

Rational Capital Budgeting in an Irrational World*

I. Introduction

The last several years have not been good ones for the capital asset pricing model (CAPM). A large volume of recent empirical research has found that (i) cross-sectional stock returns bear little or no discernible relationship to β and (ii) a number of other variables besides β have substantial predictive power for stock returns. For example, one variable that has been shown to be an important and reliable predictor is the book-to-market ratio: the higher a firm's book-to-market, the greater its conditional expected return, all else being equal.¹

In light of these empirical results, this article

* This research is supported by the National Science Foundation and the International Financial Services Research Center at Massachusetts Institute of Technology. Thanks to Maureen O'Donnell for assistance in preparing the manuscript and to Michael Barclay, Doug Diamond, Steve Kaplan, Jay Ritter, Dick Thaler, Luigi Zingales, an anonymous referee, and seminar participants at the National Bureau of Economic Research for helpful comments and suggestions.

1. Early work on the predictive power of variables other than β includes Stattman (1980), Banz (1981), Basu (1983), Keim (1983), DeBondt and Thaler (1985), Rosenberg, Reid, and Lanstein (1985), Bhandari (1988), and Jaffe, Keim, and Westerfield (1989). See Fama (1991) for a detailed survey of this literature. More recent papers that have focused specifically on book-to-market as a predictive variable include Chan, Hamao, and Lakonishok (1991), Fama and French (1992), Davis (1994), Lakonishok, Shleifer, and Vishny (1994), and Chan, Jegadeesh, and Lakonishok (1995).

This article addresses the following basic capital budgeting problem: suppose that cross-sectional differences in stock returns can be predicted based on variables other than β (e.g., book-to-market) and that this predictability reflects market irrationality rather than compensation for fundamental risk. In this setting, how should companies determine hurdle rates? I show how factors such as managerial time horizons and financial constraints affect the optimal hurdle rate. Under some circumstances, β can be useful as a capital budgeting tool, even if it is of no use in predicting stock returns.

addresses a simple, yet fundamental question: What should the academic finance profession be telling MBA students and practitioners about how to set hurdle rates for capital budgeting decisions? Can we still follow the standard textbook treatments with a clear conscience and a straight face, and march through the mechanics of how to do weighted-average-cost-of-capital or adjusted-present-value calculations based on the CAPM? Or should we abandon the CAPM for capital budgeting purposes in favor of alternative models that seem to do a better job of fitting actual stock-return data?

If one believes that the stock market is efficient and that the predictable excess returns documented in recent studies are, therefore, just compensation for risk—risk that is for some reason not well captured by β —then the answer to the question is simple. According to standard finance logic, in an efficient market the hurdle rate for an investment in any given asset should correspond exactly to the prevailing expected return on the stock of a company that is a pure play in that asset. The only operational question is, Which regression specification gives the best estimates of expected return? Thus the inevitable conclusion is that one must throw out the CAPM and in its place use the “new and improved” statistical model to set hurdle rates. For shorthand, I will label this approach to setting hurdle rates the NEER approach, for new estimator of expected return.

As an example of the NEER approach, consider a chemical company that currently has a low book-to-market ratio, and hence—according to an agreed-upon regression specification—a low expected return. If the company were considering investing in another chemical plant, this approach would argue for a relatively low hurdle rate. The implicit economic argument is this: the low book-to-market ratio is indicative of the low risk of chemical-industry assets. Given this low risk, it makes sense to set a low hurdle rate. Of course, if the chemical company’s book-to-market ratio—and hence its expected return—were to rise over time, then the hurdle rate for capital budgeting purposes would have to be adjusted upward.

This NEER approach to capital budgeting is advocated by Fama and French (1993). Fama and French couch the predictive content of the book-to-market ratio and other variables in a linear multifactor-model setting that they argue can be interpreted as a variant of the arbitrage pricing theory (APT) or intertemporal capital asset pricing model (ICAPM). They then conclude: “In principle, our results can be used in any application that requires estimates of expected stock returns. This list includes . . . estimating the cost of capital” (p. 53).

However, it is critical to the Fama-French logic that the return differentials associated with the book-to-market ratio and other predictive variables be thought of as compensation for fundamental risk. While there seems to be fairly widespread agreement that variables such as book-to-market do indeed have predictive content, it is much less clear

that this reflects anything having to do with risk. Indeed, several recent papers find that there is very little affirmative evidence that stocks with high book-to-market ratios are riskier in any measurable sense.²

An alternative interpretation of the recent empirical literature is that investors make systematic errors in forming expectations, so that stocks can become significantly over- or undervalued at particular points in time. As these valuation errors correct themselves, stock returns will move in a partially predictable fashion. For example, a stock that is overvalued relative to fundamentals will tend to have a low book-to-market ratio. Over time, this overvaluation will work its way out of the stock price, so the low book-to-market ratio will predict relatively low future returns.³

In a world where predictable variations in returns are driven by investors' expectational errors, it is no longer obvious that one should set hurdle rates using a NEER approach. The strong classical argument for such an approach rests in part on the assumption that there is a one-to-one link between the expected return on a stock and the fundamental economic risk of the underlying assets. If this assumption is not valid, the problem becomes more complicated. Think back to the example of the chemical company with the low book-to-market ratio. If the low value of the ratio—and hence the low expected return—is not indicative of risk, but rather reflects the fact that investors are currently overoptimistic about chemical-industry assets, does it make sense for rational managers to set a low hurdle rate and invest aggressively in acquiring more such assets?

The key point is that, when the market is inefficient, there can be a meaningful distinction between a NEER approach to hurdle rates, and one that focuses on a measure of fundamental asset risk—which I will call a FAR approach. A FAR approach involves looking directly at the variance and covariance properties of the cash flows on the assets in question. In the chemical company case, this might mean assigning a high hurdle rate if, for example, the cash flows from the new plant were highly correlated with the cash flows on other assets in the economy—irrespective of the company's current book-to-market ratio or the conditional expected return on its stock.

Given this distinction between the NEER and FAR approaches, the main goal of this article is to assess the relative merits of the two, and to illustrate how one or the other may be preferred in a given set of circumstances. To preview the results, I find that, loosely speaking, a

2. See, e.g., Lakonishok, Shleifer, and Vishny (1994), Daniel and Titman (1995), and MacKinlay (1995). The Daniel and Titman paper takes direct issue with the Fama-French (1993) notion that the book-to-market effect can be given a multifactor risk interpretation.

3. For evidence supporting this interpretation, see Lakonishok et al. (1994) and La-Porta et al. (1995).

NEER approach makes the most sense when either (i) managers are interested in maximizing short-term stock prices or (ii) the firm faces financial constraints (in a sense which I will make precise). In contrast, a FAR approach is preferable when managers are interested in maximizing long-run value and the firm is not financially constrained.

The fact that a FAR approach can make sense in some circumstances leads to an interesting and somewhat counterintuitive conclusion: in spite of its failure as an *empirical description* of actual stock returns, the CAPM (or something quite like it) may still be quite useful from a *prescriptive* point of view in capital budgeting decisions. This is because β —if calculated properly—may continue to be a reasonable measure of the fundamental economic risk of an asset, even if it has little or no predictive power for stock returns. However, it must be emphasized that this sort of rationale for using the CAPM does not apply in all circumstances. As noted above, when managers have short horizons or when the firm faces financial constraints, the CAPM—or any FAR-based approach—will be inappropriate to the extent that it does not present an accurate picture of the expected returns on stocks.

Before proceeding, I should also reiterate that the entire analysis that follows is based on the premise that the stock market is inefficient. More precisely, cross-sectional differences in expected returns will be assumed throughout to be driven in part by expectational errors on the part of investors. I make this assumption for two reasons. First, it strikes me that, at the least, there is enough evidence at this point for one to question market efficiency seriously and to wonder what capital budgeting rules would look like in its absence. Second, as discussed above, the efficient-markets case is already well understood and there is little to add. In any event, however, readers can judge for themselves whether or not they think the inefficient-markets premise is palatable as a basis for thinking about capital budgeting.

Of course, this is not the first paper to raise the general question of whether and how stock market inefficiencies color investment decisions. This question dates back at least to Keynes (1936, p. 151), who raises the possibility that “certain classes of investment are governed by the average expectation of those who deal on the Stock Exchange as revealed in the price of shares, rather than by the genuine expectations of the professional entrepreneur.”⁴ More recent contributions include Bosworth (1975), Fischer and Merton (1984), DeLong et al. (1989), Morck, Shleifer, and Vishny (1990) and Blanchard, Rhee, and Summers (1993). The latter two pieces are particularly noteworthy in that they stress—as does this article—the importance of managers’ time horizons and financing constraints. The contribution of this article

4. As will become clear, in terms of the language of this article Keynes is effectively expressing the concern that managers will adopt a NEER approach to capital budgeting.

relative to these earlier works is twofold: first, it provides a simple analytical framework in which the effects of horizons and financing constraints can be seen clearly and explicitly, and, second, it focuses on the question of what are appropriate *risk-adjusted* hurdle rates, thereby developing an inefficient-markets analog to textbook treatments of capital budgeting under uncertainty.⁵

The remainder of the article is organized as follows. In Section II, I examine the link between managers' time horizons and hurdle rates, leaving aside for simplicity the issue of financial constraints. This section establishes that a FAR approach is more desirable when the time horizon is longer. In Section III, I introduce the possibility of financial constraints and show how these have an effect similar to shortening the time horizon—that is, financial constraints tend to favor a NEER approach. In Section IV, I take up measurement issues. Specifically, if one decides to use a FAR approach, what is the best way to get an empirical handle on fundamental asset risk? To what extent can one rationalize the use of β —as conventionally calculated—as an attempt to implement a FAR approach? Section V briefly discusses a number of extensions and variations of the basic framework, and Section VI fleshes out its empirical implications. Section VII concludes the discussion.

II. Time Horizons and Optimal Hurdle Rates

I consider a simple 2-period capital budgeting model that is fairly standard in most respects. At time 0, the firm in question is initially all equity financed. It already has physical assets in place that will produce a single net cash flow of F at time 1. From the perspective of time 0, F is a random variable that is normally distributed. The firm also has the opportunity to invest \$1 more at time 0 in identical assets—that is, if it chooses to invest, its physical assets will yield a total of $2F$ at time 1. The decision of whether or not to invest is the only one facing the firm at time 0. If it does invest, the investment will be financed with riskless debt that is fairly priced in the marketplace. There is no possibility of issuing or repurchasing shares. (As will become clear in Section III, allowing the firm to transact in the equity market at time 0 may in some circumstances alter the results.) There are also no taxes.

The manager of the firm is assumed to have rational expectations. I denote the manager's rational time-0 forecast of F by $F^r \equiv EF$. The firm's other outside shareholders, however, have biased expectations. Their biased forecast of F is given by $F^b \equiv EF(1 + \delta)$. Thus δ is a

5. In contrast, the informal discussion in Blanchard et al. (1993) assumes risk neutrality and, hence, does not speak to the whole issue of risk-adjusted hurdle rates.

measure of the extent to which outside investors are overoptimistic about the prospects for the firm's physical assets. This bias in assessing the value of physical assets is the only way in which the model departs from the standard framework. Outside investors are perfectly rational in all other respects. For example, they perceive all variances and covariances accurately. Thus the degree of irrationality that is being ascribed to outside investors is in a sense quite mild. It would certainly be interesting to entertain alternative models of such irrationality, but in the absence of clear-cut theoretical guidance, the simple form considered here seems a natural place to start.

The shares in the firm are part of a larger market portfolio. The net cash flow payoff on the market portfolio at time 1 is given by M , which is also normally distributed. For simplicity, I assume that both the manager of the firm and all outside investors have rational expectations about M . Thus I am effectively assuming that investors make firm-level mistakes in assessing cash flows, but that these mistakes wash out across the market as a whole.⁶ I denote the price of the market portfolio at time 0 by P_M , and define $R_M \equiv M/P_M - 1$ as the realized percentage return on the market.

The final assumption is that the price of the firm's shares is determined solely by the expectations of the outside investors. This amounts to saying that even though the manager may have a different opinion, he is unable or unwilling to trade in sufficient quantity to affect the market price.

With the assumptions in place, the first thing to do is to calculate the initial market price, P , of the firm's shares, *before* the investment decision has been made at time 0. This is an easy task. Note that we are operating in a standard mean-variance framework, with the only exception being that investors have biased expectations. This bias does not vitiate many of the classical results that obtain in such a framework. First of all, investors will all hold the market portfolio, and the market portfolio will be—in their eyes—mean-variance efficient. Second, the equilibrium return required by investors in firm's equity, k , will be given by

$$k = r + \beta^r(ER_M - r), \quad (1)$$

where r is the riskless rate and β^r is the usual "rate-of-return β ," defined as

$$\beta^r \equiv \text{cov}(F/P, R_M)/\text{var}(R_M). \quad (2)$$

Thus, outside investors' required returns are determined as in the standard CAPM. Given the cash flow expectations of these investors,

6. Or, said somewhat more mildly, the manager of the firm does not disagree with outside investors' assessment of M .

the initial price of the firm's shares, P , satisfies

$$P = F^b/(1 + k). \quad (3)$$

Equation (3) is not a completely reduced form, however, because P appears in the definition of k and hence is on both sides of the equation. Rearranging terms, we obtain the following expression for P in terms of primitive parameters:

$$P = \{F^b - \beta^d(ER_M - r)\}/(1 + r), \quad (4)$$

where β^d is the "dollar β ," defined as

$$\beta^d \equiv \text{cov}(F, R_M)/\text{var}(R_M). \quad (5)$$

It is useful to compare equations (1)–(4) with the analogous expressions that would prevail in a classical setting with rational expectations. Using asterisk superscripts to denote these (unobserved) rational expectations values, we have

$$k^* = r + \beta^*(ER_M - r), \quad (6)$$

$$\beta^* \equiv \text{cov}(F/P^*, R_M)/\text{var}(R_M), \quad (7)$$

$$P^* = F^r/(1 + k^*), \quad (8)$$

and

$$P^* = \{F^r - \beta^d(ER_M - r)\}/(1 + r). \quad (9)$$

From a comparison of equations (4) and (9), it can be seen that in the reduced form, the only difference between P and P^* is the bias in the expected cash flow term.

While outside investors perceive that the firm's stock will yield an expected return of k , that is not the rational expectation of the stock's performance. Rather, the best estimate of conditional expected return, which I will denote by CER, is

$$\text{CER} = F^r/P - 1 = (1 + k)/(1 + \delta) - 1. \quad (10)$$

Thus from the perspective of a rational observer such as the firm's manager, the stock may have a CER that is either greater or less than the CAPM rate k . In this sense, the model crudely captures that empirical regularity that there are predictable returns on stocks that are not related to their β 's. These predictable returns simply reflect the biases of the outside investors. Note that when $\delta = 0$, so that outside investors have no bias, $\text{CER} = k = k^*$. When $\delta > 0$, so that the stock is overpriced, $\text{CER} < k < k^*$. And, conversely, when $\delta < 0$, $\text{CER} > k > k^*$.

We are now ready to address the question of the optimal hurdle rate. To do so, we have to be clear about the objective function that is being maximized. There are two distinct possibilities. First, one

might assume that the manager seeks to maximize the stock price that prevails at time 0, immediately after the investment decision has been made. This is the same as saying that the manager tries to maximize outside investors' perception of value. Alternatively, one might posit that the manager seeks to maximize the present value of the firm's future cash flows, as seen from his more rational perspective.

In principle, one can think of reasons why managers might tend to favor either objective. For example, if they are acting on behalf of shareholders (including themselves) who have to sell their stock in the near future for liquidity reasons, they will be more inclined to maximize current stock prices. In contrast, if they are acting on behalf of shareholders (including themselves) who will be holding for the longer term—for example, due to capital gains taxes or other frictions—they will be more inclined to maximize the present value of future cash flows. In what follows, I treat the managerial time horizon as exogenous, although in a fuller model it would be endogenously determined.⁷

Let us first consider the “short-horizon” case, in which the goal is to maximize the current stock price. It is easy to see that in this case, the “value” created by investing is simply $(P - 1)$. Intuitively, as long as the market's current valuation of the assets in question exceeds the acquisition cost, the current stock price will be increased if the assets are purchased. To translate this into a statement about hurdle rates, note that, from management's perspective, the expected cash flow on the investment is F' . Thus the short-horizon hurdle rate, defined as h^S , has the property that the gross discounted value of the investment, $F'/(1 + h^S)$, equals P . Using equation (10), it follows immediately that

PROPOSITION 1. In the short-horizon case, the manager should discount his expected cash flow F' at a hurdle rate $h^S = \text{CER}$. In other words, the manager should take a NEER approach and use the conditional expected return on the stock as the hurdle rate.

One way to think about proposition 1 is that, if the manager is interested in maximizing the current stock price, he must cater to any misperceptions that investors have. Thus if investors are overly optimistic about the prospects for the firm's assets—thereby leading to a low value of CER—the manager should be willing to invest very aggressively in these assets and hence should adopt a low hurdle rate.

Things work quite differently in the “long-horizon” case, in which

7. This distinction between maximizing current stock prices vs. maximizing management's perception of long-run value also arises in the literature on investment and financing decisions under asymmetric information. See, e.g., Miller and Rock (1985) and Stein (1989) for a fuller discussion of the forces that shape the trade-off between the two objectives. One potentially important factor has to do with agency considerations. Specifically, shareholders may—in response to agency problems—impose on managers an incentive scheme or corporate policies that have the effect of making the managers behave as if they were more concerned with maximizing current stock prices. This issue is discussed in more detail in Sec. VC below.

the manager seeks to maximize his perception of the present value of future cash flows. In this case, the “value” created by investment is $(P^* - 1)$. That is, the manager should only invest if the *rational expectations value* of the assets exceed their acquisition cost. Thus the long-horizon hurdle rate h^L has the property that $F^r/(1 + h^L) = P^*$. Using equation (8), this leads to

PROPOSITION 2. In the long-horizon case, the manager should discount his expected cash flow F^r at a hurdle rate $h^L = k^*$. In other words, the manager should take a FAR approach and choose a hurdle rate that reflects the fundamental risk of the assets in question and that is independent of outside investors’ bias δ .

Proposition 2 suggests that hurdle rates in the long-horizon case should be set in a “CAPM-like” fashion. This is very close in spirit to the standard textbook prescription. However, the one major caveat is that, unlike in the textbook world, one needs to be more careful in the empirical implementation. According to equations (6) and (7), the β^* that is needed for this CAPM-like calculation is the (unobserved) β that would prevail in a rational expectations world, as this is the correct measure of the fundamental risk borne by long-horizon investors. And given that the underlying premise throughout is that the stock market is inefficient, one cannot blithely make the usual assumption that a β calculated in the traditional way—with a regression of the firm’s *stock returns* on market returns—will provide an adequate proxy for the β^* that is called for in proposition 2. Thus there is a nontrivial set of issues surrounding the best way to measure β^* . These issues are taken up in detail in Section VI below.

III. Financing Considerations and Optimal Hurdle Rates

So far, the analysis has ignored the possibility that the firm might either issue or repurchase shares. Given the premise—that the market is inefficient and that managers know it—this is a potentially important omission. First of all, there will naturally be circumstances in which managers wish to engage in stock issues or repurchases to take advantage of market inefficiencies. Second, and more significant for our purposes, there may, in some cases, be a link between these opportunistic financing maneuvers and the optimal hurdle rate for capital budgeting.

The goal of this section is to explore these links between financing considerations and hurdle rates. I begin with a general formulation of the problem. I then consider a series of special cases that yield particularly crisp results and that highlight the most important intuition.

A. A General Formulation of the Problem

When a manager chooses an investment-financing combination in an inefficient market, there are, in general, three considerations that must

be taken into account: (1) the net present value of the investment, (2) the “market timing” gains or losses associated with any share issues or repurchases, and (3) the extent to which the investment-financing combination leads to any costly deviations from the optimal capital structure for the firm. Thus, in order to specify an overall objective function, one must spell out each of these considerations in detail.⁸

1. The Net Present Value of Investment

As will become clear, to the extent that financing considerations have any consequence at all for hurdle rates, it is to effectively shorten managers’ time horizons—that is, to make them behave in more of a NEER fashion. Therefore, to make the analysis interesting, I assume that absent financing concerns, managers take a FAR approach, and seek to maximize the present value of future cash flows.

For the purpose of doing a bit of calculus, I generalize slightly from the previous section, and allow the amount invested at time 0 to be a continuous variable K . The gross expected proceeds at time 1 from this investment are given by $f(K)$, which is an increasing, concave function. The relevant FAR-based definition of the net present value of investment is thus $f(K)P^*/F^r - K$, or equivalently, $f(K)/(1 + k^*) - K$, where k^* continues to be given by equations (6) and (7).

2. Market Timing Gains or Losses

Denote by E the dollar amount of equity raised by selling new shares at time 0. Thus, if $E > 0$, this should be interpreted as an equity issue by the firm; if $E < 0$, this should be interpreted as a repurchase. If the firm is able to transact in its own equity without any price-pressure effects, the market timing gains from the perspective of the manager are given simply by the difference between the market’s initial time-0 valuation of the shares and the manager’s time-0 valuation. For a transaction of size E , this market timing gain is simply $E(1 - P^*/P)$.⁹

Of course, it is extreme and unrealistic to assume that there are no price-pressure effects whatsoever, particularly if the implied equity transactions turn out to be large in absolute magnitude. At the same time, given the premise of investor irrationality, one does not necessarily want to go to the other extreme—represented by rational asymmetric information models such as that of Myers and Majluf (1984)—and assume that the announcement effects of a share issue or repurchase

8. One possibility that I ignore for the time being is that managers might wish to take advantage of market inefficiencies by transacting in the stock of *other* firms. This consideration is taken up in Sec. VD below, and, as will be seen, need not materially affect the conclusions of the analysis.

9. As above, I continue to assume that when the firm issues debt, this debt is fairly priced, so that there are no market timing gains or losses. This assumption can be relaxed without affecting the qualitative results that follow.

are such that they, on average, completely eliminate the potential for market timing gains.

As a compromise, I adopt a simple, relatively unstructured formulation in which the net-of-price-pressure market timing gains are given by $E(1 - P^*/P) - i(E)$. Here $i(E)$ captures the price-impact-related losses associated with an equity transaction of size E , with $i(0) = 0$. The only other restrictions I impose a priori are, first, when $E > 0$, $di/dE \geq 0$, and, conversely, when $E < 0$, $di/dE \leq 0$; second, $d^2i/dE^2 \geq 0$ everywhere. In words, equity issues tend to knock prices down, while repurchases push prices up, with larger effects for larger transactions in either direction.

The $i(E)$ function can be interpreted in terms of a couple of different underlying phenomena. First, it might be that even irrational investors do update their beliefs somewhat when they see management undertaking an equity transaction. However, in contrast to rational models based on asymmetric information, the updating is insufficient to wipe out predictable excess returns. This interpretation fits with the spirit of recent studies that suggest that the market underreacts dramatically to the information contained in both seasoned equity offerings and repurchases.¹⁰ Alternatively, in the case of share repurchases, $i(E)$ might be thought of as reflecting the premium that tendering investors require to compensate them for capital gains taxes.

3. The Costs of Deviating from Optimal Capital Structure

Finally, one must consider the possibility that a given investment-financing combination will lead to a suboptimal capital structure. For example, if a firm decides to invest a great deal and to engage in repurchases to take advantage of a low stock price, leverage may increase to the point where expected costs of financial distress become significant. To capture this possibility in a simple way, I assume that the optimal debt ratio for the firm is given by D , and that, prior to the investment and financing choices at time 0, the firm is exactly at this optimum. Thus, after it has invested an amount K and raised an amount of new equity E , it will be overleveraged by an amount $L \equiv K(1 - D) - E$. I assume that this imposes a cost of $Z(L)$.

Again, I do not put too much a priori structure on the $Z(L)$ function. By definition, things are normalized so that $Z(0) = 0$. In principle, straying in either direction from the optimum of 0 can be costly—too little debt may be a problem as well as too much debt. Moreover, to the extent that there are costs of straying, these costs are a convex function of the distance from the optimum. As with the $i(E)$ function,

10. See, e.g., Cheng (1995), Loughran and Ritter (1995), and Spiess and Affleck-Graves (1995) on seasoned equity offerings, and Ikenberry, Lakonishok, and Vermaelen (1995) on share repurchases.

this implies that $dZ/dL \geq 0$ when $L > 0$, and, conversely, $dZ/dL \leq 0$ when $L < 0$; also, $d^2Z/dL^2 \geq 0$ everywhere.

4. Optimal Investment and Financing Decisions

Taking all three considerations together, the manager's objective function is

$$\max f(K)P^*/F^r - K + E(1 - P^*/P) - i(E) - Z(L), \quad (11)$$

subject to

$$L \equiv K(1 - D) - E.$$

The first-order conditions for this problem are

$$df/dK - [1 + (1 - D)dZ/dL]F^r/P^* = 0, \quad (12)$$

$$(1 - P^*/P) - di/dE + dZ/dL = 0. \quad (13)$$

A little algebra shows that optimal investment therefore satisfies

$$\begin{aligned} df/dK &= DF^r/P^* + (1 - D)(F^r/P + di/dE) \\ &= D(1 + k^*) + (1 - D)(1 + \text{CER} + di/dE). \end{aligned} \quad (14)$$

B. Case-by-Case Analysis

Since the intuition underlying equation (14) may not be immediately apparent, it is useful to go through a series of special cases to build an understanding of the various forces at work.

1. Capital Structure Is Not a Binding Constraint

The first, simplest limiting case to consider is one in which $dZ/dL = 0$ —that is, there are no marginal costs or benefits to changing leverage, other than those that come directly from issuing or repurchasing shares at time 0. This condition would clearly hold in a world with no taxes and no costs of financial distress, where, were it not for the mispricing of the firm's stock, the Modigliani-Miller theorem would apply. But more generally, one need not make such strong assumptions for this case to be (approximately) relevant. All that is really required is that the Z function be flat in the neighborhood of the optimal solution. For example, if price-pressure effects are significant, and diminishing returns to investment set in quickly, the firm will only choose to make small investment and financing adjustments, and thus will never try to push capital structure very far from its initial position of $L = 0$. If, in addition, the Z function happens to be flat in this region near 0, the firm will be left in a position where, for example, incremental increases in leverage would have only a trivial impact on costs of financial distress.




FIG. 1.—Investment and financing policies when capital structure is not a constraint.

When $dZ/dL = 0$, equations (12) and (13) tell us that investment and financing decisions are fully separable. Intuitively, this is because capital structure can at the margin adjust costlessly to take up the slack between the two. The optimal behavior for the firm in this case is spelled out in the following proposition:

PROPOSITION 3. When capital structure is not a binding constraint, and the manager has long horizons, the optimal policies are always to set the hurdle rate at the FAR value of k^* , as in proposition 2, and to issue stock if the CER $< k^*$, but repurchase stock if the CER $> k^*$.

Figure 1 illustrates the optimal investment and financing policies. As can be seen, the two are completely decoupled. When the stock price is low and the CER is high, the firm repurchases shares. However, because capital structure is fully flexible, the repurchase does not affect its hurdle rate. Rather, the firm adjusts to the repurchase purely by taking on more debt. Therefore, at the margin, investment should be evaluated vis-à-vis fairly priced debt finance, exactly as in Section II above.

Conversely, when the stock price is high and the CER is low, the firm issues shares. However, it does not have to plow the proceeds of the share issue into investment. These proceeds can be used to pay down debt or accumulate cash. So there is no reason that the issuance of “cheap stock” should lower the hurdle rate for investment.

2. Binding Capital Structure Constraint, No Price-Pressure Effects

The next case to examine is one in which the capital structure constraint is binding, but where price-pressure effects are absent—that is, one in which $dZ/dL \neq 0$ and $di/dE = 0$. In this case, equation (14) simplifies to

$$df/dK = D(1 + k^*) + (1 - D)(1 + \text{CER}). \quad (15)$$

Based on equation (15), we have:

PROPOSITION 4. When the capital structure constraint is binding and there are no price-pressure considerations, the optimal hurdle rate has the following properties: the hurdle rate is between the NEER and FAR values of CER and k^* , respectively; as D approaches 0, the hurdle rate converges to CER, as in proposition 1; as D approaches 1, the hurdle rate converges to k^* , as in proposition 2.

The intuition behind proposition 4 is very simple. When capital structure imposes a binding constraint, one cannot, in general, separate investment and financing decisions. This is perhaps easiest to see in the case where $\delta < 0$, so that the stock is undervalued and the firm would like to repurchase shares. For each dollar that is devoted to investment, there is less cash available to engage in such repurchases, holding the capital structure fixed. Indeed, in the extreme case where $D = 0$ —that is, where the incremental investment has zero debt capacity—each dollar of investment leaves one full dollar less available for repurchases. Hence, in this case, the opportunity cost of investment is simply the expected return on the stock, as in the NEER approach.

Thus, in the limiting case where $D = 0$, financial constraints force managers who would otherwise take a long-run view into behaving exactly as if they were interested in maximizing short-term stock prices. This is simply because, in order to leave capital structure undisturbed, any investment must be fully funded by an immediate stock issue, so all that matters is the market's current assessment of whether the investment is attractive or not.

In the intermediate cases, where $0 < D < 1$, investment need only be partially funded by a stock issue. This implies that the hurdle rate moves less than one-for-one with the CER on the stock. At the other extreme, when $D = 1$ and investment can be entirely debt financed, the hurdle rate remains anchored at the FAR value of k^* , irrespective of the CER on the stock.¹¹ Figure 2 illustrates the relationship between the hurdle rate and the CER on the stock for different values of D .

3. Binding Capital Structure Constraint and Price-Pressure Effects

The final case to consider is the most general one where the capital structure constraint is binding and where there are price-pressure ef-

11. Of course, one should not take this limiting case too literally, given that I have also assumed that the firm can issue fairly priced (i.e., riskless) debt.




FIG. 2.—Optimal hurdle rates with binding capital structure constraint and no price-pressure effects.

fects. This case is most usefully attacked by breaking it down into subcases.

Stock is undervalued: $\delta < 0$. When $\delta < 0$, it is easy to show that $L > 0$. That is, the firm will choose to be overlevered relative to the static optimal capital structure of $L = 0$. However, the sign of E is ambiguous, the firm may either issue or repurchase shares. This ambiguity in E arises because there are two competing effects: on the one hand, the fact that $\delta < 0$ makes a repurchase attractive from a market timing standpoint; on the other, given that the firm is investing, it needs to raise some new equity if it does not wish to see its capital structure get too far out of line. Depending on which effect dominates, there can either be a *net share repurchase* or a *share issue*. In the case of share repurchase ($E < 0$), it is straightforward to verify:

PROPOSITION 5. When the capital structure constraint is binding, there are price-pressure considerations, $\delta < 0$ and $E < 0$, then the optimal hurdle rate has the following properties: the hurdle rate is always between the NEER and FAR values; the stronger the price-pressure effects—that is, the larger is di/dE in absolute magnitude—the lower the hurdle rate, all else equal, and therefore the closer the hurdle rate is to the FAR value of k^* .

Proposition 5 says that this case represents a well-behaved middle ground between the two more extreme cases covered in propositions 3 and 4. When price-pressure effects are strong, the outcome is closer to that in proposition 3, where capital structure is not a binding constraint—the hurdle rate is set more according to a FAR approach. This is because price pressure leads the firm to limit the scale of its repurchase activity. Consequently, capital structure is not much distorted, and there is less influence of financial constraints on investment. Of course, when price-pressure effects are very weak, we converge back to the case described in proposition 4. The outcome with an equity issue ($E > 0$) is a bit more counterintuitive:

PROPOSITION 6. When the capital structure constraint is binding, there are price-pressure considerations, $\delta < 0$ and $E > 0$, then the optimal hurdle rate has the following properties: The hurdle rate no longer necessarily lies between the NEER and FAR values; in particular, it may exceed them both, though it will never be below the lower of the two, namely the FAR value of k^* . The stronger are price-pressure effects—that is, the larger di/dE is in absolute magnitude—the higher the hurdle rate, all else equal.

Thus here is a situation—the first we have encountered so far—where the hurdle rate does not lie between the NEER and FAR values. However, this result has nothing really to do with the market irrationality that is the focus of this article. Rather, it is just a variant on the Myers-Majluf (1984) argument that when investment requires an equity issue, and such an equity issue knocks stock prices down, there will typically be underinvestment. Indeed, the effect is seen most cleanly by assuming that there is no irrationality whatsoever—that is, that $\delta = 0$ —so that the NEER and FAR values coincide. Inspection of equation (14) tells us that optimal investment will satisfy $df/dK = (1 + k^*) + (1 - D)di/dE$. In other words, the hurdle rate is a markup over the NEER/FAR value of k^* , with the degree of this markup determined by the magnitude of the price-pressure effect.

Stock is overvalued: $\delta > 0$. When $\delta > 0$, there is no ambiguity about the sign of E . This is because both market timing considerations and the need to finance investment now point in the same direction, implying a desire to sell stock. So $E > 0$. This gives rise to a result very similar to that seen just above:

PROPOSITION 7. When the capital structure constraint is binding, there are price-pressure considerations, $\delta > 0$ and $E > 0$, then the optimal hurdle rate has the following properties: The hurdle rate no longer necessarily lies between the NEER and FAR values. In particular, it may exceed them both, though it will never be below the lower of the two, namely, the NEER value of CER. The stronger the price-pressure effects—that is, the larger di/dE is in absolute magnitude—the higher the hurdle rate, all else equal.

C. Conclusions on the Effects of Financing Considerations

The analysis of this section has shown that financing considerations shape the optimal hurdle rate in three distinct ways. The first factor that matters is the shape of the Z function, which measures the degree to which deviations in capital structure are costly. When such deviations are inconsequential, this tends to favor a FAR-based approach to setting hurdle rates. In contrast, when such deviations are costly, the optimal hurdle rate is pushed in the NEER direction.

The second factor is the debt capacity, D , of the new investment. This second factor interacts with the first. In particular, the lower is D , the more pronounced an effect capital structure constraints have in terms of driving the hurdle rate toward the NEER value.

The third factor is the extent to which share issues or repurchases have price-pressure consequences. In terms of the NEER-FAR dichotomy, the impact of this third factor is somewhat more ambiguous than that of the other two. One can make a clear-cut statement only in the case where the firm engages in a stock repurchase; here price-pressure considerations unambiguously move the hurdle rate closer to the FAR value of k^* . However, when the firm issues equity, all that one can say for sure is that price pressure exerts an upward influence on the hurdle rate; it no longer follows that the hurdle rate is pushed closer to the FAR value.

The overall message of this section is that, while one can certainly argue in favor of a FAR-based approach to capital budgeting, the argument is somewhat more delicate than it might have appeared in Section II, and it does not apply in all circumstances. In order for a FAR-based approach to make sense, not only must managers have long horizons, but they must be relatively unconstrained by their current capital structures.

IV. Implementing a FAR-Based Approach: Measuring β^*

Part of the appeal of a FAR-based approach to capital budgeting is that it appears to be very close to the textbook CAPM method. However, as noted in Section II above, the one hitch is that, in order to implement a FAR-based approach, one needs to know β^* , which is the value of β that would prevail in a rational expectations world; that is, the fundamental risk of the assets in question. And given the underlying premise of the article—that the stock market is inefficient—one cannot simply assume that a β calculated using observed stock returns will yield a good estimate of β^* . Thus the following question arises: as a practical matter, how close to β^* can one expect to get using the standard regression methodology for calculating β ?

A. Theoretical Considerations

In order to clarify the issues, it is useful to begin with a more detailed analytical comparison of the value of a β computed from actual stock return data—which I will continue to denote by β^r —versus that of β^* . To do so, I will generalize somewhat from the setting of the previous sections by entertaining the possibility that there is mispricing of the market as a whole as well as mispricing of individual stocks. In addition, and somewhat trivially, I will allow for more than one period's worth of stock returns.

Note that in any period t , for any stock i , we can always make the following decomposition:

$$R_{it} \equiv R_{it}^* + N_{it}, \quad (16)$$

where R_{it} is the *observed* return on the stock, R_{it}^* is the return that would prevail in a rational expectations world—that is, the portion of the observed return due to “fundamentals”—and N_{it} is the portion of the observed return due to “noise.” We can also make a similar decomposition for the observed return on the market as a whole, R_{Mt} :

$$R_{Mt} \equiv R_{Mt}^* + N_{Mt}. \quad (17)$$

Clearly, as a general matter, the β calculated from observed stock returns, $\beta_i^r = \text{cov}(R_{it}, R_{Mt})/\text{var}(R_{Mt})$ will not coincide with $\beta_i^* = \text{cov}(R_{it}^*, R_{Mt}^*)/\text{var}(R_{Mt}^*)$. To get a better intuitive handle on the sources of the difference between β_i^r and β_i^* , it is helpful to consider a simple case where both R_{it}^* and N_{it} are generated by one-factor processes, as follows:

$$R_{it}^* = \beta_i^* R_{Mt}^* + \epsilon_{it}, \quad (18)$$

$$N_{it} = \theta_i N_{Mt} + \mu_{it}, \quad (19)$$

where $\text{cov}(\epsilon_{it}, \mu_{it}) = 0$.¹² In this formulation, θ_i represents the sensitivity of stock i 's noise component to the noise component on the market as a whole—that is, θ_i is a “noise β ” for stock i .

It is now easy to calculate β_i^r :

$$\beta_i^r = [\beta_i^* \text{var}(R_M^*) + \theta_i \text{var}(N_M) + (\beta_i^* + \theta_i) \text{cov}(R_M^*, N_M)] / [\text{var}(R_M^*) + \text{var}(N_M) + 2\text{cov}(R_M^*, N_M)]. \quad (20)$$

From (20), one can see how various parameters influence the relative magnitudes β_i^r and β_i^* . The most important conclusion for our purposes

12. Note that the 2-period model used in Secs. II and III above does not quite conform to this specification. This is because all mispricing is assumed to disappear after the first period, which in turn implies that there are not enough degrees of freedom to also assume that $\text{cov}(\epsilon_{it}, \mu_{it}) = 0$. However, the simple specification of eqq. (18) and (19) is merely an expositional device that allows one to illustrate the important effects more clearly.

is that it is not obvious a priori that one will be systematically larger or smaller than the other. Indeed, in some circumstances, they will be exactly equal. For example, if $\text{var}(N_M) = 0$, so that all noise is firm specific and washes out at the aggregate level, then $\beta_i^* = \beta_i^*$. Alternatively, the same result obtains if there is marketwide noise, but $\theta_i = \beta_i^*$. Although these cases are clearly special, they do illustrate a more general point: a stock may be subject to very large absolute pricing errors—in the sense of $\text{var}(N_i)$ being very large—and yet one might in principle be able to retrieve quite reasonable estimates of β_i^* from stock-price data.¹³ Whether this is true in practice is, then, a purely empirical question.

B. Existing Evidence

In order to ascertain whether a β estimated from stock-price data does in fact do a good job of capturing the sort of fundamental risk envisioned in β^* , one needs to develop an empirical analog of β^* . A natural, though somewhat crude, approach would be as follows. Suppose one posits that the rational expectations value of a stock is the present value of the expected cash flows to equity, discounted at a constant rate. Suppose further that cash flows follow a random walk, so today's level is a sufficient statistic for future expectations. In this very simple case, it is easy to show that for any given stock i :

$$\beta_i^* = \text{cov}(\Delta F_i/F_i, \Delta M/M)/\text{var}(\Delta M/M), \quad (21)$$

where F_i and M are the cash flows accruing to stock i and the market as a whole, respectively. This is a quantity that can be readily estimated and then compared to the corresponding β 's estimated from stock prices.

In fact, there is an older literature, beginning with Ball and Brown (1969) and Beaver, Kettler, and Scholes (1970), that undertakes a very similar comparison. In this literature, the basic hypothesis being tested is whether "accounting β 's" for either individual stocks or portfolios are correlated with β 's estimated from stock returns.¹⁴ In some of this work—notably Beaver and Manegold (1975)—accounting β 's are defined in a way that is very similar to equation (21), with the primary

13. A second implication of (20) is that if the marketwide noise is stationary, one might be able to obtain better estimates of β_i^* by using longer-horizon returns. At sufficiently long horizons, the variance of R_M^* will dominate the other terms in (20), ultimately leading β_i^* to converge to β_i^* .

14. The motivation behind this earlier literature is quite different from that here, however. In the 1970s work, market efficiency is taken for granted, and the question posed is whether accounting measures of risk are informative, in the sense of being related to market-based measures of risk (which are assumed to be objectively correct). In addition to the papers mentioned in the text, see also Gonedes (1973, 1975) for further examples of this line of research.

exception being that an accounting net income number is typically used in place of a cash flow.

Subject to this one accounting-related caveat, Beaver and Manegold's (1975) results would seem to indicate that there is indeed a fairly close correspondence between stock market β 's and fundamental risk. For example, with 10-stock portfolios, the Spearman correlation coefficient between accounting and stock-return β 's varies from about .70 to .90, depending on the exact specification used.

The bottom line is that both theoretical considerations and existing empirical evidence suggest that, at the least, it may not be totally unreasonable to assume simultaneously that stocks are subject to large pricing errors and that a β estimated from stock returns can provide a good measure of the fundamental asset risk variable β^* needed to implement a FAR-based approach to capital budgeting.¹⁵

V. Extensions and Variations

The analysis in Sections II and III above has made a number of strong simplifying assumptions. In some cases, it is easy to see how the basic framework could be extended so as to relax these assumptions; in other cases, it is clear that the problem becomes substantially more complex and that more work is required.

A. *Alternative Measures of Fundamental Risk*

One assumption that has been maintained until now is that the underlying structure of the economy is such that β^* is the appropriate summary statistic for an asset's fundamental risk. This need not be the case. One can redo the entire analysis in a world where there is a multifactor representation of fundamental risk, such as that which emerges from the APT or the ICAPM. In either case, the spirit of the conclusions would be unchanged—these alternative risk measures would be used instead of β^* to determine FAR-based hurdle rates. Whether or not such FAR-based hurdle rates would actually be used for capital budgeting purposes—as opposed to NEER-based hurdle rates—would continue to depend on the same factors identified above, namely managers' time horizons and financing constraints.

The harder question this raises is how can one know a priori which is the right model of fundamental risk. For once one entertains the premise that the market is inefficient, it may become difficult to use

15. Of course, this statement may be reasonable on average and at the same time be more appropriate for some categories of stocks than for others. To take just one example, some stocks—e.g., those included in the S&P 500—might have more of a tendency to covary excessively with a market index. This would tend to bias measured β 's for these particular stocks toward one and thereby present a misleading picture of their fundamental risk. More empirical work would clearly be useful here.

empirical data in a straightforward fashion to choose between, say, a β^* representation of fundamental risk and a multifactor APT-type representation. Clearly, one cannot simply run atheoretical horse races and see which factors better predict expected returns. For such horse races may tell us more about the nature of market inefficiencies than about the structure of the underlying fundamental risk. In particular, a book-to-market “factor”—à la Fama and French (1993)—may do well in prediction equations, but given the lack of a theoretical model, it would seem inappropriate to unquestioningly interpret this factor as a measure of fundamental risk. (Unless, of course, one’s priors are absolute that the market is efficient, in which case the distinction between NEER and FAR vanishes, and everything in this article becomes irrelevant.)

B. Managers Are Not Sure They Are Smarter than the Market

Thus far, the discussion has proceeded as if a manager’s estimate of future cash flow is always strictly superior to that of outside shareholders. However, a less restrictive interpretation is also possible. One might suppose that outside shareholders’ forecast of F_t , while containing some noise, also embodies some information not directly available to managers. In this case, the optimal thing for a rational manager to do would be to put some weight on his own private information and some weight on outside shareholders’ forecast. That is, the manager’s rational forecast, F^r , would be the appropriate Bayesian combination of the manager’s private information and the market forecast. The analysis would then go forward exactly as before. So one does not need to interpret FAR-based capital budgeting as dictating that managers completely ignore market signals in favor of their own beliefs; rather it simply implies that they will be less responsive to such market signals than with NEER-based capital budgeting.

C. Agency Considerations

Suppose we have a firm that is financially unconstrained and whose shareholders all plan to hold onto their shares indefinitely. The analysis above might seem to suggest that such a firm should adopt a FAR-based approach to capital budgeting. But this conclusion rests in part on the implicit assumption that the manager who makes the cash flow forecasts and carries out the capital budgeting decisions acts in the interests of shareholders. More realistically, there may be agency problems, and managers may have a desire to overinvest relative to what would be optimal. If this is the case, and if the manager’s forecast F^r is not verifiable, shareholders may want to adopt ex ante capital budgeting policies that constrain investment in some fashion.

One possibility—though not necessarily the optimal one—is for shareholders simply to impose on managers NEER-based capital bud-

geting rules. An advantage of NEER-based capital budgeting in an agency context is that it brings to bear some information about F that is verifiable. Specifically, under the assumptions of the model above, shareholders can always observe whether or not NEER-based capital budgeting is being adhered to simply by looking at market prices. In contrast, if the manager is left with the discretion to pursue FAR-based capital budgeting, there is always the worry that he will overinvest and explain it away as a case where the privately observed F' is very high relative to the forecast implicit in market prices. Of course, if shareholders have long horizons, there is also a countervailing cost to imposing NEER-based capital budgeting to the extent that market forecasts contain not only some valid information about F , but biases as well.

This discussion highlights the following limitation of the formal analysis: while I have been treating managers' horizons as exogenous, they would, in a more complete model, be endogeneously determined. Moreover, in such a setting, managerial horizons might not correspond to those of the shareholders for whom they are working. If agency considerations are important, shareholders may choose ex ante to set up corporate policies or incentive schemes that effectively foreshorten managerial horizons, even when this distorts investment decisions.¹⁶

D. Portfolio Trading in the Stock of Other Firms

To this point, I have ignored the possibility that managers might wish to take advantage of market inefficiencies by transacting in stocks other than their own. To see why this possibility might be relevant for capital budgeting, consider a manager who perceives that his firm's stock is underpriced, say, because it has a high book-to-market ratio. On the one hand, as discussed above, this might lead the manager to engage in repurchases of his own stock. And to the extent that such repurchases push capital structure away from its optimum level of $L = 0$, they will spill over and affect investment decisions—in this particular case, raising the hurdle rate from its FAR value in the direction of the higher NEER value.

On the other hand, the capital structure complications associated with own-stock repurchases lead one to ask whether there are other ways for the manager to make essentially the same speculative bet. For example, he might create a zero net-investment portfolio, consisting of

16. This is already a very familiar theme in the corporate finance literature, particularly that on takeovers. For example, it has been argued that it can be in the interest of shareholders to remove impediments to takeovers as a way of improving managerial incentives, even when the resulting foreshortening of managerial horizons leads to distorted investment. See, e.g., Laffont and Tirole (1988), Scharfstein (1988) and Stein (1988). Note, however, that these earlier papers make the point without invoking market irrationality, but rather simply appeal to asymmetries of information between managers and outside shareholders.

long positions in other high book-to-market stocks and short positions in low book-to-market stocks. An apparent advantage of this approach is that it does not alter his own firm's capital structure.

For the purposes of this article, the bottom line question is whether the existence of such portfolio trading opportunities changes the basic conclusions offered in Section III above. The answer is, it depends. In particular, the pivotal issue is whether the other trading opportunities are sufficiently attractive and available that they completely eliminate managers' desire to distort capital structure away from the first-best of $L = 0$. If so, capital structure constraints will become irrelevant for hurdle rates, leading to strictly FAR-based capital budgeting. If not, the qualitative conclusions offered in Section III will continue to hold, with binding capital structure constraints pushing hurdle rates in the direction of NEER values.

Ultimately, the outcome depends on a number of factors that are not explicitly modeled above. First, while "smart money" managers can presumably exploit some simple inefficiencies—like the book-to-market effect—by trading in other stocks using only easily available public data, it seems plausible that they can do even better by trading in their own stock. If this is the case, there will be circumstances in which the existence of other trading opportunities does not eliminate the desire to transact in own-company stock, and the basic story sketched in Section III will still apply. A second unmodeled factor that is likely to be important is the extent to which firms exhibit risk aversion with respect to passive portfolio positions. If such risk aversion is pronounced, it will again be the case that the existence of other trading opportunities is not a perfect substitute for transactions in own-company stock.¹⁷

E. Richer Models of Irrationality

Finally, and perhaps most fundamentally, another area that could use further development is the specification of investors' misperceptions about key parameters. I have adopted the simplest possible approach here, assuming that all investors are homogeneous and that their only misperception has to do with the expected value of future firm cash flows. In reality, there are likely to be important heterogeneities across outside investors. Moreover, estimates of other parameters—such as variances and covariances—may also be subject to systematic biases. It would be interesting to see how robust the qualitative conclusions of this article are to these and related extensions.

17. One can imagine a number of reasons for such risk aversion at the corporate level. For example, Froot, Scharfstein, and Stein (1993) develop a model in which capital market imperfections lead firms to behave in a risk-averse fashion, particularly with respects to those risks—such as portfolio trading—that are uncorrelated with their physical investment opportunities.

VI. Empirical Implications

Traditional efficient-markets-based models conclude that a firm's investment behavior ought to be closely linked to its stock price. And, indeed, a substantial body of empirical research provides evidence for such a link.¹⁸ At the same time, however, a couple of recent papers have found that once one controls for fundamentals like profits and sales, the incremental explanatory power of stock prices for corporate investment, while statistically significant, is quite limited in economic terms, both in firm-level and aggregate data (Morck et al. 1990; Blanchard et al. 1993). Thus it appears that, relative to these fundamental variables, the stock market may be something of a sideshow in terms of its influence on corporate investment.

This sideshow phenomenon is easy to rationalize in the context of the model presented above. If the market is inefficient, and if managers are for the most part engaging in FAR-based capital budgeting, one would not expect investment to track stock prices nearly as closely as in a classical world. Perhaps more interestingly, however, this paper's logic allows one to go further in terms of empirical implications. Rather than simply saying the theory is consistent with existing evidence, it is also possible to generate some novel cross-sectional predictions.

These cross-sectional predictions flow from the observation that not all firms should have the same propensity to adopt FAR-based capital budgeting practices. In particular, FAR-based capital budgeting should be more prevalent among either firms with very strong balance sheets (who, in terms of the language of the model are presumably operating in a relatively flat region of the *Z* function) or those whose assets offer substantial debt capacity. In contrast, firms with weak balance sheets and hard-to-collateralize assets—for example, a cash-strapped software development company—should tend to follow NEER-based capital budgeting. Thus the testable prediction is that the cash-strapped software company should have investment that responds more sensitively to movements in its stock price than, say, an AAA-related utility with lots of tangible assets.

A similar sort of reasoning can be used to generate predictions for the patterns of asset sales within and across industries. For concreteness, consider two airlines, one financially constrained, the other not. Now suppose that a negative wave of investor sentiment knocks airline-industry stock prices down and thereby drives conditional expected returns up. The constrained airline, which uses NEER-based capital budgeting, will raise its hurdle rates, while the unconstrained airline, which uses FAR-based capital budgeting, will not. This divergence in the way the two airlines value physical assets might be ex-

18. See, i.e., Barro (1990) for an overview and a recent empirical treatment of the relationship between stock prices and investment.

pected to lead the constrained airline to sell some of its planes to the unconstrained airline. Conversely, if there is a positive sentiment shock, the prediction goes the other way—the constrained airline will cut its hurdle rates, and become a net buyer of assets.¹⁹

VII. Conclusions

Is β dead? The answer to this question would seem to depend on the job that one has in mind for β . If the job is to predict cross-sectional differences in stock returns, then β may well be dead, as Fama and French (1992) argue. But if the job is to help in determining hurdle rates for capital budgeting purposes, then β may be only slightly hobbled. Certainly, any argument in favor of using β as a capital budgeting tool must be carefully qualified, unlike in the typical textbook treatment. Nonetheless, in the right circumstances, the textbook CAPM approach to setting hurdle rates may ultimately be justifiable.

This defense of β as a capital budgeting tool rests on three key premises. First, one must be willing to assume that the cross-sectional patterns in stock returns that have been documented in recent research—such as the tendency of high book-to-market stocks to earn higher returns—reflect pricing errors, rather than compensation for fundamental sources of risk. Second, the firm in question must have long horizons and be relatively unconstrained by its current capital structure. And finally, it must be the case that even though there are pricing errors, a β estimated from stock returns is a satisfactory proxy for the fundamental riskiness of the firm's cash flows.

This article was intended as a first cut at the problem of capital budgeting in an inefficient market, and, as such, it leaves many important questions unanswered. There are at least three broad areas where further research might be useful. First, there are the pragmatic risk-measurement issues raised in Section IV, namely, just how well do stock return β 's actually reflect the fundamental riskiness of underlying firm cash flows? Are stock return β 's more informative about fundamental risk for some classes of companies than for others? Does lengthening the horizon over which returns are computed help matters? Here it would clearly be desirable to update and build on some of the work done in the 1970s.

Second, as discussed in Section V, there is potentially quite a bit more that can be done in terms of refining and extending the basic

19. I use the example of airlines because of a very interesting recent paper by Pulvino (1995), who documents exactly this sort of pattern of asset sales in the airline industry—financially unconstrained airlines significantly increase their purchases of used aircraft when prices are depressed. As Shleifer and Vishny (1992) demonstrate, this pattern can arise purely as a consequence of liquidity constraints and thus need not reflect any stock market inefficiencies. Nonetheless, in terms of generating economically large effects, such inefficiencies are likely to give an added kick to their story.

conceptual framework. And finally, as seen in Section VI, the theory developed here gives rise to some new empirical implications, having to do with cross-sectional differences in the intensity of the relationship between stock prices and corporate investment.

References

- Ball, R., and Brown, P. 1969. Portfolio theory and accounting. *Journal of Accounting Research* 7 (Spring): 300–323.
- Banz, Rolf W. 1981. The relationship between return and market value of common stocks. *Journal of Financial Economics* 9 (March): 3–18.
- Barro, Robert J. 1990. The stock market and investment. *Review of Financial Studies* 3 (March): 115–31.
- Basu, Sanjoy. 1977. Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis. *Journal of Finance* 32 (June): 663–82.
- Beaver, William; Kettler, Paul; and Scholes, Myron. 1970. The association between market determined and accounting determined risk measures. *Accounting Review* (October): 654–82.
- Beaver, William, and Manegold, James. 1975. The association between market-determined and accounting-determined measures of systematic risk: Some further evidence. *Journal of Financial and Quantitative Analysis* (June): 231–84.
- Bhandari, Laxmi Chand. 1988. Debt/equity ratio and expected common stock returns: Empirical evidence. *Journal of Finance* 43 (June): 507–28.
- Blanchard, Olivier; Rhee, Changyong; and Summers, Lawrence. 1993. The stock market, profit, and investment. *Quarterly Journal of Economics* 108 (February): 115–36.
- Bosworth, Barry. 1975. The stock market and the economy. *Brookings Papers on Economic Activity*, pp. 257–300.
- Chan, Louis K. C.; Hamao, Y.; and Lakonishok, Josef. 1991. Fundamentals and stock returns in Japan. *Journal of Finance* 46 (December): 1739–64.
- Chan, Louis K. C.; Jegadeesh, Narasimhan; and Lakonishok, Josef. 1995. Evaluating the performance of value versus glamour stocks: The impact of selection bias. *Journal of Financial Economics* 38 (July): 269–96.
- Cheng, Li-Lan. 1995. The motives, timing and subsequent performance of seasoned equity issues. Ph.D. thesis, Massachusetts Institute of Technology, Department of Economics.
- Daniel, Kent, and Titman, Sheridan. 1995. Evidence on the characteristics of cross-sectional variation in stock returns. Working paper. Chicago: University of Chicago.
- Davis, James L. 1994. The cross-section of realized stock returns: The pre-COMPUSTAT evidence. *Journal of Finance* 49 (December): 1579–93.
- DeBondt, Werner F. M., and Thaler, Richard H. 1985. Does the stock market overreact. *Journal of Finance* 40 (July): 793–805.
- DeLong, J. Bradford; Shleifer, Andrei; Summers, Lawrence H.; and Waldmann, Robert J. 1989. The size and incidence of the losses from noise trading. *Journal of Finance* 44 (July): 681–96.
- Fama, Eugene F. 1991. Efficient capital markets: II. *Journal of Finance* 46 (December): 1575–1617.
- Fama, Eugene F., and French, Kenneth R. 1992. The cross-section of expected stock returns. *Journal of Finance* 47 (June): 427–65.
- Fama, Eugene F., and French, Kenneth R. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33 (February): 3–56.
- Fischer, Stanley, and Merton, Robert. 1984. Macroeconomics and finance: The role of the stock market. *Carnegie Rochester Conference Series on Public Policy* 21:57–108.
- Froot, Kenneth A.; Scharfstein, David S.; and Stein, Jeremy C. 1993. Risk management: Coordinating corporate investment and financing policies. *Journal of Finance* 48 (December): 1629–58.

- Gonedes, Nicholas J. 1973. Evidence on the information content of accounting numbers: Accounting-based and market-based estimates of systematic risk. *Journal of Financial and Quantitative Analysis* (June): 407–43.
- Gonedes, Nicholas J. 1975. A note on accounting-based and market-based estimates of systematic risk. *Journal of Financial and Quantitative Analysis* (June): 355–65.
- Ikenberry, David; Lakonishok, Josef; and Vermaelen, Theo. 1995. Market underreaction to open market share repurchases. Working paper. N.p.
- Jaffe, Jeffrey; Keim, Donald B.; and Westerfield, Randolph. 1989. Earnings yields, market values, and stocks returns. *Journal of Finance* 44 (March): 135–48.
- Keim, Donald B. 1983. Size-related anomalies and stock return seasonality. *Journal of Financial Economics* 12 (June): 13–32.
- Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. London: Macmillan.
- Laffont, Jean Jacques, and Tirole, Jean. 1988. Repeated auctions of incentive contracts, investment, and bidding parity with an application to takeovers. *Rand Journal of Economics* 19 (Winter): 516–37.
- Lakonishok, Josef; Shleifer, Andrei; and Vishny, Robert W. 1994. Contrarian investment, extrapolation, and risk. *Journal of Finance* 49 (December): 1541–78.
- La Porta, Rafael; Lakonishok, Josef; Shleifer, Andrei; and Vishny, Robert. 1994. Good news for value stocks: Further evidence on market efficiency. Working paper. Cambridge, Mass.: Harvard University.
- Loughran, Tim, and Ritter, Jay R. 1995. The new issues puzzle. *Journal of Finance* 50 (March): 23–51.
- McKinlay, A. Craig. 1995. Multifactor models do not explain deviations from the CAPM. *Journal of Financial Economics* 38 (May): 3–28.
- Miller, Merton M., and Rock, Kevin. 1985. Dividend policy under asymmetric information. *Journal of Finance* 40 (September): 1021–52.
- Morck, Randall; Shleifer, Andrei; and Vishny, Robert. 1990. The stock market and investment: Is the market a sideshow? *Brookings Papers on Economic Activity*, pp. 157–215.
- Myers, Stewart, and Majluf, N. 1984. Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics* 13 (June): 187–221.
- Pulvino, Todd. 1995. Do asset fire-sales exist? An empirical investigation of commercial aircraft transactions. Working paper. Cambridge, Mass.: Harvard University.
- Rosenberg, Barr; Reid, Kenneth; and Lanstein, Ronald. 1985. Persuasive evidence of market inefficiency. *Journal of Portfolio Management* 11:9–17.
- Scharfstein, David. 1988. The disciplinary role of takeovers. *Review of Economic Studies* 55 (April): 185–99.
- Shleifer, Andrei, and Vishny, Robert. 1992. Liquidation values and debt capacity: A market equilibrium approach. *Journal of Finance* 47 (September): 1343–66.
- Spiess, D. Katherine, and Affleck-Graves, John. 1995. Underperformance in long-run stock returns following seasoned equity offerings. *Journal of Financial Economics* 38 (July): 241–67.
- Stattman, Dennis. 1980. Book values and stock returns. *Chicago MBA: A Journal of Selected Papers* 4:25–45.
- Stein, Jeremy C. 1988. Takeover threats and managerial myopia. *Journal of Political Economy* 96 (February): 61–80.
- Stein, Jeremy C. 1989. Efficient capital markets, inefficient firms: A model of myopic corporate behavior. *Quarterly Journal of Economics* 104 (November): 655–69.