

# GRAHAM SMART MEGGINSON

THIRD



EDITION



## CORPORATE FINANCE

LINKING THEORY TO WHAT COMPANIES DO


equal to  $\bar{\sigma}^2$  and that, across any pair of stocks (say, stock  $i$  and stock  $j$ ), the average covariance is  $\bar{\sigma}_{ij}$ . Then the portfolio variance equation can be written as shown at the bottom of Figure 5.7.

As the number  $N$  of stocks in the portfolio becomes very large, the variance term  $N(1/N)^2 \bar{\sigma}^2$  approaches zero. This means the average variance of individual stocks has no impact on portfolio variance. As  $N$  increases, the second term in the equation converges to  $\bar{\sigma}_{ij}$ , indicating that what really determines the risk of a large portfolio is the average covariance between all pairs of securities. A large portfolio consisting of securities that are, on average, only weakly correlated with each other will have a lower variance than a portfolio that consists of highly correlated securities.

Figure 5.8 plots the relationship between the number of securities in a portfolio and the portfolio's variance given by this equation. For investors, the figure contains both good and bad news. The good news is that as the number of securities in the portfolio increases, the portfolio's variance declines. Given the proliferation of low-cost mutual funds available today, investors can construct portfolios containing hundreds of securities, thereby reducing the variance of their investment portfolio to some degree. The bad news is that the marginal risk reduction benefit of adding more securities to the portfolio decreases as the number of securities in the portfolio increases. Not even a very well diversified portfolio can eliminate all risk.

**FIGURE 5.8**  
Effect of Diversification on Portfolio Variance

Adding more securities to a portfolio lowers the portfolio's volatility, but the incremental benefit of adding more securities declines as the number of securities rises.



Because some risks systematically affect almost all securities, there is a limit to the risk reduction achievable by adding more securities to a portfolio. The average covariance term  $\bar{\sigma}_{ij}$  represents this limit. No matter how diversified a portfolio becomes, its variance cannot fall below the average covariance of securities in the portfolio. Financial economists give this type of risk special names: undiversifiable risk, **systematic risk**, or market

**JOB INTERVIEW QUESTION**

What is the difference between systematic and unsystematic risk?

Smart  
Ethics  
Video



Utpal Bhattacharya, Indiana  
University

"The cost of equity goes up if  
insider trading laws are not  
enforced."

See the entire interview at  
**SMARTFinance**

risk. Similarly, the risk that diversification eliminates is called diversifiable risk, **unsystematic risk**, idiosyncratic risk, or unique risk.

In real-world terms, what exactly is systematic risk? This is a difficult question to answer, and we explore it in more depth in the next chapter. For now, we just say that systematic risks are those that are common across all types of securities. Fluctuations in gross domestic product, inflation, oil prices, or interest rates can be thought of as systematic risks, and so might certain political factors. For example, the legal system governing investors and markets in a given country can influence systematic risk because that system determines the level of protection given to minority shareholders, creditors, and ordinary investors. When investors perceive that the legal system protects their interests, their willingness to trade and invest in securities increases and so the returns they require for bearing risk decline.

If investors can cheaply eliminate some risks through diversification, then we should not expect a security to earn higher returns for risks that can be eliminated through diversification. Investors can expect compensation only for bearing systematic risk (i.e., risk that cannot be diversified away). Refer back to the Example following Equation 5.4, which showed that the average return on Merck stock was about the same over 10 years as the average return on the S&P 500 *even though* Merck stock was much more volatile than the index. An undiversified investor who held only Merck stock had to bear twice as much volatility as an investor who owned the S&P 500, even though both investors earned the same reward. This is not to say that Merck was or is a bad investment. The point is that holding Merck (or any other individual company's stock) *in isolation* is a poor investment strategy. Undiversified portfolios are generally suboptimal because they expose investors to unsystematic risk without offering higher returns.

### Measuring the Systematic Risk of an Individual Security

The previous section demonstrated two important facts. First, the formula for portfolio variance shows that each security contributes to a portfolio's risk through two channels, the security's own variance and its covariance with all other securities in the portfolio. In diversified portfolios, only the second channel matters. This implies that an individual stock's variance may be a poor measure of its risk. The variance of a stock captures its total volatility, some of which is unsystematic and some of which is systematic. Second, because diversification eliminates unsystematic risk, the market provides no reward for bearing it. As a consequence, though we still expect to see a positive relationship in the market between risk and return, we can no longer be confident that a positive relationship will exist between returns on an individual asset and its variance. Again, a stock's variance captures both its systematic and unsystematic fluctuations, but only the systematic component should be correlated with returns.

We need a new measure for an individual asset's risk, one that captures only the systematic component of its volatility. Remember, the primary contribution to portfolio risk from a single asset comes from its covariance with all the other assets in the portfolio. Imagine that an investor holds a fully diversified portfolio—literally, a portfolio containing every asset available in the economy. How would this investor determine the contribution of a single security to the portfolio's risk? One way to do that would be to measure the covariance between a single asset and the portfolio. Recall, though, the difficulty that nonstandard units cause for interpreting covariance calculations. A standardized measure would be preferable, and finance theory gives us just such a measure in the concept of *beta*:

$$\beta_i = \frac{\sigma_{im}}{\sigma_m^2}$$

(Eq. 5.14)

The **beta** of asset  $i$  ( $\beta_i$ ) equals the covariance of the asset's returns with the returns on the overall portfolio, divided by the portfolio's variance. As you will see in the next chapter, the portfolio we refer to here is known as "the market portfolio," a value-weighted portfolio of all available assets.<sup>14</sup> A security's beta gives us a standardized measure of its covariance with all other assets, or a measure of its systematic risk. If the market rewards only systematic risk and if beta captures the systematic risk of an individual asset, then we should observe a positive relationship between values of beta and returns in the market.

Observe that the formula for an asset's beta closely resembles that of the correlation coefficient:

$$\beta_i = \frac{\sigma_{im}}{(\sigma_m)(\sigma_m)};$$

$$\rho_i = \frac{\sigma_{im}}{(\sigma_i)(\sigma_m)}$$

The equations are identical except in one respect: The denominator of the correlation coefficient multiplies the standard deviations of the asset and the market, whereas the denominator of the beta formula squares the standard deviation of the market. This small adjustment to the denominator makes the interpretation of beta a little different from that of the correlation coefficient. First, unlike  $\rho$ , beta has no maximum or minimum value. Second, beta indicates how much the individual asset's return moves, on average, when the market moves by 1%. For example, if a stock has a beta of 1.5, then, when the market return increases by 1%, the stock return will (on average) increase by 1.5%.

### EXAMPLE

Now that we understand the beta measure of a stock's risk, how does it compare to the measure we started with, standard deviation? Comparing the monthly returns on each of the four stocks listed in Table 5.3 to returns on the overall stock market, suppose you calculate the following statistics:

Stock	Covariance with Market
Mead	0.0031
Boise	0.0026
Nike	0.0011
Arrow	-0.0003

If the variance of market returns were 0.0028, then the betas of the four stocks would be as follows:

Mead 1.11    Boise 0.93    Nike 0.39    Arrow -0.11

These betas contain several surprises. First, based on comparison of the standard deviations of each stock in Table 5.3, we concluded that Nike was the riskiest security. Comparing the betas, however, suggests that Nike is less risky than either Mead or Boise Cascade. Recall

(continued)

<sup>14</sup>The modifier "value-weighted" means that the fraction invested in a particular security is equal to that security's total market value as a percentage of the market value of all securities. For example, if the total market value of all securities in the market is \$10 trillion and if the total market value of a certain company's stock equals \$100 billion, then the fraction of that stock in a value-weighted portfolio would be 0.01, or 1%.

## 6.4 THE CAPITAL ASSET PRICING MODEL (CAPM)

### The Security Market Line

The basic CAPM was developed almost simultaneously during the mid-1960s by William Sharpe (1964), John Lintner, and Jan Mossin (1966); it was quickly embraced by academic researchers and, in time, by practitioners as well. The reason for the CAPM's widespread acceptance is not hard to understand—for the first time, researchers and practitioners had a model that generated specific predictions about the risk–return characteristics of individual assets, and this relation was driven by how each asset *covaries* with the market portfolio.

The formal development of the CAPM requires several assumptions about investors and markets. Rather than present a detailed list of these assumptions, we present the logic of the CAPM as it flows from the material we have covered so far.

1. Investors are risk averse and require higher returns on riskier investments.
2. Because investors can diversify, they care only about the systematic (or undiversifiable) risk of any investment.
3. The market offers no reward for bearing unsystematic risk because it can be diversified away.
4. Some portfolios are better than others. Portfolios that maximize expected return at any level of risk are efficient portfolios.
5. If investors can borrow and lend at the risk-free rate, then there exists a single risky portfolio that dominates all others. Only portfolios consisting of the risk-free asset and the optimal risky portfolio are efficient.
6. If investors have homogeneous expectations then they will agree on the composition of the optimal portfolio. In equilibrium, the optimal portfolio will be the market portfolio.
7. The central insight of the CAPM is that if all investors hold the market portfolio then—when evaluating the risk of any specific asset—they will be concerned with the covariance of that asset with the overall market. The implication is that any measure of an asset's systematic risk exposure must capture how it covaries with the rest of the market. An asset's beta provides a quantitative measure of this risk, and therefore the CAPM predicts a positive, linear relationship between expected return and beta. In the CAPM, beta risk (or market risk) is the only risk that is priced.

The **capital asset pricing model (CAPM)** indicates that the expected return on a specific asset,  $E(R_i)$ , equals the risk-free rate plus a premium that depends on the asset's beta,  $\beta_i$ , and the expected risk premium on the market portfolio,  $E(R_m) - R_f$ :

$$E(R_i) = R_f + \beta_i [E(R_m) - R_f] \quad (6.1)$$

Recall that beta measures an asset's correlation with a broader portfolio—in this case, the market portfolio. The higher the beta of a security, the greater the security's exposure to systematic risk and the higher the expected return it must offer investors. Although there are three variables ( $R_f$ ,  $\beta_i$ , and  $E(R_m)$ ) on the right-hand side of the CAPM equation, beta changes from one security to the next. For that reason, analysts classify the CAPM as a **single-factor model**, meaning that just one variable explains differences in returns across securities.

Figure 6.8 plots the CAPM equation on a diagram with the expected return on the y-axis and beta on the x-axis. The intercept of this line is  $R_f$ , and its slope is  $E(R_m) - R_f$ . According to the CAPM, the equilibrium expected returns of all securities must plot on

#### JOB INTERVIEW QUESTION

How would you estimate the expected return of a stock?



In the CAPM, you recall, the market portfolio is a value-weighted combination of all assets in the economy. At present, we are unaware of any market index that attempts to incorporate every type of asset. When using the CAPM, most practitioners and academics use the returns on a broad-based stock index as a proxy for the true market portfolio. Accordingly, rather than try to estimate the expected risk premium on the market portfolio, analysts usually focus on the expected equity risk premium: the difference in expected returns between a well-diversified portfolio of common stocks and a risk-free asset such as a U.S. Treasury bill.

Since 1900, the average real return on stocks outpaced the average real return on U.S. Treasury bills by about 5.4% per year. But in the CAPM, what matters is not the actual equity risk premium from the past but rather the expected equity risk premium looking forward. Though many analysts trust the historical evidence and simply plug in a value close to 6% for the term  $E(R_m - R_f)$ , a naive reliance on long-run historical averages is the only approach for estimating the expected risk premium. Using an unbiased estimate is important because an error in the risk premium translates directly into an error in a project's discount rate and thus in its NPV.

One variable that analysts can use to obtain a forward-looking estimate of the equity risk premium is the market's aggregate earnings yield, which is the reciprocal of the price-to-earnings ratios. For example, to calculate the earnings yield for the S&P 500, add up the earnings of all 500 companies and divide by the aggregate market value of these companies. Corporate earnings fluctuate with the business cycle, so analysts usually try to adjust, or *normalize*, these temporary effects before using the earnings yield to estimate the equity risk premium. In the United States, the long-run average value of the earnings yield is about 7%, a little less than the average real return on stocks. It should come as no surprise that the earnings yield is closely related to the real return on stocks. After all, stocks represent a claim on corporate earnings.

A second forward-looking method for estimating the equity risk premium is the dividend growth model. Recall that this model calculates the present value of a perpetuity of a dividend stream growing at a constant rate,  $g$ :

$$P_0 = \frac{D_1}{r - g}$$

Rearranging this equation shows that the required return on the stock equals the sum of the dividend yield and the dividend growth rate:

$$r = \frac{D_1}{P_0} + g$$

To use this model when estimating the equity risk premium, we must think in aggregate, macroeconomic terms. In other words,  $r$  represents the required return on the stock market rather than the required return on a single stock. The ratio  $D_1/P_0$  represents the aggregate dividend yield, and  $g$  represents the (real) growth rate of aggregate dividends. From 1872 to 1950, the expected equity risk premium derived from this model almost exactly matched the actual risk premium measured using average historical returns (a little more than 4%). From 1950 to the present, however, the average real return on equities was much higher than predicted by the dividend growth model.<sup>8</sup>

<sup>8</sup>The opposite has been true since 2000: real equity returns have been *lower* on average than the dividend growth model would predict.

The market portfolio is a value-weighted combination of all assets in the economy. At present, we are unaware of any market index that attempts to incorporate every type of asset. When using the CAPM, most practitioners and academics use the returns on a broad-based stock index as a proxy for the true market portfolio. Accordingly, rather than try to estimate the expected risk premium on the market portfolio, analysts usually focus on the expected equity risk premium: the difference in expected returns between a well-diversified portfolio of common stocks and a risk-free asset such as a U.S. Treasury bill.

Con  
Re  
Quest

JOB INTERVIEW  
QUESTION  
What do you believe is the equity risk premium? Why?

efficiency is more important because efficient capital markets incorporate all relevant information into financial asset prices, which in turn helps ensure that promising investments receive funding.

The concept of efficient capital markets is one of the most influential contributions that financial economics has made to modern economic thought. The **efficient markets hypothesis (EMH)**, as formally presented by Eugene Fama in 1970, has revolutionized financial thought, practice, and regulation. The EMH asserts that financial asset prices fully reflect all available information. What do we mean by “all available information”? The answer to this question varies, and we discuss three distinct versions of the efficient markets hypothesis.

### The Three Forms of Market Efficiency

The EMH presents three increasingly stringent definitions of efficiency based on the information that market prices reflect: weak-form, semistrong-form, and strong-form efficiency.

**Weak-Form Efficiency** In markets characterized by **weak-form efficiency**, asset prices incorporate all information from the historical record—that is, all information about price trends or repeating patterns that occurred in the past. This implies that trading strategies based on analyses of historical pricing trends or relationships cannot consistently yield market-beating returns.

Prices in a weak-form efficient market will be unpredictable and will change only in response to the arrival of new information. In technical terms, this means that prices follow a **random walk**: they wander aimlessly, with no connection to past price changes and no tendency to return to a mean value over time.

**Semistrong-Form Efficiency** The second form of market efficiency, **semistrong-form efficiency**, asserts that asset prices incorporate *all publicly available information*. The key point about this form of efficiency is that the prices need only reflect information from *public sources* (e.g., newspapers, press releases, computer databases).

There is both a “stock” and a “flow” aspect to the information processing capabilities of semistrong-form efficient markets: First, the *level* of asset prices should correctly reflect all pertinent historical, current, and predictable future information that investors can obtain from public sources. Second, asset prices should *change* fully and instantaneously in response to relevant new information.


**Strong-Form Efficiency** In markets characterized by **strong-form efficiency**, asset prices reflect *all* information, both public and private. This extreme form of market efficiency implies that important company-specific information will be fully incorporated in asset prices with the very first trade after the information is generated.

In strong-form efficient markets, most insider trading would be unprofitable and there would be no benefit to ferreting out information on publicly traded companies. Any data morsel so obtained would already be reflected in stock and bond prices.

Table 10.1 on page 358 describes the three forms of market efficiency and summarizes the key implications of each form.

### Does Empirical Evidence Support Market Efficiency?

Ultimately, whether financial markets are efficient is an empirical question. For more than a quarter of a century, the efficient market hypothesis enjoyed overwhelming support among financial economists. However, in recent years a large body of empirical evidence challenging the EMH has caused many former “true believers” to take a fresh look at the efficiency question. It also seems likely that the paralysis and near-collapse of global



Malcolm Baker, Harvard University  
 “There is increasing evidence that stock markets aren't perfectly efficient.”

See the entire interview at  
**SMART Finance**

(SF) or for  
 nsures that  
 nufactured  
 ?  
 oser to the  
 the same

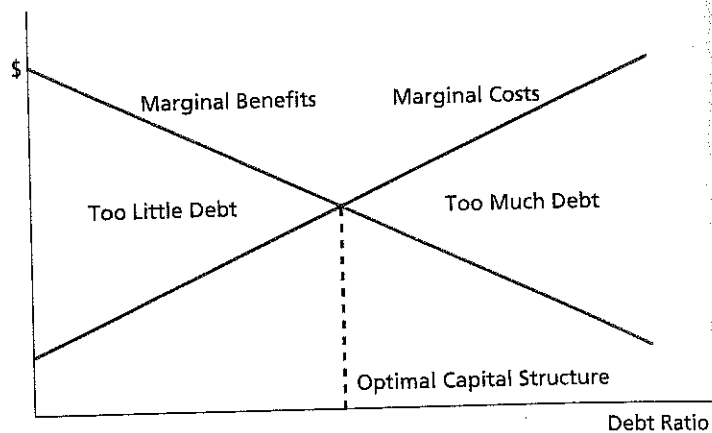
**MARKET!**

of financial  
 et to rapidly  
 formational


**FIGURE 13.2**  
**Weighing Debt's Benefits and Costs to Find Optimal Capital Structure**

The optimal amount of debt occurs where the marginal cost and marginal benefit curves intersect. At that point, total firm value is at its peak, and the weighted average cost of capital (WACC) is at its minimum.

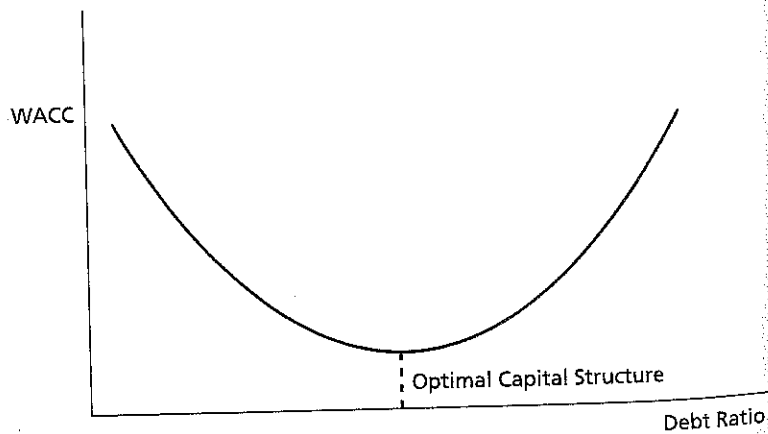
**Panel A: Optimal Capital Structure**



**Panel B: Total Firm Value**



**Panel C: WACC**





combination of equity and debt, with the debt interest sheltering cash flows from taxes. Even so, most firms do not finance their activities exclusively with debt. This suggests that managers see debt as having costs that at some point offset debt's tax advantages. Based on observing what companies actually do, the optimal capital structure for most firms is apparently one that contains some debt, but not too much.

How do managers trade-off the benefits and costs of debt to establish a target capital structure that maximizes firm value? Figure 13.2 offers a conceptual answer to this question. The blue line in Panel A shows that the marginal benefit of borrowing an additional dollar falls as the firm's overall debt ratio rises. The red line indicates that costs associated with using debt rise as leverage increases. We will explain in the next section why marginal benefits fall and marginal costs rise as debt increases, but for now you can just take the benefits and costs in Figure 13.2 as given. As in any cost-benefit analysis, the optimum occurs when marginal benefits and marginal costs are equal. Therefore, a manager facing these cost and benefit curves would choose a debt level where the two curves intersect. To the left of that point, the firm has too little debt in the sense that marginal benefits exceed marginal costs, so adding more debt would increase firm value. At higher debt levels, debt's marginal costs exceed its benefits, so adding leverage decreases firm value.

Panel B shows the relation between total firm value and leverage. If a firm has no debt, its value equals  $V_U$ . From that point, if the firm adds debt to its capital structure, its value begins to rise. At some point, firm value reaches a peak, and from that point, adding more debt decreases the value of the firm. The graph shows that, at the same point where the marginal benefit and cost curves in Panel A intersect, firm value reaches its peak. The point at which firm value begins to fall as leverage rises is exactly when debt's marginal costs first exceed its marginal benefits.

At the end of this chapter, we will demonstrate how to find the optimal debt ratio. But how much difference does finding the right capital structure really make in the overall value of the firm? In a recent study, van Binsbergen, Graham, and Yang (2008) estimate that, for the average firm, appropriate debt choices can increase firm value by about 5%. In some companies, like the one described at the beginning of Chapter 14, the increase in value may be 10% or more.

Panel C of Figure 13.2 demonstrates how a firm's weighted-average-cost of capital (WACC) changes as leverage rises. Here, the relation is U-shaped. A firm with no leverage can reduce its WACC by substituting debt for equity, but, eventually, the firm reaches a point where further increases in debt cause the WACC to increase. Naturally, managers want to find the debt ratio that minimizes the cost of capital because doing so maximizes firm value. Therefore, the optimum point in Panel C is the same optimum debt ratio in Panels A and B.

In the next section, we explore in more detail why debt's marginal benefits fall and its marginal costs rise as a firm uses more debt in its capital structure. To begin, we revisit the tax advantage of debt, taking into account some important features of the tax code that we have neglected thus far.



### Concept Review Questions

1. How large would the costs of debt have to be in order to justify a firm's decision to operate with 100% equity?
2. If a firm is operating well below its optimum debt level, then what market forces might prompt it to use more debt?

## The Investment Banker's Role in Equity Issues

We now turn to the services that investment banks provide to issuing companies, with particular attention to U.S. practices.<sup>1</sup> We focus on common stock offerings, though the procedures for selling bonds and preferred stocks are similar. Investment banks play several different roles throughout the securities offering process, and this section describes the evolution of these roles over the course of an issue. We also describe how issuers compensate IBs for the services they provide.

Although firms can issue stock without the assistance of investment bankers, in practice almost all firms hire IBs to help issue equity. Firms can choose an investment banker in one of two ways. The most common approach is a **negotiated offer**, where the issuing firm negotiates the terms of the offer directly with one investment bank. Alternatively, in a **competitively bid offer**, the firm announces the terms of its intended equity sale and then investment banks bid for the business. Intuition suggests that competitive bidding should be cheaper, but the empirical evidence is mixed. One clear sign that competitive offers are not better and cheaper is that the vast majority of equity sales are negotiated. If the costs of negotiated deals were truly higher, then why would so many firms choose that approach?


Firms issuing securities often hire more than one investment bank. In these cases, one of the banks is usually named the **lead underwriter**, or *book-runner*, while the other leading banks are called **co-managers**. Chen and Ritter (2000) argue that firms often prefer to issue securities with several co-managers because doing so increases the number of stock analysts that will follow the firm after the offering. Firms believe that a larger analyst following leads to greater liquidity and higher stock values. Cliff and Denis (2004) verify the importance of attracting the coverage of top-rated analysts by showing that issuing firms willingly allow their IPO share price to be set low enough to attract excess demand and high trading volume, since this will indirectly compensate the underwriters' star analysts.

Investment bankers sell equity under two types of contracts. In a **best-efforts offering**, the investment bank merely promises to give its best effort to sell the firm's securities at the agreed-upon price but makes no guarantee about the ultimate success of the offering. If there is insufficient demand, the firm withdraws the issue from the market. Best-efforts offerings are most commonly used for small, high-risk companies, and the IB receives a commission based on the number of shares sold.

In contrast, in a **firm-commitment offering** the investment bank agrees to underwrite the issue, meaning that the bank guarantees (**underwrites**) the offering price. The IB actually purchases the shares from the firm and resells them to investors. This arrangement requires the investment bank to bear the risk of inadequate demand for the issuer's shares, but banks mitigate this risk in two ways. First, the lead underwriter forms an **underwriting syndicate** consisting of many investment banks. These banks collectively purchase the firm's shares and market them, thus spreading the risk across the syndicate. Second, underwriters go to great lengths to determine the demand for a new issue before it comes to market, and they generally set the issue's *offer price* and take possession of the securities no more than a day or two before the issue date. These steps help ensure that the investment bank faces only a small risk of being unable to sell the shares that it underwrites.

In firm-commitment offerings, investment banks receive compensation for their services via the **underwriting spread**, the difference between the price at which the banks purchase shares from firms (the **net price**) and the price at which they sell the shares to institutional and individual investors (the **offer price**). In some offerings, underwriters receive additional compensation in the form of warrants that grant the right to buy shares

<sup>1</sup>Ljungqvist, Jenkinson, and Wilhelm (2003) and DeGeorge, Derrien, and Womack (2007) document an increasing tendency for security issues around the world to conform to U.S. standards.



David Chen, Co-head of M&A  
at Goldman Sachs in the  
Americas  
has insight at the  
intersection of industry and  
capital markets.

See the entire interview at  
**Q&A Finance**