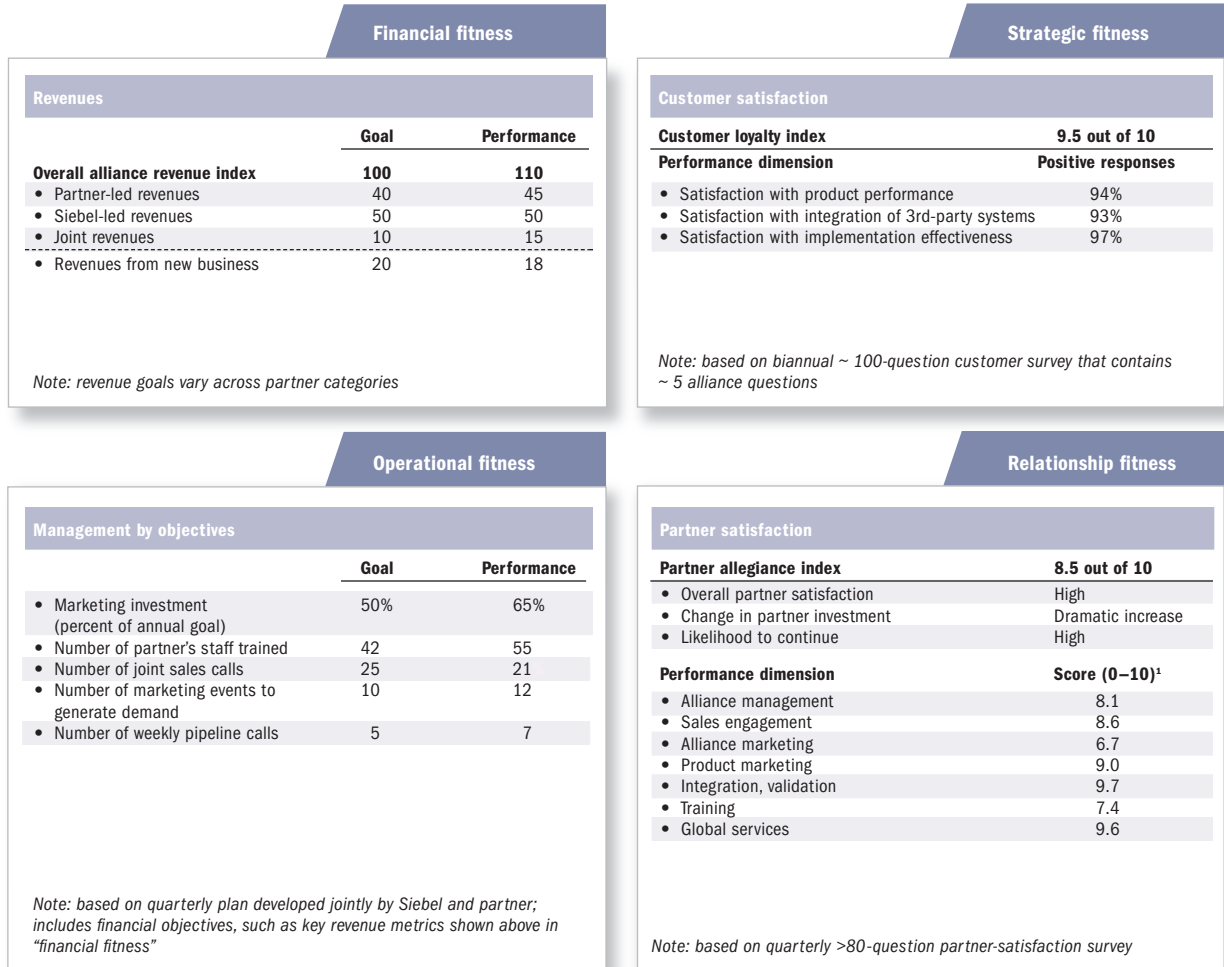
activities outside the alliance or to other, similar alliances.

It is essential, both at the alliance and the parent level, to take a balanced view of performance. To achieve such a balance, we have found it useful to include four dimensions of performance fitness: financial, strategic, operational, and relationship. Financial and strategic metrics show how the alliance is performing and whether it is meeting its goals—but may not provide enough insight into exactly what, if anything, isn't going well. Operational and relationship metrics can help uncover the first signs of trouble and reveal the causes of problems. Together, the four dimensions of performance create an integrated picture that has proved invaluable to the relatively few companies, such as Siebel Systems, that have used them to measure the health of alliances (Exhibit 2).

1. *Financial fitness*: Metrics such as sales revenues, cash flow, net income, return on investment, and the expected net present value of an alliance measure its financial fitness. Most alliances should also monitor progress in reducing overlapping costs, achieving purchasing discounts, or increasing revenues. In addition, financial fitness can include partner-specific metrics such as transfer-pricing revenues and sales of related products by the parent companies. At one international oil industry joint venture, the partners tracked not only revenues and consolidation synergies on a quarterly basis, but also the costs of goods sold to and from the parents, as well as estimates of profitability on those parent-related transactions.

2. *Strategic fitness*: Nonfinancial metrics such as market share, new-product launches, and customer loyalty can help executives measure

Exhibit 2. Siebel Systems' alliance scorecard



¹ Siebel uses this information to calculate gap score (importance of dimension to partner – Siebel performance = gap score); gaps of 2.0 or higher require action plan by alliance manager; performance in applying this plan is monitored by Siebel and senior executives of partner's company.

Source: Siebel Systems; McKinsey analysis

the strategic fitness of a deal; other metrics could, for example, track the competitive positioning and access to new customers or technologies resulting from it. Devising strategic metrics can take imagination. The international semiconductor research consortium SEMATECH, for instance, tracks the number of employees from member companies who are working on its research initiatives in order to assess whether it is transferring knowledge to its partners.

3. *Operational fitness*: The number of customers visited and staff members recruited, the quality of products and manufacturing throughput are examples of operational fitness metrics that call for explicit goals linked to the performance reviews and compensation of individuals. For example, executives at one health care company define operationally-fit alliances as those hitting 60 to 80 percent of their key operating milestones. Any figure higher than 80 percent

indicates that the goals weren't sufficiently ambitious.

4. *Relationship fitness*: Questions about the cultural fit and trust between partners, the speed and clarity of their decision making, the effectiveness of their interventions when problems arise, and the adequacy with which they define and deliver their contributions all fall under the heading of relationship fitness. To measure it, Siebel Systems developed a sophisticated partner-satisfaction survey, sent each quarter to key managers of alliance partners, that contains more than 80 questions about issues such as alliance management and partners' loyalty to Siebel. The company uses this information to spot problems and to develop detailed action plans to address them.

The weight placed on each type of metric and the amount of detail it includes depend on the size and aims of the alliance. A consolidation joint venture whose main goal is to reduce costs, for instance, should focus heavily on financial and operational metrics. But managers of an alliance entering a new market expect negative financial returns in the early stages and should give more weight to strategic goals such as increasing market share and penetrating distribution channels. Smaller, short-term alliances might have simple scorecards with only four or five metrics; larger ventures with substantial assets or revenues deserve something more detailed.

Scorecard results provide important clues to what might be going wrong with an alliance, but uncovering the true problem often requires further investigation. For example, a large media company found that the hundreds of millions of dollars it had invested in alliances were at risk when close scrutiny revealed that five of its ten most important deals were losing

money. In addition, two joint ventures with an international media company were found to have been troubled by flawed deal structures from the start. Further probing found that three unprofitable alliances could be renegotiated, saving \$23 million a year, and that redefining each joint venture partner's contributions and responsibilities could save another \$45 million a year. Subsequently, the company established a corporate-level alliance unit to keep a critical eye on all of its ventures.

At a time when alliances are increasingly important, continuing to ignore their performance is simply not an option. Instead, managers should systematically measure the performance of each individual alliance to ensure that the maximum value is derived and management is able to intervene when a deal veers off track. Experience has proved that the effort pays substantial dividends. **MoF**

Jim Bamford (Jim_Bamford@McKinsey.com) is a consultant and **David Ernst** (David_Ernst@McKinsey.com) is a principal in McKinsey's Washington, D.C. office. Copyright © 2002, McKinsey & Company. All rights reserved.

The authors would like to thank David Dorton, Tammy Halevy, C. Brent Hastie, and Eric A. Kutcher for their contribution to this article.

¹ We define alliances as a broad range of collaborative arrangements involving shared objectives; shared risk, reward, or both; and a significant degree of coordination or integration. Alliances involve more shared decision-making than do arm's-length contracts and lack the full control and integration of mergers and acquisitions.

² One study found that 51 percent of the alliances reviewed had essentially no performance metrics at all and that only 11 percent had sufficient metrics. See Jeffrey H. Dyer, Prashant Kale, and Harbir Singh, "How to make strategic alliances work," *Sloan Management Review*, Summer 2001, Volume 42, Number 4.

³ For the complete discussion of alliance portfolio and strategic alliance performance, see "Managing an alliance portfolio." *The McKinsey Quarterly*, 2002, No. 3.

The real cost of equity

The inflation-adjusted cost of equity has been remarkably stable for 40 years, implying a current equity risk premium of 3.5 to 4 percent

Marc H. Goedhart, Timothy M. Koller, and Zane D. Williams

As central as it is to every decision at the heart of corporate finance, there has never been a consensus on how to estimate the cost of equity and the equity risk premium.¹

Conflicting approaches to calculating risk have led to varying estimates of the equity risk premium from 0 percent to 8 percent—although most practitioners use a narrower range of 3.5 percent to 6 percent. With expected returns from long-term government bonds currently about 5 percent in the US and UK capital markets, the narrower range implies a cost of equity for the typical company of between 8.5 and 11.0 percent. This can change the estimated value of a company by more than 40 percent and have profound implications for financial decision making.

Discussions about the cost of equity are often intertwined with debates about where the stock market is heading and whether it is over- or undervalued. For example, the run-up in stock prices in the late 1990s prompted two contradictory points of view. On the one hand, as prices soared ever higher, some investors expected a new era of higher equity returns driven by increased future productivity and economic growth. On the other hand, some analysts and academics suggested that the rising stock prices meant that the risk premium was declining. Pushed to the extreme, a few analysts even argued that the

premium would fall to zero, that the Dow Jones industrial average would reach 36,000 and that stocks would earn the same returns as government bonds. While these views were at the extreme end of the spectrum, it is still easy to get seduced by complex logic and data.

We examined many published analyses and developed a relatively simple methodology that is both stable over time and overcomes the shortcomings of other models. We estimate that the real, inflation-adjusted cost of equity has been remarkably stable at about 7 percent in the US and 6 percent in the UK since the 1960s. Given current, real long-term bond yields of 3 percent in the US and 2.5 percent in the UK, the implied equity risk premium is around 3.5 percent to 4 percent for both markets.

The debate

There are two broad approaches to estimating the cost of equity and market risk premium. The first is historical, based on what equity investors have earned in the past. The second is forward-looking, based on projections implied by current stock prices relative to earnings, cash flows, and expected future growth.

The latter is conceptually preferable. After all, the cost of equity should reflect the return expected (required) by investors. But forward-


looking estimates are fraught with problems, the most intractable of which is the difficulty of estimating future dividends or earnings growth. Some theorists have attempted to meet that challenge by surveying equity analysts, but since we know that analyst projections almost always overstate the long-term growth of earnings or dividends,² analyst objectivity is hardly beyond question. Others have built elaborate models of forward-looking returns, but such models are typically so complex that it is hard to draw conclusions or generate anything but highly unstable results. Depending on the modeling assumptions, recently published research suggests market risk premiums between 0 and 4 percent.³

Unfortunately, the historical approach is just as tricky because of the subjectivity of its assumptions. For example, over what time period should returns be measured—the previous 5, 10, 20, or 80 years or more? Should average returns be reported as arithmetic or geometric means? How frequently should average returns be sampled? Depending on the answers, the market risk premium based on historical returns can be estimated to be as high as 8 percent.⁴ It is clear that both historical and forward-looking approaches, as practiced, have been inconclusive.

Overcoming the typical failings of economic models

In modeling the behavior of the stock market over the last 40 years,⁵ we observed that many real economic variables were surprisingly stable over time (including long-term growth in corporate profits and returns on capital) and that much of the variability in stock prices related to interest rates and inflation (Exhibit 1). Building on these findings, we

Exhibit 1. US median P/E vs. inflation




Source: McKinsey analysis

developed a simple, objective, forward-looking model that, when applied retrospectively to the cost of equity over the past 40 years, yielded surprisingly stable estimates.

Forward-looking models typically link current stock prices to expected cash flows by discounting the cash flows at the cost of equity. The implied cost of equity thus becomes a function of known current share values and estimated future cash flows (see sidebar, “Estimating the cost of equity”). Using this standard model as the starting point, we then added three unique characteristics that we believe overcome the shortcomings of many other approaches:


1. *Median stock price valuation.* For the US, we used the value of the median company in the S&P 500 measured by P/E ratio as an estimate of the market’s overall valuation at any point in time. Most researchers have used the S&P 500 itself, but we argue that the S&P 500 is a value-weighted index that has been distorted at times by a few highly valued companies, and therefore does not properly

Exhibit 2. Return on book equity (ROE)



Source: McKinsey analysis

Exhibit 3. Annual estimates of the real cost of equity



Source: McKinsey analysis

reflect the market value of typical companies in the US economy. During the 1990s, the median and aggregate P/E levels diverged sharply. Indeed by the end of 1999, nearly 70 percent of the companies in the S&P 500 had P/E ratios below that of the index as a whole. By using the median P/E ratio, we believe we generate estimates that are more representative for the economy as a whole. Since UK indices have not been similarly distorted, our estimates for the UK market are based instead on aggregate UK market P/E levels.

2. *Dividendable cash flows.* Most models use the current level of dividends as a starting point for projecting cash flows to equity. However, many corporations have moved from paying cash dividends to buying back shares and finding other ways to return cash to shareholders, so estimates based on ordinary dividends will miss a substantial portion of what is paid out. We avoid this by discounting not the dividends paid but the cash flows available to shareholders after new investments


have been funded. These are what we term “dividendable” cash flows to investors that might be paid out through share repurchases as ordinary dividends, or temporarily held as cash at the corporate level.

We estimate dividendable cash flows by subtracting the investment required to sustain the long-term growth rate from current year profits. This investment can be shown to equal the projected long-term profit growth (See sidebar, “Estimating the cost of equity”) divided by the expected return on book equity. To estimate the return on equity (ROE), we were able to take advantage of the fact that US and UK companies have had fairly stable returns over time. As Exhibit 2 shows, the ROE for both US and UK companies has been consistently about 13 percent per year,⁶ the only significant exception being found in UK returns of the late 1970s.

3. *Real earnings growth based on long-term trends.* The expected growth rate in cash flow

The stability of the implied inflation-adjusted cost of equity is striking. Despite a handful of recessions and financial crises over the past 40 years . . . equity investors have continued to demand about the same cost of equity in inflation-adjusted terms.

Exhibit 4. Decomposition of the inflation-adjusted cost of equity



Source: McKinsey analysis

and earnings was estimated as the sum of long-term real GDP growth plus expected inflation. Corporate profits have remained a relatively consistent 5.5 percent of US GDP over the past 50 years. Thus, GDP growth rates are a good proxy for long-term corporate profit growth. Real GDP growth has averaged about 3.5 percent per year over the last 80 years for the US and about 2.5 percent over the past 35 years for the UK. Using GDP growth as a proxy for expected earnings growth allows us to avoid using analysts' expected growth rates.

We estimated the expected inflation rate in each year as the average inflation rate experienced over the previous five years.⁷ The nominal growth rates used in the model for each year were the real GDP growth combined with the contemporary level of expected inflation for that year.

Results

We used the above model to estimate the inflation-adjusted cost of equity implied by stock market valuations each year from 1963 to 2001 in the US and from 1965 to

2001 for the UK (Exhibit 3). In the US, it consistently remains between 6 and 8 percent with an average of 7 percent. For the UK market, the inflation-adjusted cost of equity has been, with two exceptions, between 4 percent and 7 percent and on average 6 percent.

The stability of the implied inflation-adjusted cost of equity is striking. Despite a handful of recessions and financial crises over the past 40 years including most recently the dot.com bubble, equity investors have continued to demand about the same cost of equity in inflation-adjusted terms. Of course, there are deviations from the long-term averages but they aren't very large and they don't last very long. We interpret this to mean that stock markets ultimately understand that despite ups and downs in the broad economy, corporate earnings and economic growth eventually revert to their long-term trend.

We also dissected the inflation-adjusted cost of equity over time into two components: the inflation-adjusted return on government bonds and the market risk premium. As Exhibit 4 demonstrates, from 1962 to 1979 the expected

Estimating the cost of equity

To estimate the cost of equity, we began with a standard perpetuity model:

$$P_t = \frac{CF_{t+1}}{k_e - g} \quad (1)$$

where P_t is the price of a share at time t , CF_{t+1} is the expected cash flow per share at time $t + 1$, k_e is the cost of equity, and g is the expected growth rate of the cash flows. The cash flows, in turn, can be expressed as earnings, E , multiplied by the payout ratio:

$$CF = E(\text{payout ratio})$$

Since the payout ratio is the share of earnings left after reinvestment, replacing the payout ratio with the reinvestment rate gives:

$$CF = E(1 - \text{reinvestment rate})$$

The reinvestment rate, in turn, can be expressed as the ratio of the growth rate, g , to the expected return on equity:

$$\text{reinvestment rate} = \frac{g}{ROE}$$

And thus the cash flows can be expressed as:

$$CF = E \left(1 - \frac{g}{ROE} \right) \quad (2)$$

We then combined formulas (1) and (2) to get the following:

$$\frac{P_t}{E_{t+1}} = \frac{1 - \frac{g}{ROE}}{k_e - g} \Rightarrow k_e = \frac{E_{t+1}}{P_t} \left(1 - \frac{g}{ROE} \right) + g \quad (3)$$

If the inflation embedded in k_e and g is the same, we can then express equation 3 as:

$$k_{er} = \frac{E_{t+1}}{P_t} \left(1 - \frac{g}{ROE} \right) + g_r \quad (4)$$

Where k_{er} and g_r are the inflation-adjusted cost of equity and real growth rate, respectively. We then solved for k_{er} for each year from 1963 through 2001, using the assumptions described in the text of the article.

inflation-adjusted return on government bonds appears to have fluctuated around 2 percent in the US and around 1.5 percent in the UK. The implied equity risk premium was about 5 percent in both markets.⁸ But in the 1990s, it appears that the inflation-adjusted return on both US and UK government bonds may have risen to 3 percent, with the implied equity risk premium falling to 3 percent and 3.6 percent in the UK and US respectively.

We attribute this decline not to equities becoming less risky (the inflation-adjusted cost of equity has not changed) but to investors demanding higher returns in real terms on government bonds after the inflation shocks of the late 1970s and early 1980s. We believe

that using an equity risk premium of 3.5 to 4 percent in the current environment better reflects the true long-term opportunity cost for equity capital and hence will yield more accurate valuations for companies. **MoF**

Marc H. Goedhart (Marc_Goedhart@McKinsey.com) is associate principal in McKinsey's Amsterdam office, **Timothy M. Koller** (Tim_Koller@McKinsey.com) is a principal in McKinsey's New York office, and **Zane D. Williams** (Zane_Williams@McKinsey.com) is a consultant in McKinsey's Washington, D.C., office, Copyright © 2002 McKinsey & Company. All rights reserved.

¹ Defined as the difference between the cost of equity and the returns investors can expect from supposedly risk-free government bonds.

² See Marc H. Goedhart, Brendan Russel, and Zane D. Williams, "Prophets and profits?" *McKinsey on Finance*, Number 2, Autumn 2001.

³ See, for example, Eugene Fama and Kenneth French, "The Equity Premium," *Journal of Finance*, Volume LVII, Number 2, 2002; and Robert Arnott and Peter Bernstein, "What Risk Premium is 'Normal'," *Financial Analysts Journal*, March/April, 2002; James Claus and Jacob Thomas, "Equity premia as low as three percent?" *Journal of Finance*, Volume LVI, Number 5, 2001.

⁴ See, for example, *Ibbotson and Associates*, Stock, Bonds, Bills and Inflation: 1997 Yearbook.

⁵ See Timothy Koller and Zane Williams, "What happened to the bull market?" *McKinsey on Finance*, Number 1, Summer 2001.

⁶ One consequence of combining a volatile nominal growth rate (due to changing inflationary expectations) with a stable ROE is that the estimated reinvestment rate varies tremendously over time. In the late 1970s, in fact, our estimates are near 100 percent. This is unlikely to be a true representation of actual investor expectations at the time. Instead, we believe it likely that investors viewed the high inflation of those years as temporary. As a result, in all of our estimates, we capped the reinvestment rate at 70 percent.

⁷ This assumption is the one that we are least comfortable with, but our analysis seems to suggest that markets build in an expectation that inflation from the recent past will continue (witness the high long-term government bond yields of the late 1970s).

⁸ There is some evidence that the market risk premium is higher in periods of high inflation and high interest rates, as was experienced in the late 1970s and early 1980s.

McKinsey on Finance is a quarterly publication written by experts and practitioners in McKinsey & Company's Corporate Finance & Strategy Practice. It offers readers insights into value-creating strategies and the translation of those strategies into stock market performance. This and archive issues of *McKinsey on Finance* are available on line at <http://www.corporatefinance.mckinsey.com>

McKinsey & Company is an international management consulting firm serving corporate and government institutions from 85 offices in 44 countries.

Editorial Board: Marc Goedhart, Bill Javetski, Timothy Koller, Michelle Soudier, Dennis Swinford

Editorial Contact: McKinsey_on_Finance@McKinsey.com

Editor: Dennis Swinford

Managing Editor: Michelle Soudier

External Relations: Joan Horrwich

Editorial Consultant: Janet Bush

Design and Layout: Kim Bartko

Copyright © 2002 McKinsey & Company. All rights reserved.

Cover images, left to right: ©Rob Colvin/Artville; © Garrian Manning/Artville; © Stephanie Carter/Artville; © Todd Davidson/The Image Bank, © Jack Star/PhotoLink.

This publication is not intended to be used as the basis for trading in the shares of any company or undertaking any other complex or significant financial transaction without consulting with appropriate professional advisers.

No part of this publication may be copied or redistributed in any form without the prior written consent of McKinsey & Company.