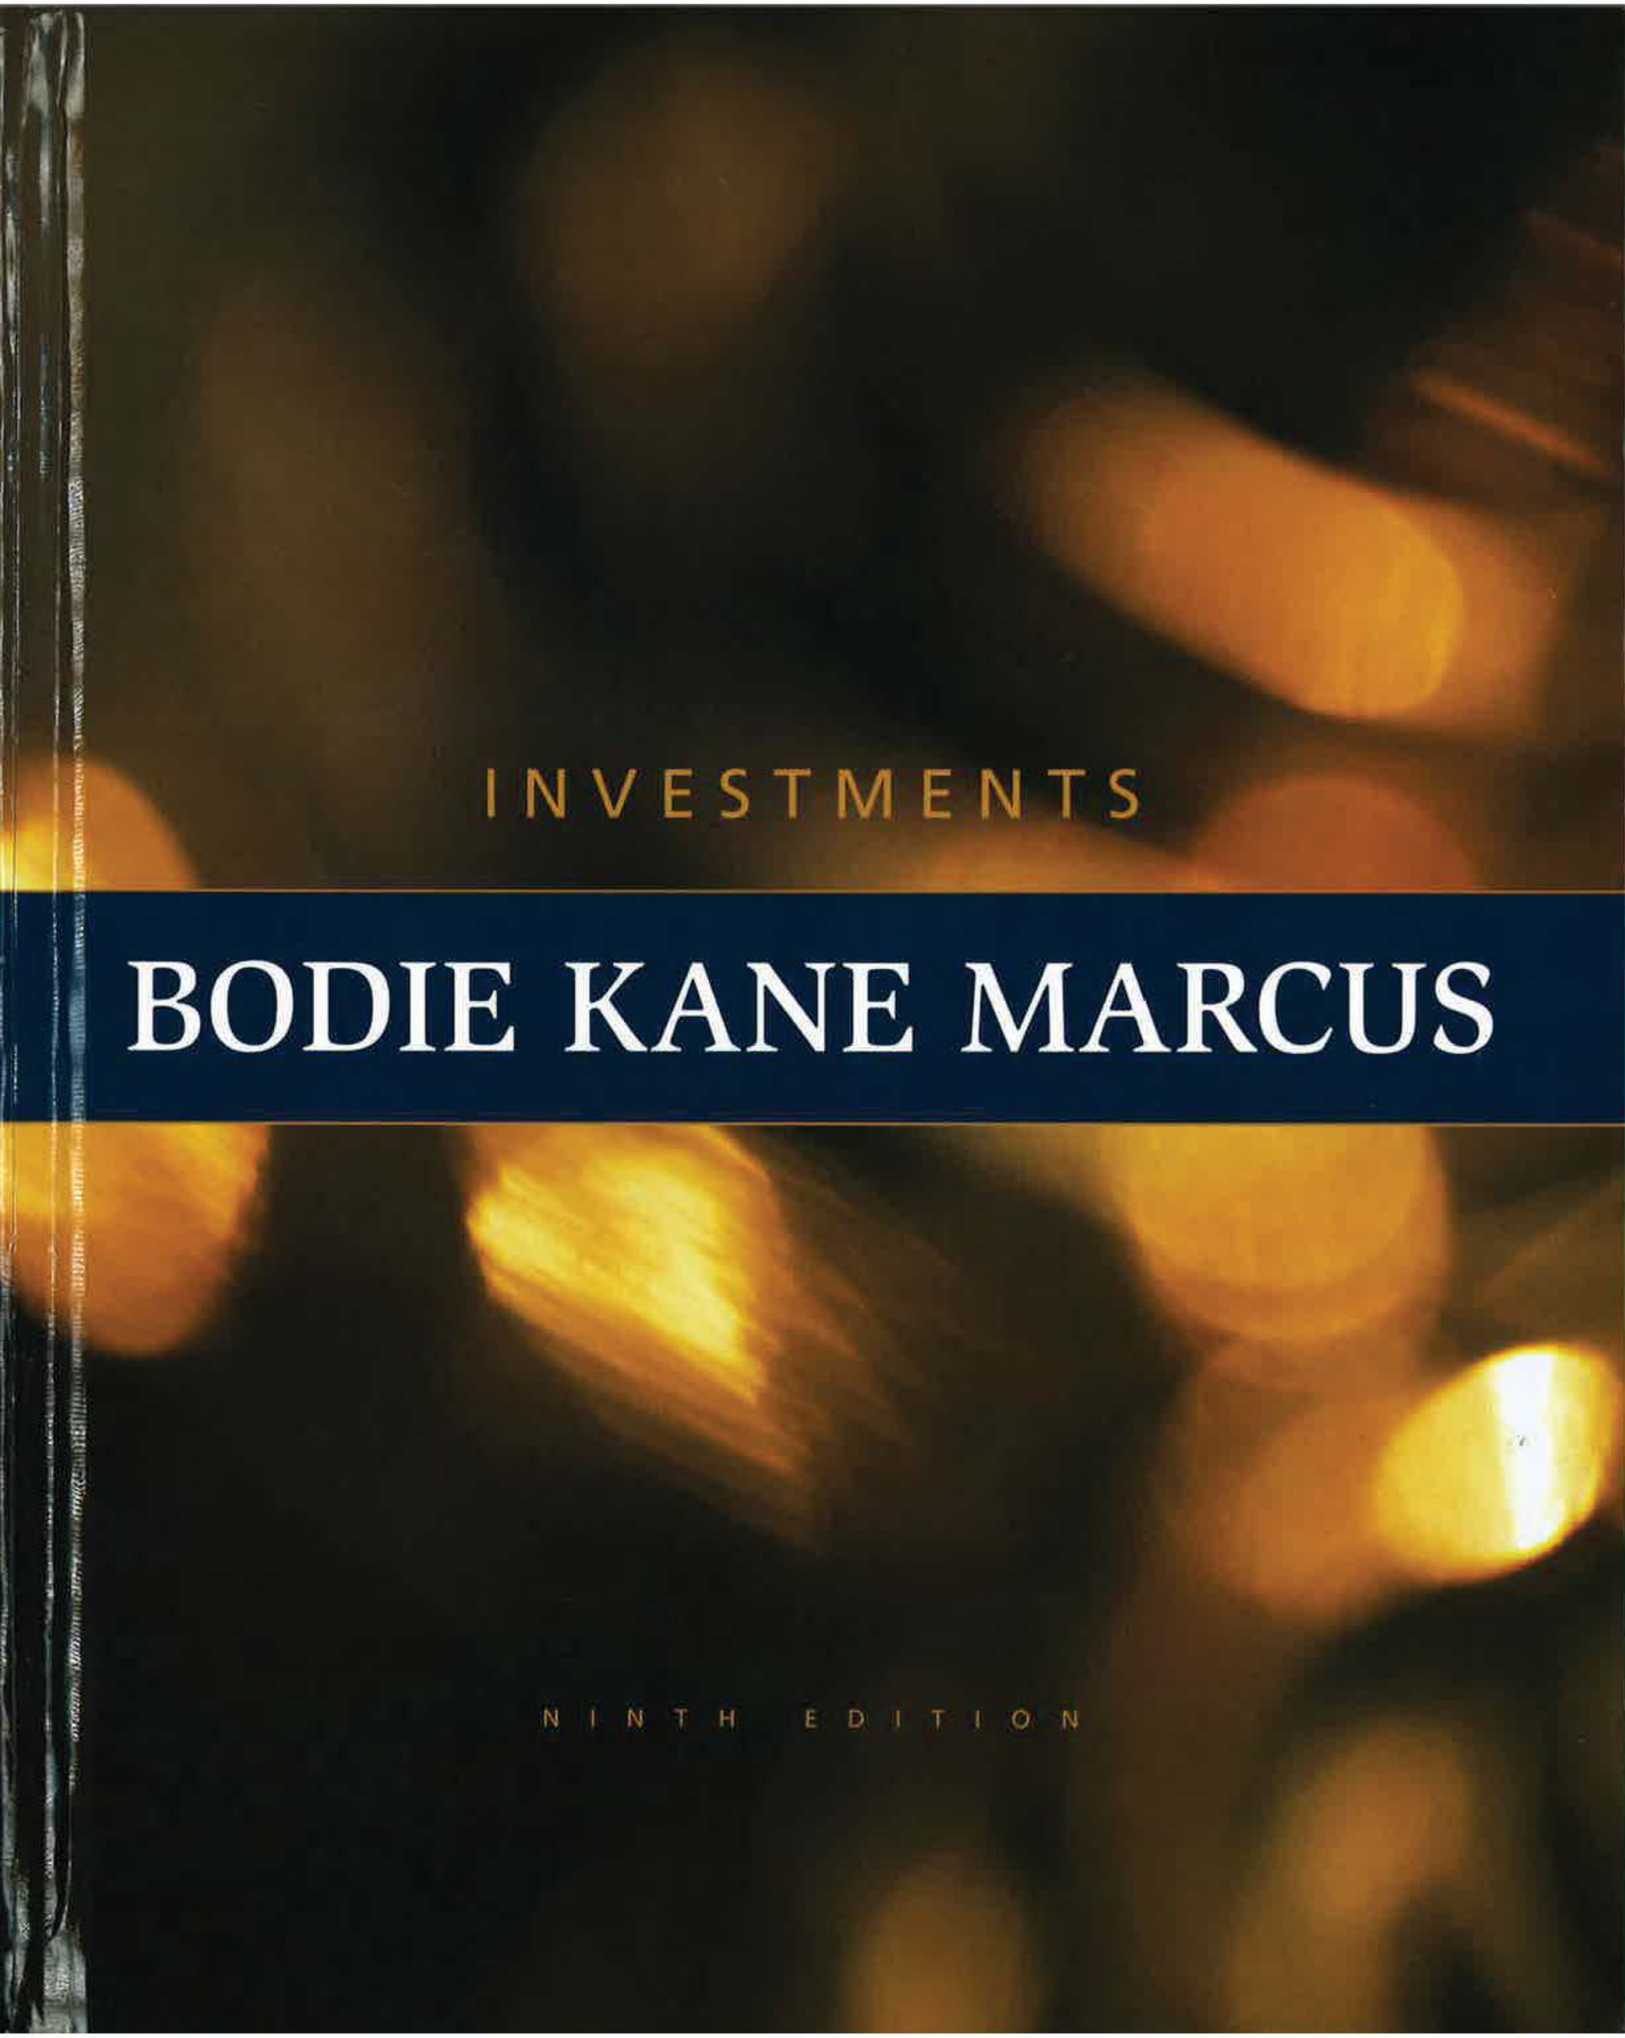

INVESTMENTS

BODIE KANE MARCUS

NINTH EDITION

shouldn't confuse an efficient market, where all available information is reflected in prices, with a perfect foresight market. Even "all available information" is still far from complete information, and generally rational market forecasts will sometimes be wrong; sometimes, in fact, they will be very wrong.

11.3 Event Studies

The notion of informationally efficient markets leads to a powerful research methodology. If security prices reflect all currently available information, then price changes must reflect new information. Therefore, it seems that one should be able to measure the importance of an event of interest by examining price changes during the period in which the event occurs.

An **event study** describes a technique of empirical financial research that enables an observer to assess the impact of a particular event on a firm's stock price. A stock market analyst might want to study the impact of dividend changes on stock prices, for example. An event study would quantify the relationship between dividend changes and stock returns.

Analyzing the impact of any particular event is more difficult than it might at first appear. On any day, stock prices respond to a wide range of economic news such as updated forecasts for GDP, inflation rates, interest rates, or corporate profitability. Isolating the part of a stock price movement that is attributable to a specific event is not a trivial exercise.

The general approach starts with a proxy for what the stock's return would have been in the absence of the event. The **abnormal return** due to the event is estimated as the difference between the stock's actual return and this benchmark. Several methodologies for estimating the benchmark return are used in practice. For example, a very simple approach measures the stock's abnormal return as its return minus that of a broad market index. An obvious refinement is to compare the stock's return to those of other stocks matched according to criteria such as firm size, beta, recent performance, or ratio of price to book value per share. Another approach estimates normal returns using an asset pricing model such as the CAPM or one of its multifactor generalizations such as the Fama-French three-factor model.

Many researchers have used a "market model" to estimate abnormal returns. This approach is based on the index models we introduced in Chapter 9. Recall that a single-index model holds that stock returns are determined by a market factor and a firm-specific factor. The stock return, r_t, during a given period t, would be expressed mathematically as

$$r_t = a + br_{Mt} + e_t \tag{11.1}$$

where r_{Mt} is the market's rate of return during the period and e_t is the part of a security's return resulting from firm-specific events. The parameter b measures sensitivity to the market return, and a is the average rate of return the stock would realize in a period with a zero market return.[7] Equation 11.1 therefore provides a decomposition of r_t into market and firm-specific factors. The firm-specific or abnormal return may be interpreted as the unexpected return that results from the event.

Determination of the abnormal return in a given period requires an estimate of e_t. Therefore, we rewrite Equation 11.1:

$$e_t = r_t - (a + br_{Mt}) \tag{11.2}$$

[7] We know from Chapter 9 that the CAPM implies that the intercept a in Equation 11.1 should equal $r_f(1 - \beta)$. Nevertheless, it is customary to estimate the intercept in this equation empirically rather than imposing the CAPM value. One justification for this practice is that empirically fitted security market lines seem flatter than predicted by the CAPM (see Chapter 13), which would make the intercept implied by the CAPM too small.

Equation 11.2 has a simple interpretation: The residual, e_t, that is, the component presumably due to the event in question, is the stock's return over and above what one would predict based on broad market movements in that period, given the stock's sensitivity to the market.

The market model is a highly flexible tool, because it can be generalized to include richer models of benchmark returns, for example, by including industry as well as broad market returns on the right-hand side of Equation 11.1, or returns on indexes constructed to match characteristic such as firm size. However, one must be careful that regression parameters in Equation 11.1 (the intercept a and slope b) are estimated properly. In particular, they must be estimated using data sufficiently separated in time from the event in question that they are not affected by event-period abnormal stock performance. In part because of this vulnerability of the market model, returns on characteristic-matched portfolios have become more widely used benchmarks in recent years.

Example 11.3 Abnormal Returns

Suppose that the analyst has estimated that $a = .05\%$ and $b = .8$. On a day that the market goes up by 1%, you would predict from Equation 11.1 that the stock should rise by an expected value of $.05\% + .8 \times 1\% = .85\%$. If the stock actually rises by 2%, the analyst would infer that firm-specific news that day caused an additional stock return of $2\% - .85\% = 1.15\%$. This is the abnormal return for the day.

We measure the impact of an event by estimating the abnormal return on a stock (or group of stocks) at the moment the information about the event becomes known to the market. For example, in a study of the impact of merger attempts on the stock prices of target firms, the announcement date is the date on which the public is informed that a merger is to be attempted. The abnormal returns of each firm surrounding the announcement date are computed, and the statistical significance and magnitude of the typical abnormal return are assessed to determine the impact of the newly released information.

One concern that complicates event studies arises from *leakage* of information. Leakage occurs when information regarding a relevant event is released to a small group of investors before official public release. In this case the stock price might start to increase (in the case of a "good news" announcement) days or weeks before the official announcement date. Any abnormal return on the announcement date is then a poor indicator of the total impact of the information release. A better indicator would be the **cumulative abnormal return,** which is simply the sum of all abnormal returns over the time period of interest. The cumulative abnormal return thus captures the total firm-specific stock movement for an entire period when the market might be responding to new information.

Figure 11.1 (earlier in the chapter) presents the results from a fairly typical event study. The authors of this study were interested in leakage of information before merger announcements and constructed a sample of 194 firms that were targets of takeover attempts. In most takeovers, stockholders of the acquired firms sell their shares to the acquirer at substantial premiums over market value. Announcement of a takeover attempt is good news for shareholders of the target firm and therefore should cause stock prices to jump.

Figure 11.1 confirms the good-news nature of the announcements. On the announcement day, called day 0, the average cumulative abnormal return (CAR) for the sample of takeover candidates increases substantially, indicating a large and positive abnormal return on the announcement date. Notice that immediately after the announcement date the CAR no longer increases or decreases significantly. This is in accord with the efficient market hypothesis. Once the new information became public, the stock prices jumped almost immediately in response to the good news. With prices once again fairly set, reflecting the effect of the new information, further abnormal returns on any particular day are equally likely to be positive or negative. In fact, for a sample of many firms, the average abnormal return should be extremely close to zero, and thus the CAR will show neither upward nor downward drift. This is precisely the pattern shown in Figure 11.1.

The pattern of returns for the days preceding the public announcement date yields some interesting evidence about efficient markets and information leakage. If insider trading rules were perfectly obeyed and perfectly enforced, stock prices should show no abnormal returns on days before the public release of relevant news, because no special firm-specific information would be available to the market before public announcement. Instead, we should observe a clean jump in the stock price only on the announcement day. In fact, Figure 11.1 shows that the prices of the takeover targets clearly start an upward drift 30 days before the public announcement. It appears that information is leaking to some market participants who then purchase the stocks before the public announcement. Such evidence of leakage appears almost universally in event studies, suggesting at least some abuse of insider trading rules.

Actually, the SEC also can take some comfort from patterns such as that in Figure 11.1. If insider trading rules were widely and flagrantly violated, we would expect to see abnormal returns earlier than they appear in these results. For example, in the case of mergers, the CAR would turn positive as soon as acquiring firms decided on their takeover targets, because insiders would start trading immediately. By the time of the public announcement, the insiders would have bid up the stock prices of target firms to levels reflecting the merger attempt, and the abnormal returns on the actual public announcement date would be close to zero. The dramatic increase in the CAR that we see on the announcement date indicates that a good deal of these announcements are indeed news to the market and that stock prices did not already reflect complete knowledge about the takeovers. It would appear, therefore, that SEC enforcement does have a substantial effect on restricting insider trading, even if some amount of it still persists.

Event study methodology has become a widely accepted tool to measure the economic impact of a wide range of events. For example, the SEC regularly uses event studies to measure illicit gains captured by traders who may have violated insider trading or other securities laws.[8] Event studies are also used in fraud cases, where the courts must assess damages caused by a fraudulent activity.

Example 11.4 Using Abnormal Returns to Infer Damages

Suppose the stock of a company with market value of $100 million falls by 4% on the day that news of an accounting scandal surfaces. The rest of the market, however, generally did well that day. The market indexes were up sharply, and on the basis of the usual relationship between the stock and the market, one would have expected a 2% gain on the stock. We would conclude that the impact of the scandal was a 6% drop in value, the difference between the 2% gain that we would have expected and the 4% drop actually observed. One might then infer that the damages sustained from the scandal were $6 million, because the value of the firm (after adjusting for general market movements) fell by 6% of $100 million when investors became aware of the news and reassessed the value of the stock.

[8]For a review of SEC applications of this technique, see Mark Mitchell and Jeffry Netter, "The Role of Financial Economics in Securities Fraud Cases: Applications at the Securities and Exchange Commission," *The Business Lawyer* 49 (February 1994), pp. 545–90.