ROUTLEDGE REVIVALS

Early Warning Indicators of Corporate Failure

A Critical Review of Previous Research and Further Empirical Evidence

Richard Morris

EARLY WARNING INDICATORS OF CORPORATE FAILURE

Early Warning Indicators of Corporate Failure

A critical review of previous research and further empirical evidence

RICHARD MORRIS
Professor of Accounting
University of Liverpool

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

First published 1997 by Ashgate Publishing

Reissued 2018 by Routledge
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN
711 Third Avenue, New York, NY 10017, USA

Routledge is an imprint of the Taylor & Francis Group, an informa business

A Library of Congress record exists under LC control number : 97073608

ISBN 13: 978-1-138-31200-5 (hbk)
ISBN 13: 978-0-429-45850-7 (ebk)

Contents

Acknowledgements *x*
Preface *xi*

Executive summary 1

1 The background 20

Introduction 20
The meaning of 'failure' 24
The meaning of 'prediction' 24
Methodological issues 26
Informational market efficiency 30
User needs 31
Assessing whether or not a business is a going concern 34
Summary 41
Notes 43
Appendix 1.1: Sampling bias error 48

PART I: Previous research 51

2 Normative theories of corporate failure 53

Introduction 53
Disequilibrium models of corporate failure 54
Financial models of corporate failure 64
Valuation theories 69
Agency models of corporate failure 79
Management theories of corporate failure 82
Summary 82
Notes 84
Appendix 2.1: Informational decomposition measures 85
Appendix 2.2: The gambler's ruin model 88

3 **Positive theories of corporate failure: I – Univariate models** 90

Introduction 90
The interpretation of financial statements 91
Traditional ratio analysis 95
Systematic ratio analysis 108
Summary 110
Notes 110

4 **Positive theories of corporate failure: II – Multivariate models** 113

Introduction 113
Control and validation procedures 114
Regression analysis 116
The discriminant model 133
Survival models 141
Summary 150
Notes 151

5 **Positive theories of corporate failure: III – Iterative models** 153

Introduction 153
Subjectively weighted variables 154
Recursive partitioning 156
Artificial intelligence 158
Neural networking 159
Summary 164
Notes 165

6 **Positive theories of corporate failure: IV – Early warning studies** 166

Introduction 166
Share price behaviour 168
Laboratory experiments 185
Summary 191
Notes 192

7 Positive theories of corporate failure: V – Case study research 195

Introduction 195
Turnaround studies 196
Case study research 199
Summary 202
Notes 202

8 The explanatory variables: I – Financial ratios 204

Introduction 204
Common characteristics in ratios 205
Statistical distributions 207
Industry specific ratios 211
The effects of applying different accounting policies 212
Summary 216
Notes 217

9 The explanatory variables: II – Non-financial ratio indicators 219

Introduction 219
Measures used in economics and finance studies 220
Measures used in management and corporate strategy studies 235
The acquisition/failure alternative 237
Summary 239
Notes 239
Appendix 9.1: Company liquidations 1972-1996 242

PART II: The empirical studies 243

10 The data 245

Introduction 245
The sample companies 245
The variable sets 247
Summary 249
Notes 250

11 Univariate analysis 251

Introduction 251
Univariate financial ratio models 251
Univariate qualitative variable models 254
Share price variables 257
Entropy models 259
Gambler's ruin models 261
Summary 267
Notes 268

12 Multivariate analysis: Logit and survival models 270

Introduction 270
Logit models – qualitative variables 272
Logit models – financial ratios 274
'Rolling' logit models 282
The impact of misclassification costs 285
Taffler variable models 286
The use of matched pairing as a control technique 288
Survival analysis 289
Summary 290
Notes 292

13 Multivariate analysis: Iterative models 294

Introduction 294
Neural network models 294
Comparison with logit models' predictions 300
Summary 301
Notes 302

14 Share price behaviour models 303

Introduction 303
Share price returns and event identification 304
Market identification of failing status 308
The information content of failure identification models 309
Summary 317
Notes 318
Appendix 14.1: The El Hennawy and Morris failure identification models 319

15 Case study analysis 321

Introduction 321
The causes of failure 323
The causes of failure, 1973-1983 323
The causes of failure, 1988-1991 332
Summary 340
Notes 341
Appendix 15.1: Short case histories of 25 failed companies 341

16 Summary and conclusions 373

Introduction 373
The background 373
The empirical findings 375
Summary 378

Appendix: Sample and control companies 379
Bibliography 383
Glossary 406
Subject index 422

Acknowledgements

The data on which the empirical studies reported in Part Two of this book were based proved far more difficult to collect than was originally envisaged. The assurances given by various parties that comprehensive files and records were readily available regrettably proved well wide of the mark. Unfortunately this was not immediately evident, and a great deal of effort was therefore required to construct the necessary data files. In the circumstances, thanks are due to all those former students who helped to collect the necessary material, but especially to Ian Finney. The assistance of Donna McDougall of Strathclyde University is also gratefully acknowledged.

The empirical studies reported in Part Two of the book were the result of extensive collaboration, and consequently joint authorship of chapters 11-15 should be acknowledged as follows:

11 S. Blackman, S. Kaletzis, R. Morris and N. Samios
12 S. Blackman, S. Kaletzis, R. Morris and N. Samios
13 E. Edmond and R. Morris
14 D. Brookfield and R. Morris
15 E. Almond, K. Cleaver and R. Morris.

The helpful comments of two anonymous referees (one a practitioner, the other an academic) are also acknowledged. As a result, the text has been altered to try to accommodate their suggestions (e.g. a section now deals with the assessment of whether or not a business is a 'going concern'; and a glossary is provided). The opportunity has also been taken to explain various issues in more detail (e.g. why the studies focus on the experience of listed rather than private companies; and why it is inappropriate to try to control for size when calculating share price residuals, since the former is a potential explanatory factor in the failure identification models being studied.).

Finally, thanks should be expressed to the Research Board of the Institute of Chartered Accountants in England and Wales, which made a grant available to meet the cost of collecting the data.

Preface

Over the past thirty years businessmen have become increasingly aware that there are procedures which claim to be able to distinguish failing from non-failing firms. However, this poses some awkward questions, namely:

> If indeed there is a relatively straightforward way of discriminating between financially sound and financially distressed companies, why is it – given the large potential payoffs – that analysts are not already using or mimicking the procedures?
>
> If there is a widespread belief that the discriminatory procedures *forecast* accurately, then they should be self-fulfilling. This should mean that once they enter the public domain they will no longer have any predictive power: i.e. they should no longer be able to distinguish between failing and non-failing firms, except immediately before the distressed companies go bankrupt.
>
> To what extent are the models merely capturing information that can be inferred by using relatively crude methods of analysis: i.e. how much of the news they appear to contain is incremental?

To be fair, academics do usually enter caveats about the use of their models, although the qualifications are often buried in the small print. Instead, the impression is generally given that their procedures are extremely good at the discriminating between failing and non-failing businesses.

In fact, what the models really seem to be indicating is (rather unsurprisingly) that companies which go bankrupt overwhelmingly report low profits and high borrowings immediately before their demise. But what is generally not pointed out is that the reverse is *not* the case: i.e. not all companies reporting low profits and high borrowings collapse. Indeed, it is only a minority of companies exhibiting these signs of financial distress which eventually go bankrupt. Thus out-of-sample even the best performing models – derived using the latest statistical and computer based procedures – report unacceptable misclassification rates for non-failing companies of around 20 per cent.

However, the search for a model which will give its author an advantage (albeit a short-lived one) is likely to continue. Moreover, no analyst can afford to ignore a supposedly predictive device which is referred to by his/her rivals.

In order to shed more light on the subject, and to provide answers to the three questions posed above, it is necessary to survey the extensive literature on corporate bankruptcy, to which academics in a variety of disciplines have made contributions. In this way it should be possible to adopt a critical stance and review both the arguments put forward and the empirical evidence gathered from a number of different sources.

Given that this book is aimed at both practitioners and academics, its written style is deliberately not that which would be appropriate for an academic treatise. Rather, wherever possible an attempt has been made to try to explain the issues and procedures – even straightforward bivariate regression, for instance – in terms which hopefully an interested practitioner might understand. Equally it is intended that the reader should be able to 'pick-and-mix': i.e. select topics which are of special interest and skim through the rest.

For the record, the results of the empirical studies reported in Part Two suggest that, when allowance is made for potential sampling bias and overfitting, the ability of the models to discriminate between bankrupt and surviving companies is considerably less than is generally claimed, and much more like that reported for well tried models out-of-sample. Nevertheless, a careful study of the evidence yielded by various models and by case study research is likely to improve understanding of the phenomena which lead to bankruptcy. Moreover, society in general should benefit if researchers are able to extract some incremental information from data in the public domain but which has previously not been fully exploited.

Executive summary

Chapter 1: The background

Finding ways of trying to identify failing companies as early as possible is clearly a matter of considerable interest to investors, creditors and auditors, especially as upwards of a third of newly established private companies collapse within five years of incorporation. The rate of failure amongst listed companies is much lower – around 2 per cent per annum – but for a variety of reasons (e.g. the availability of accounting and share price data) most studies of corporate bankruptcy relate to quoted companies. Yet given that many, well publicised models purport to 'predict' potential bankruptcies, it is puzzling that their forecasts do not appear to be reflected in the behaviour of share prices. If they were, one might reasonably expect the latter to fall sharply in reasonably efficient capital markets as soon as potential bankruptcy is identified, and the companies to be forced into receivership. In fact, this is generally not what seems to happen. Consequently, unless one believes that securities markets are grossly inefficient – which seems unlikely given their highly competitive nature – it would appear more plausible that there are defects in the models themselves.

In fact, close scrutiny of the models, derived using a variety of techniques, indicates that they frequently exhibit high misclassification rates outside the sample period. Thus, while immediately before their demise they often correctly identify over 90 per cent of companies which collapse, they typically also diagnose an unacceptable 20 per cent of surviving companies as prima facie failures. This in turn appears in large part to be because the procedures used frequently do not adequately allow for the true incidence of failures in a population; and because heterogeneous data are pooled from a broad cross section of companies over periods of time when underlying economic circumstances are changing. As a result, the models tend to identify the lowest common denominator of failing businesses, such as low profits, high levels of borrowing, and the relatively small size of financially distressed companies: i.e. they focus on the *symptoms* rather than the underlying *causes* of bankruptcy. Moreover, they tend to perform best when a company is beyond redemption and the news content of the 'predictions' is minimal. In the circumstances, the models are probably best used as a shorthand procedure for summarising data about a business.

In any bankruptcy study it is necessary to consider the meaning of the terms 'failure' and 'prediction'. The former can embrace various types of financial distress, ranging from bankruptcy at one extreme to a decline in profitability at the other. 'Prediction' for its part can refer to an ability to identify an event before it occurs; or instead an ability to discriminate correctly afterwards. The models are inevitably derived using historical data, although they are generally tested on 'hold out' samples to see how well they might forecast future failures.

A general weakness of failure identification models is that there is usually little or no economic theory underpinning them which could indicate *why* certain companies might be expected to fail and others survive and prosper. Rather, they tend to be derived on an ad hoc basis. Further problems which give rise to unjustified inferences are inadequate allowance for the fact that only 2 per cent of listed companies in a population are likely to fail in any one year; and various other 'sample selection biases' (e.g. data is less easily available for defunct businesses, giving rise to 'survivorship bias'; and there is inadequate allowance for industry and general economic factors because of the use of a 'matched pairing' technique).

Even in highly competitive financial markets, investors and creditors have strong incentives to identify financially distressed companies, and in particular to be the first to get the news. However, any advantage gained is likely to be short lived. Nevertheless, the search for early warnings of financial distress is reflected in the pressure placed on auditors to flag up impending difficulties in their reports on the annual accounts of client companies. In the circumstances, how to assess whether or not a business is a going concern is a matter of considerable concern to accountants in public practice, especially as the risk of litigation has increased. This has given rise to considerable discussion of the issue, and there is a growing amount of guidance in law (affecting both directors and auditors), the code on corporate governance, the Accounting Standards Committee's SSAP 2, the Accounting Standards Board's recommendations on preparing the Operating and Financial Review, and the Auditing Standards Board's SAS 130. Nevertheless, the incidence of going concern qualifications in audit reports is still extremely low, and the accounts of most companies which eventually go bankrupt are not so qualified. Furthermore, such warnings, when given, are not generally issued until the companies concerned are effectively beyond redemption. On the other hand, many companies whose accounts receive a going concern qualification in fact survive. Nevertheless, despite this evidence there have recently been suggestions that the signals given by failure identification models could be used as the pretext for entering such qualifications. This would seem dangerous, given the weaknesses in such models, which it seems are often not fully appreciated by practitioners.

Part I: Previous research

Chapter 2: Normative theories of corporate failure

As mentioned previously, most bankruptcy prediction models are derived on an ad hoc basis with little theoretical underpinning. However, there are a number of theories which inform a general understanding of corporate failure.

The first group of such theories views financial distress as the result of disequilibrating shocks. Some (such as chaos and catastrophe theories) can equally well be applied to the natural sciences, the idea being that an unexpected event disturbs an equilibrium and has 'knock on' effects. Obvious examples are the way the body reacts after an accident or with the onset of a disease; or how the countryside recovers its equilibrium after a storm or a fire. In the same way, economic systems are frequently knocked out of kilter by unanticipated shocks: e.g. a hike in oil prices; the outbreak of a war; a sudden change in exchange parities or interest rates; a prolonged strike by employees; or the collapse of a bank or a major company in an industry.

Chaos and catastrophe theories provide a general framework for studying the ways in which systems adjust to an unanticipated event, but they offer relatively little in the way of statistical procedures for analysing the position, certainly with respect to bankruptcy. However, the entropy (or informational decomposition) approach has been used in a number of studies, the idea in terms of its application to financial distress being to examine changes in balance sheet structures over time. The suggestion is usually that one might reasonably expect the proportions of current/non-current assets and claims to alter more for failing companies than for their surviving counterparts, reflecting the disequilibrating shock to which they are subject.

In industrial economics, academics have developed theories to explain why there are changes in market structure (i.e. why a few companies come to dominate in some industries). This in turn has led them to examine possible causes (e.g. the existence of so-called 'barriers to entry'; and the 'exit characteristics' in an industry: i.e. through mergers or bankruptcies). Another factor identified is the financial structure of a business. While in a 'perfect market' setting the level of gearing should be of little consequence, in practice firms and individuals are unable to borrow at will, and bankruptcy risk is a factor which has to be considered, even where there are sizeable 'clienteles' of investors who might be prepared to develop investment portfolios to offset the risks inherent in holding any one security.

Another aspect examined by industrial economists is the growth characteristics of an industry, the implication being that there is a (probably changing) optimal size for a firm operating in a particular industry. If a firm is too small, it is unlikely to survive. Similarly, economic geographers have argued that in some (if not all) industries, location is a critical factor in determining a company's costs, and hence its ability to compete with rival firms.

Various financial models of corporate failure have been developed. The simplest is to view the firm as an option in the hands of the shareholders, giving them the right to buy the business back from its creditors at a future specified time (e.g. when loan stock can be redeemed) if there is then a positive equity value. Clearly, as with all options, the value will be greater the higher the variability of expected cash flows. Another way of explaining the bankruptcy phenomenon is to consider the compound probability of a company running out of cash (i.e. of a net cash outflow in one period being followed by net cash outflows in successive periods). Where a company is subject to 'capital rationing' (i.e. it cannot raise new capital or borrow easily), such a path (the 'gambler's ruin' scenario) will lead to a firm's failure. However, where there is no capital rationing, a company should be able to raise new capital regardless of its operating cash flow position so long as investors believe that the value of the business is positive. In practice, even listed companies are likely to be subject to some degree of capital rationing, and the incidence of inflation is another factor which is likely to make it more difficult for firms to survive when faced with a succession of negative cash outflows over time. In the circumstances, one might therefore expect firms with highly variable operating cash flows (or, as a proxy, profits) to be more likely to fail than those with more-or-less constant cash flows or profits. Models have been developed with some success to test this argument, relating the variability of cash flows to an opening 'cash reservoir'. Moreover, the more conventional failure identification models, developed on a more ad hoc basis, can be rationalised and justified in terms of the gambler's ruin hypothesis.

Essentially the option pricing and gambler's ruin models argue that a firm will not go bankrupt until such time as the going concern value only equals the break up value of its net assets. It is therefore necessary to consider how firms are valued in the market place. The appropriate setting, however, is to view a company's securities in a portfolio context, and it is relatively easy to show that combining assets together will reduce risk to an investor, except for that element of risk which is 'market related'. This can be done within the framework of the 'capital asset pricing model' (CAPM) and/or 'arbitrage pricing theory' (APT), the principles of which are relatively easy to understand.

More recently, academics have begun to explore so-called 'agency' models of corporate failure. These attempt to analyse the nature of the contractual relationships between various parties (such as shareholders, creditors and managers). This of special interest, given that financial distress is often resolved by various interests being redefined (e.g. bankers accepting equity shares in exchange for cancelling debts outstanding).

Finally, the management and business strategy literature has attempted to popularise basic concepts in industrial economics. Writers have in fact tried to identify the key causes of financial distress and then undertake case study analysis to see what evidence there is of their existence in practice. The main variables discussed are shortcomings in management (a portmanteau label which, with the advantage of hindsight, can be used to cover most eventualities); inadequate financial controls; slow reaction to changes in the economic environment; the incidence of mistakes by management; and the fact that the symptoms of distress (e.g. declining profits, increasing indebtedness) become more evident as failure approaches.

Chapter 3: Positive theories of corporate failure: I – Univariate models

The most widely applied models used to identify companies at risk of failure are so-called 'univariate' models. These involve the analyst examining a series of variables (usually financial ratios) one-by-one. However, the 'traditional' basis for interpreting financial statements is full of potential flaws which are rarely identified in standard text books. In particular, it is essential to identify an appropriate bench mark against which to compare a ratio; and it is also important to remember that the figures relate to a legal rather than an economic entity. Moreover, while it is widely acknowledged that economic events can be accounted for in a variety of equally acceptable ways, the reverse is far less commonly recognised, although anyone with practical experience will be aware of the fact: namely, that various combinations of economic circumstances can give rise to similar accounting numbers. It follows that unless an analyst is careful, it is extremely easy to draw incorrect inferences from figures reported in a company's financial statements.

Traditional textbooks do not generally examine the nature of accounting ratios: e.g. there are commonalities and interrelationships between them; there is an implicit assumption of linear proportionality; the statistical distributions tend not to be symmetric (i.e. they are not 'normal'); and – perhaps most important of all, and well understood by practitioners and professional analysts – the means and distributions of particular ratios tend to vary between industries and even between different types of firm operating within the same sector. (These matters are further discussed in chapter 8.)

Traditional ratio analysis usually focuses on three matters: a company's long term financial position (i.e. primarily its financial structure); its short term financial position (i.e. its liquidity); and its profitability and efficiency. With respect to long term financial position, financial statement indicators only paint part of the picture. It is probably best to focus on a company's cost structure (i.e. its mix of fixed and variable costs), the variability of its net revenues, and the redemption terms and security available to support additional borrowing. As for the short term position, the main aim should be to try so far as possible to construct a crude cash (or working capital) budget, much as internal management will do. In such circumstances, financial ratios should only be used for screening purposes – and particular care should be exercised when referring to ratios where both numerator and denominator have been drawn from a balance sheet as the coefficients are particularly susceptible to distortion. 'Mixed' ratios (i.e. where one of the figures represents a flow through time) may be more helpful, but there are still potential problems (e.g. the 'position' figure taken from a balance sheet is unrepresentative; lack of linear proportionality; etc). Overall, it is probably safer to concentrate on the picture revealed by a flow of funds statement and try to use this as the basis for projecting future likely cash flow outcomes.

As for assessing profitability and efficiency, it is important to examine both proportional and absolute changes. Moreover, care is needed when examining margin ratios, as it is necessary to allow both for the impact of the economic factors which help to determine outcomes; and for the incidence of particular accounting conventions, either over time within a single firm or as between different companies. At first sight it may appear best to refer exclusively to 'rate of return on capital employed' (ROCE) ratios. However, regardless of which of several versions of this statistic are used, it is important to remember that the numbers can be quite seriously affected by the accounting conventions used. In particular, other things being equal the so-called 'concept' of prudence (i.e. conservatism) employed by accountants can increase reported returns for firms which are declining or only growing slowly; whereas those which are expanding will tend to be penalised. This results from overdepreciation (in economic terms) of relatively new wasting fixed assets and underdepreciation of relatively old ones, and the phenomenon will be accentuated when price levels are increasing. Similarly, firms which invest heavily in (say) R&D will tend to be penalised; and the basis of valuing stocks and work in progress is another potentially distorting factor. Certainly, as anyone with practical experience knows, such factors (the incidence of which differs from industry-to-industry and even from firm-to-firm) makes it very difficult to make meaningful cross sectional comparisons on the basis of figures alone; and it also indicates the potential dangers of trying to construct models based on ratios, especially where data have to be aggregated across industry sectors.

In fact, the first bankruptcy identification models devised were based on univariate ratios, the earliest matched pairing experiment dating from 1932. The best known model developed along these lines was that of Beaver in 1966, and its prima facie discriminatory power was impressive, being not all that inferior to the more complex models devised subsequently. Other univariate ratio models have since been developed, but they have tended to focus on specific explanatory variables, notably operating cash flows. However, insufficient allowance seems to have been made for sampling bias (especially for the fact that the incidence of failures in the population is nearer 2 per cent rather than the 50 per cent assumed). Moreover, in a follow up study, Beaver seemed to find that share prices reflected impending problems if anything a little earlier than the accounting indicators (and probably quite a lot earlier if allowance is made for the lag between year end and the publication of the financial statements).

Chapter 4: Positive theories of corporate failure: II – Multivariate models

Most failure identification models that have been developed are not univariate in nature but multivariate: i.e. the status of potential bankrupt/non-bankrupt firms is determined in terms of a series of variables, thus allowing for simultaneous interactions between them. The variables are usually in the form of financial ratios (the properties of which are examined in chapter 8), but are also sometimes qualitative in nature (see chapter 9). As for the multivariate models themselves, they can be derived in two main ways: using statistical procedures or iterative (i.e. search) techniques. The former are discussed in this chapter, the latter in chapter 5.

One of the weaknesses of bankruptcy identification studies is the potential inadequacy of the control and validation procedures used. In particular, the matched pairing technique used and pooling of data over time tend to make models 'sample specific'; and as mentioned previously there is the difficulty of sampling bias, the true incidence of bankruptcy in the populations being studied usually being around 2 per cent rather than 50 per cent. Even where attempts are made to adjust for such bias, it is commonplace to allow simultaneously for differential costs of misclassifying failing and non-failing businesses. This latter adjustment generally offsets the impact of the sampling bias and has the effect of making it appear that the models perform well on the sample from which they are derived. However, when tested on hold out samples, and especially subsequently when they are applied in a practical context, the models seem to perform far less well.

The simplest type of multivariate model that can be derived is to apply regression analysis, with the dependent variable being a dichotomous fail/non-fail classification. However, it appears that some of basic requirements of the regression model (e.g. linearity) are violated. A more realistic procedure in the circumstances is logit regression, which potentially has the advantage that it should enable meaningful probabilities of failure/non-failure to be calculated. Yet the computational requirements are such that for the model to work properly it is necessary to work with large samples with the distinguishing characteristics. This is rarely possible with bankruptcy studies, which greatly reduces the advantage of using the logit model. On the other hand, it is relatively easy to adjust for sampling bias and make the model valid where the prior probability of failure is around 2 per cent rather than the 50 per cent usually assumed in order to use the matched pairing technique.

The logit model has been increasingly applied in recent years in bankruptcy studies, and various refinements have been introduced (e.g. to formulate models at various intervals before failure; to explore the incremental information generated year-by-year as bankruptcy approaches; and to develop multilogit models using data over a number of years to derive the discriminatory function). Other refinements include attempts to identify a number of states of financial distress and not just the two extremes, fail and non-fail. However, to be valid this would once again appear to require large samples of companies with the different distinguishing characteristics, a requirement which cannot generally be met. Another variation is to develop a 'rolling logit' model, which includes the dependent variable score of the preceding year's model as an explanatory variable for the current period.

Although regression would seem to be the most obvious statistical model to use for bankruptcy identification purposes, before the comparatively recent popularity of logit it was another but essentially equivalent statistical technique that was generally applied, discriminant analysis (DA). In fact, the technical requirements for using this approach frequently seem to have been violated, but this has not prevented its widespread use for deriving bankruptcy prediction models from different sets of explanatory variables (e.g. based on cash flows and inflation adjusted accounting numbers). The best known models are Altman's ZETA in the US and Taffler's Z-score models in the UK, and they have now been in commercial use for around 20 years. The former is applied by over 80 commercial clients and the latter by more than 40. But while they are known to identify correctly almost all failing listed companies, they also crucially misclassify around one fifth of surviving companies as being potentially bankrupt.

In recent years there have been attempts to develop 'survival models' in a failure prediction context. Whereas other models attempt to predict *which* companies in a population will go bankrupt within a given period of time (effectively a few months), duration models estimate the length of time for *all* companies up until their 'deaths'. There are various versions of the model, the simplest being that which assumes the probabilities of survival are strictly proportional. However, a more severe problem in practice is 'censoring' (i.e. inevitably the deaths of most companies in a population will not be known, but frequently, and more importantly, for all except cohorts of newly formed private companies it will usually be impractical to obtain and include all data since each individual firm's birth). Despite these difficulties, the approach has been used in a small number of studies, assuming explicitly or implicitly for long established businesses a common birth date. Despite its unsuitability for listed company studies, the models appear to perform about as well (or as badly) as their regression, logit, or discriminant counterparts. However, the most interesting application of the technique, given the high attrition rate amongst private companies, is its recent use by industrial economists, who have studied the survival patterns of newly created small businesses in various countries.

Chapter 5: Positive theories of corporate failure: III – Iterative models

Various iterative (or search) procedures have been employed to develop multivariate models, and in general they have performed as well (or as badly, when sampling bias is allowed for!) as those derived using statistical techniques. The best known are probably those which have been developed to produce credit scores, with the variables included generally being those most frequently referred to by analysts. The weightings applied to each explanatory factor are then found by a process of trial and error until the ability to discriminate between failing and surviving firms is maximised.

A development of the above approach is 'recursive partitioning', whereby there is a sequential search for the combination of weighted explanatory variables which best separates failing from non-failing businesses. Experiments along these lines have worked as well as the statistical procedures in correctly classifying bankrupt and non-bankrupt firms; and the technique has also been used to determine bond ratings.

A natural evolution of recursive partitioning is to develop 'artificial intelligence' models on a computer, using IF...THEN...ELSE statements. This has been done with as much success as other models, and a further step has been to employ 'neural networking' procedures. These non-linear models, which are supposed to mimic thought processes, are becoming increasingly popular in a variety of applications, including a number in the field of finance. The algorithms are expensive in terms of computer time, the number of iterations needed to adjust the model until it minimises forecasting errors often exceeding a thousand cycles. Moreover, it is not possible for outsiders to identify precisely how the initial weights attached to the half a dozen or so explanatory factors alter as the models are continuously refined. This is because there are complex interactions with so-called 'hidden' variables. The process works with the model being derived from a 'training set' sample and then being tested on a hold out sample. A number of such models have been developed in recent years to try to discriminate between failing and non-failing companies, including some in the UK. Generally they perform at least as well as, and often slightly better than, more conventional statistically derived multivariate bankruptcy identification models. However, there is a hint of 'overfitting', and the error rates out-of-sample would still appear to be of such magnitude as to make them of limited practical value to investors, creditors and auditors.

Chapter 6: Positive theories of corporate failure: IV – Early warning studies

As indicated previously, if failure prediction models have the strong discriminatory power that is usually claimed for them, it is puzzling that share prices and credit ratings do not immediately reflect the forecasts, forcing the companies into bankruptcy straight away. The implication is either that financial markets are seriously inefficient in absorbing information; or, alternatively, that claims with respect to failure identification models are somewhat overstated. The purpose of this chapter is therefore to examine the evidence which indicates how decision makers react to the signals conveyed by the models. This can be done by studying, on the one hand, share price behaviour; and on the other how analysts react to such signals in tightly controlled laboratory experiments.

The methodology for studying share price reactions to specific news is well developed. The dependent variable is essentially security returns, which can be explained in terms of a number of independent variables – generally market- and industry-wide factors and the firm-specific event or news release that is the real focus of interest. However, there are many technical problems which have to be overcome when undertaking such studies if the possibility of drawing inappropriate inferences is to be avoided. Certainly there are a number of stock market anomalies which have been brought to light in recent years which have yet to be satisfactorily explained, but despite this the overall picture is that the markets are generally efficient in almost immediately impounding information perceived as having predictive power.

With respect to bankruptcy prediction models there are some additional problems (e.g. it is especially difficult to allow for market- and industry-wide factors). However, studies of share price behaviour are really rather different from most other 'event studies' inasmuch as the focus of interest is less on share price movements around the specific date of a news announcement, but rather the trend of relative share prices over a much longer period leading up to failure. Consequently it is far less important in such studies to try to allow for each of the many factors which must be taken into account when undertaking a more conventional event study.

The evidence from previous studies, even allowing for the inevitable imperfections in the research methodologies applied, is fairly conclusive. It appears that the market gradually marks down the relative share prices of companies which eventually go bankrupt, beginning on average some 2-3 years before their eventual demise. This is slightly ahead of the signals transmitted by most failure prediction models. Moreover, it also appears that the share prices of surviving companies which are signalled as prima facie failures react in a not dissimilar way. Overall, the evidence does not seem to support the argument that bankruptcy prediction models are imparting a significant element of news to the market, but rather that they appear to be capturing information that has for the most part already been impounded in share prices.

Behavioural research involving laboratory experiments has been extensively used in financial reporting and auditing contexts to study the reactions of analysts and decision makers to specific situations and news announcements. In practice, the major problems with such experiments are on the one hand to ensure that the scenarios are sufficiently realistic; and on the other that the controls are adequate. In terms of bankruptcy prediction, subjects have tended to perform well in distinguishing between failing and non-failing firms, although it is noticeable that their discriminatory ability declines if they are unaware of the proportions of failed and non-failed businesses in the population that is the subject of the study. Moreover, in most experiments the subjects perform slightly less well than mechanistic bankruptcy prediction models; and there is evidence of 'hindsight bias'.

Finally, there have been several behavioural research studies focusing on the way in which information can be presented. In this context it appears that a useful shorthand way of summarising key accounting ratios in a bankruptcy setting is to represent them in terms of human faces, with a smile suggesting a healthy financial outlook and a frown impending problems.

Chapter 7: Positive theories of corporate failure: V – Case study research

An alternative approach to the study of corporate failure is to engage in case study analysis. This has the advantage that it should be possible to examine not just the symptoms of failure but also the causes. Equally, the interaction between different variables can be identified.

Most studies of this type are 'turnaround' or 'sharpbender' studies, which with the advantage of hindsight identify how some companies have recovered from financial distress while others have not. The framework for such analysis is usually the somewhat eclectic theory of corporate strategy, with writers such as Argenti and Slatter identifying a series of factors which can often contribute to corporate failure.

The case studies themselves tend to be largely descriptive in nature, and the general conclusion unsurprisingly is that financial distress cannot usually be explained in terms of one or two variables. Rather, it is the result of a conjunction of events, some of them controllable by management, others not, and to some extent each potential failure can be viewed as 'situation specific'.

Chapter 8: The explanatory variables: I – Financial ratios

There are significant commonalities amongst accounting ratios, indicating that generally they do not exclusively measure just one financial characteristic of a company. Consequently it is necessary to screen the potential variables to try to ensure that overlaps between them are minimal; and that they are stable over time in capturing specific economic characteristics. This can be done in a variety of ways, including factor analysis and multidimensional scaling.

Conventional interpretation procedures seem to assume that the relationships between accounting numbers are strictly linear. In fact, this is the exception rather than the rule, and it helps to explain why observed ratio values are not usually normally distributed. This can be a problem when developing ratio based models. However, more important in terms of discriminant analysis is multivariate normality, a requirement that is rarely met.

A particular problem is the fact that the average values for individual ratios tend to vary substantially between quite narrow industry categories. One way of trying to handle this would be to measure variations from such an industry average, but this has rarely been attempted. 'Creative accounting' is another factor which can undermine the validity of specific ratios, and there is certainly evidence of manipulation when companies are in financial distress – although whether it fools analysts is highly dubious.

Chapter 9: The explanatory variables: II – Non-financial ratio indicators

Financial ratios are not the only variables which might explain corporate failure. It is clear, in fact, that failures increase when an economy goes into recession; and, moreover, certain industries seem to suffer more than others. Various procedures can be used to try to allow for these factors, although their use is the exception rather than the rule in most studies. However, special models have been developed for private as opposed to listed companies; and there have also been attempts to explain financial distress by geographical location.

A variety of non-accounting firm specific variables have been, or could be, used in bankruptcy identification studies. Examples are the firm's age since incorporation; the degree of diversification in its activities; changes in lines of business; changes in company name; rates of organic growth; records of acquisitions and disposals; the existence of closure and redundancy costs; dividend policy; years since a dividend was last declared; years since a profit was last reported; years since sales last increased; share price returns; bond yields; and published risk indicators. Other non-accounting measures can relate to directors (e.g. the proportions of shares they hold; changes in their holdings; changes in the board; and changes in directors' remuneration); to the accounting year end date (e.g. changes in year end; and the lag between year end date and publication of the accounts); to changes in accounting policy; to the auditors (e.g. changes in the auditors; qualifications in their reports; changes in the lag between the year end and the date of the auditors' report; and changes in auditors' remuneration); and to indebtedness (e.g. with respect to changes in debt covenants or the register of charges; and bond and credit ratings). Various studies have been undertaken using one or more of these qualitative variables in bankruptcy identification models, and frequently they have been found to have explanatory power.

Another variable that has been studied in this context is the characteristics of the chairman's report. Various techniques exist for textual analysis, the aim being either to assess readability or the extent to which a report will be understood. As one might intuitively expect, the reports of failing companies are more complex than for non-failing counterparts as the chairmen attempt to explain the position. However, it is unclear if, when sampling bias is allowed for, the existence of a complex report necessarily means that a company is financially distressed. It seems unlikely.

As has been mentioned previously, writers on management theory have often identified key factors which can lead to corporate failure (e.g. weaknesses in accounting and control systems, a slow response rate to changes in the environment, and poor overall management, leading to too many mistakes when decisions are made). Various proxy variables have been used to measure the incidence of these factors, usually by undertaking questionnaire surveys.

Another area where non-financial indicators are relevant concerns studies of the acquisition/failure alternative, a number of which have been undertaken in the UK context. They suggest it may be possible to identify financially distressed firms which will be rescued by a take-over rather than go into receivership, although as usual the picture is not entirely clear because of the need to adjust for sampling bias.

Part II: The empirical studies

Chapter 10: The data

The data used for the empirical research studies reported in chapters 11-14 comprised three sets of 111, 75 and 61 matched pairs of listed companies. The first two were determined by the ability to obtain detailed accounting and qualitative data for periods of five and ten years respectively before the bankruptcy of the failing company. In the end, 19 accounting and 16 qualitative variables were used as the primary basis for developing the models. The third sample represented those pairs of companies for which five years share price return data were available. The periods covered were all within the time frame 1973-1983.

In addition, the circumstances of 25 of the failed companies were examined in some depth so that the case studies reported in chapter 15 could be undertaken. For control purposes, a further 21 listed companies which failed between 1988 and 1991 were selected, and these were used both as an inter-temporal hold out sample to test the various models previously derived; and for comparative purposes with respect to the case study analysis.

Chapter 11: Univariate analysis

Logit models were used to assess the discriminatory power of a number of individual financial ratios. These measured various company attributes – namely profitability, liquidity, gearing, size and asset turnover. When the intercept term was suppressed, only profitability seemed to have much explanatory power in distinguishing between failing and non-failing firms. However, the introduction of a constant term greatly improved the performance of the models, although (as might have been expected) profit ratios were still the main indicators of likely bankruptcy, discriminatory power in the last two years of a failing company's life being relatively good with between 70 per cent and 80 per cent of companies correctly classified. The discriminatory power of individual liquidity and gearing indicators was rather less good, but it was still reasonably strong.

Comparing the incidence of qualitative indicators over successive five year event windows for the ten year data sample of companies showed clearly that audit qualifications and changes in lines of business, in registered charges, in auditors, in company name, and in financial year end were all more likely to occur with bankrupt firms. Moreover, failed firms were significantly smaller than their non-failed counterparts. By contrast, there was no evidence that the failed companies were more likely to change their accounting policies than their non-failed counterparts.

Rather surprisingly there was little difference in *annual* residual share price returns between bankrupt and non-failed returns. However, *cumulative* residual returns on shares in companies which failed or were 'signalled' as failures are significantly worse than on those in companies which survived or were 'signalled' as non-failures (see chapter 14).

Informational decomposition models performed reasonably well on matched pair data, and not much worse than many multivariate models tested on hold out samples. However, in overall terms their discriminatory power seems to be no better than that of the univariate logit models using a single profit ratio. By contrast, the gambler's ruin models (also derived using a matched pairing procedure) did rather better, classifying correctly with an almost 90 per cent accuracy rate in the final year before failure and a success rate generally over 75 per cent two years before bankruptcy.

But the major problem with all the univariate models is that if allowance is made both for the overrepresentation of failed firms in the sample and for the costs to decision makers of misclassifying surviving companies as bankrupt, the operational usefulness of the models is greatly reduced. On the other hand, this is a defect which equally afflicts multivariate models and undermines their practical usefulness.

Chapter 12: Multivariate analysis: Logit and survival models

Logit models were first derived for qualitative variables for the 74 paired companies ten year data sample employed in the single variable study, using the same five year windows (see chapter 11). The models fitted well, and the results reinforced the conclusions of the single variable study. Moreover, misclassification rates before allowing for sampling bias were very similar to those recorded in other bankruptcy identification studies, and were clearly lowest immediately before failure.

The next step was to develop logit models from the 111 matched pairs of companies using a number of different financial ratios. These proxied for five independent variables which might capture the following characteristics: profitability, liquidity, gearing, size, asset turnover, and asset proportions. The various models were then tested to see which performed best. The final version fitted well, with several of the individual ratios having strong explanatory power. Misclassification rates in the years leading up to bankruptcy were impressively low and comparable to those reported in similar studies undertaken previously.

However, when the prior probability of failure was altered to be more realistic the models did not work as well, and the misclassification rates, especially for non-failed companies, were much higher.

Similar results were obtained when the models were reworked from a reduced sample of 70 matched pairs of companies, using the remaining 41 as a hold out set. However, misclassification errors were substantially reduced for the 2 per cent failure probability model when misclassification costs were allowed for. The models were then tested on three other hold out samples, taken some years after the original study period: one of 21 companies which failed between 1988 and 1991, and two of 100 non-failed companies, in 1983-84 and 1993-94 respectively. As expected, the models' ability to classify companies correctly declined quite sharply – for bankrupt companies for the 50 per cent failure probability model; and for surviving companies for the 2 per cent failure probability model.

Further tests were undertaken to see how well models derived on one year's data classified failed and non-failed companies in other years; and how 'rolling' logit models performed. The results were again much as expected, with inclusion of the previous year's dependent variable scores in the 'rolling' logit models improving classificational accuracy. Comparisons were also made with logit models developed using the variables included in Taffler's discriminant models, and classificational accuracy was found to be similar to that achieved with the corresponding models described previously.

The validity of applying the matched pairing procedure was also briefly examined. Surviving companies were randomly assigned to two groups and then matched to each other by year and industry. Logit models were then developed, the expectation being that discriminatory power would be negligible. Surprisingly it was not. Moreover, further replications of the experiment produced similar results. This would seem to imply that the matched pairing technique itself introduces a degree of misclassification error into failure prediction studies. Clearly further research is required on this.

Finally, although survival models are not really appropriate for studies of listed companies which have very different birth dates (see chapter 4), out of curiosity the data were run through the appropriate statistical packages to see what happened. This was done first assuming equal numbers of bankrupt to surviving companies in a population and then a more realistic ratio of 2:98.

Chapter 13: Multivariate analysis: Iterative models

Two experiments involving the use of neural networking (NN) procedures were undertaken. Each was derived from a 'training set' of 41 matched pairs of companies and validated on a hold out sample of 20 pairs.

In the first study, without allowing for sampling bias the misclassification rates for non-failed companies were much lower than for bankrupt companies. By contrast, those for the financially distressed companies showed a gradual worsening as bankruptcy approached, reflecting the fact that the five independent variables used became increasingly unstable. Moreover, misclassification rates on the hold out set tended to rise when the number of processing elements (or 'neurons') exceeded two, suggesting a degree of 'overfitting'. However, the errors were reduced when the forecasts of the models were averaged over time. Overall, it was found that the models were not insensitive to the starting values chosen to initiate the simulations, but despite this the results were not dissimilar to those reported for other NN studies.

A different variable set was used to develop the second group of NN models, with share price returns being included this time as an explanatory factor. Although the models produced slightly more accurate classifications on the hold out data, this was at the expense of increasing errors on the training set.

To assess the relative forecasting accuracy of the NN models, corresponding logit models were derived. On the whole, the NN models seemed to produce slightly fewer misclassification errors, but their superiority over the logit models in this respect was by no means clear cut.

Chapter 14: Share price behaviour models

For this part of the study, monthly share price residuals were calculated as the net returns for each matched pair in the 61 twinned company sample. This should provide an adequate measure of market response, given the difficulties of estimating systematic risk in a bankruptcy context and the fact that all that is required in such a setting is a profile of average residual returns over a long window of time. Moreover, it would be inappropriate to try to allow for the impact of firm size on the calculation of residual returns as this factor is already allowed for in the bankruptcy identification models.

Various potential event dates can be identified, but none is entirely satisfactory because of heterogeneity in the sample. However, given that the focus of attention is merely average residual share price behaviour over a period of several years before failure, the choice should not be that critical. All that is really required is consistency in definition.

The behaviour patterns of the net returns for the whole sample were first examined, and on average they showed a clear downward spiral from 2-3 years before the failing company's last financial year end, the cumulative negative returns being very large during the final twelve months. This is consistent with findings in previous studies, where it has been interpreted as suggesting that the market marks down relative share prices slightly before bankruptcy identification models suggest financially distressed companies are potentially bankrupt.

The next step was to try to see whether the market marked down share prices of companies in line with the *signals* transmitted by bankruptcy identification models, rather than just in terms of their ultimate fate. This was initially done using the 50:50 and 20:80 failure probability logit models derived using financial ratio data and referred to in chapter 12. In fact, the results suggest that generally market behaviour is consistent with the signals transmitted rather than a company's ultimate fate, with analysts probably identifying the prima facie risk of failure slightly before it is reflected in bankruptcy identification models.

One difficulty with this approach is that the failure identification models employed were derived historically using data for the same period over which share price returns were studied. Not only can this lead to overfitting in terms of the classification of companies as prima facie failures and non-failures, but also analysts would not at the time be able to apply such well-specified models for diagnostic purposes. Consequently it was decided to repeat the experiment using two other bankruptcy identification models. These were devised some years earlier, and their discriminatory power seemed to be comparable with that of other rival models. The number of misclassifications was, as expected, considerably higher because the models were being applied out-of-sample. Interestingly, however, average share price reaction was again consistent with the signals of a company's prima facie status as a potential failure or survivor, except towards the very end of the life of a company which ultimately went bankrupt. Moreover, this result was generally robust when the procedure was replicated on randomly selected subsamples of matched pairs of companies.

Overall, the results are not inconsistent with the view that failure prediction models are of limited practical usefulness, since all they may really be doing is capturing information that analysts are already using to revise the probabilities they attach to likely bankruptcy.

Chapter 15: Case study analysis

The evidence reviewed above suggests that at best failure identification models for listed companies probably contain only a limited amount of new information for analysts. This is in part because they seem not only to measure a lowest common denominator, but also they tend to identify symptoms rather than causes.

In order to throw further light on these matters, the characteristics of a number of companies which failed in the sample period, 1973-1983, and in a later period, 1988-1991, were analysed. The main focus was 25 companies which went bankrupt during the first period, and the key factors which characterised their declines were identified. These could then be compared against a similar analysis of the main contributory causes of the collapse of 21 companies which failed in the second period.

For both samples, the years when bankruptcy occurred were for the most part in periods of deep and prolonged economic recession. It was not surprising, therefore, that there were certain common characteristics between the two groups: e.g. the extent of the downturns took all companies – and not just those which failed – by surprise; some industries were hit harder than others; firms most vulnerable were those where demand fell sharply, which had significant borrowings, and which faced high levels of unavoidable fixed operating costs. These factors were inevitably reflected in the key accounting indicators (namely, sharply declining profits and high gearing ratios). But it was also clear that in most industries there is a constant jockeying for competitive position. As a result, it was those companies which had become the weakest players at the onset of the recession which found it most difficult to survive. In particular, their plight was often brought about by a specific (but not especially unusual) misjudgement by management, which in better times would probably not have been so catastrophic.

But what was particularly interesting was that the combinations of factors which characterised failing companies in the two samples differed quite markedly. Thus in the first period, most of the bankruptcies arose as a result of an extensive shake-out in British manufacturing industry, when overcapacity in a number of key sectors had to be shed. This painful process was sometimes accelerated by the prolonged high foreign exchange value of sterling in the early 1980s. By contrast, in the second period many of the victims were firms which had grown rapidly in the boom years in the mid 1980s. In particular, companies most severely hit were those whose growth was on the back of the property boom, the collapse of which took the banks as much by surprise as anyone else and left them trying to decide if and when to call in their debts and precipitate the failure of their clients.

Overall, the implication is that bankruptcy tends to be very much 'situation specific' and is usually the result of a particular conjunction of events, most of which could not be accurately forecast either by management or by outside analysts.

Chapter 16: Summary and conclusions

A priori reasoning suggests that it ought to be very difficult to devise failure identification models for listed companies which will consistently be able before the event to signal probable bankruptcy. Yet despite this, the impression is often given that the models perform extremely well.

In fact, a close examination of previous empirical work suggests that academics have spent much of the past 60 years applying a number of statistical and iterative procedures to various different sets of independent explanatory variables, but for the most part without greatly improving discriminatory power. In part this is probably because the models tend to identify a 'lowest common denominator' distinction between failing and surviving firms – an inevitable consequence, really, of having to pool data over time for companies operating in different industries. This probably helps to explain why out-of-sample most models seem to generate rather high misclassification rates, often identifying around 20 per cent of surviving firms as prima facie failures. This error rate is such as to explain why analysts do not seem to react as one might expect if the models' predictions could be relied on. Moreover, the evidence from share price reaction studies is not inconsistent with such an interpretation, implying that the models are probably doing no more than summarising data already in the public domain.

The empirical studies on British data reported in this book tend to reinforce the above arguments.[1] Thus the classificatory power of balance sheet decomposition measures and of gambler's ruin and multivariate models does not appear to be all that different, even in the latter case when the independent variable set is altered to comprise qualitative variables reported over a five year window. Far more important is the effect of making allowance for an appropriate prior probability of failure, which greatly weakens a model's ability to discriminate accurately. Further, it is clear that misclassification rates rise quite sharply when models are applied out-of-sample – hardly surprising, in fact, as economic conditions will have changed and firms from different industries will be under threat.

Further insight into the failure process is offered by case study analysis, and this confirms what seems to be intuitively obvious – namely, that listed companies which collapse tend to be victims of an unfortunate conjunction of a variety of events which, for the most part, are largely unanticipated.

That said, it seems likely that analysts will continue to refer to failure identification models, not because they believe their use will give them a momentary advantage over their rivals, but rather because failure to do so could possibly put them at a short term disadvantage.

Note

[1] The results are summarised in Table 16.1 on p. 376.

1 The background

Introduction

The argument

Finding ways of trying to identify failing companies as early as possible is clearly a matter of considerable significance to businessmen and other interested parties. For instance, if an investor or creditor is able to predict a company on the path to bankruptcy before anyone else, he or she will be able to liquidate the investment or obtain settlement of a debt and so minimise losses. Similarly, it is vitally important for an auditor in preparing his or her report to be able to assess whether or not a company is a going concern.

In fact, the rate of failure amongst new small businesses has always been high, upwards of a third of newly established companies collapsing within five years of incorporation.[1] By contrast, the rate of failure amongst listed companies is much lower, the attrition rate being rather less than 2 per cent per annum. Nevertheless, it is a matter of some concern that the number of bankruptcies, both of listed and unlisted companies, has increased in recent years as the British economy has suffered a series of destabilising shocks (see Figure 1.1). Moreover, it is noticeable that over time financial distress appears to be experienced in different industry sectors.[2]

Against this background there has inevitably been a growing urgency on the part of investors, bankers, trade creditors, company directors and auditors to try to find better ways of trying to identify firms likely to go bankrupt,[3] a demand which researchers have sought to satisfy by developing a number of different procedures which aim to give early warning of financial distress.[4] But for a variety of reasons (see p. 27) this research has overwhelmingly concentrated on predicting the fate of *listed* companies rather than their far more numerous unquoted counterparts, where however the risk of failure is far greater.

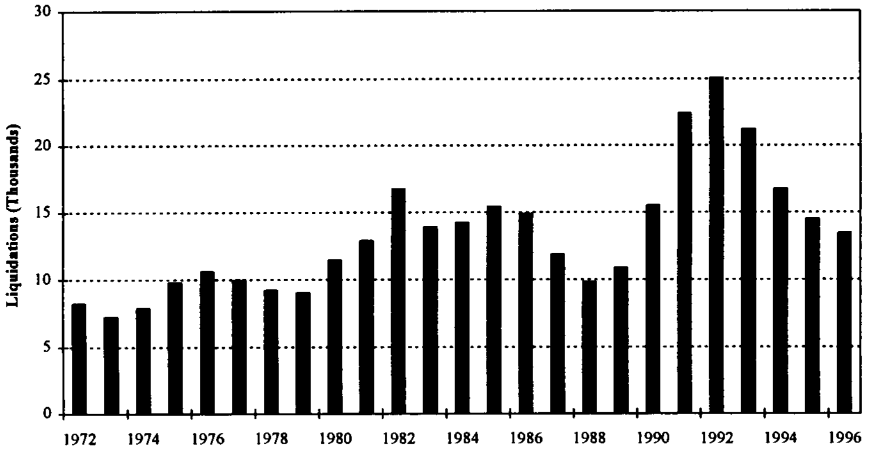

Figure 1.1 Company liquidations[5]

Source: Annual Abstracts of Statistics

What is puzzling is that if there is a well established method of identifying failing listed companies in advance of their final collapse, one might reasonably expect investors and creditors to use the procedure and immediately act upon its predictions. Consequently, as soon as a new and accurate forecasting approach has been established, one would expect it to be universally adopted. The result should be that listed companies forecast as very likely to go bankrupt ought to fail immediately.

In the circumstances it is therefore a little perplexing that many experts in the field seem to claim that they can successfully predict with remarkable accuracy which listed companies are very likely to fail and which are not.[6] Indeed, it is not unusual to find that success rates of over 90 per cent are claimed, not just immediately a new prediction model has been derived, but consistently thereafter when the diagnostic procedure has been well publicised, which stretches credulity to the limit.

Of course, investors and creditors, and the agents who work on their behalf, will search endlessly for a novel procedure that might give them a narrow (and presumably short-lived) advantage. If successful, analysts and their clients stand to make a lot of money – or, at least, in the case of the latter, not to lose heavily! It is therefore quite easy to believe that each innovation which improves the accuracy of predictions will be well worthwhile. But what is more difficult to accept is that a new approach will continue to be successful in terms of earning abnormal risk-adjusted returns after its existence becomes known and analysts are able to mimic its forecasts. Its prophecies should then become self-fulfilling. In fact, the argument and evidence presented in this book appear to provide an answer to the conundrum. Basically the 'prediction' models, however derived, seem to come up with a very similar (and unsurprising) answer: namely, that immediately prior to bankruptcy, the accounts of failing listed companies show high levels of borrowing and low profitability. The problem is that, outside the sample period and/or when allowance is made for sample selection bias, a relatively high proportion of *non*-failing listed companies – some 20 per cent – are also identified as prima facie failures.

Effectively this seems to imply that in any one year a UK analyst referring to one of the models that have been developed would correctly identify the dozen or so listed industrial companies which will go bankrupt within 12 months, but will incorrectly classify around 120 of the remaining 600 as likely to fail. Clearly analysts who might refer to the models to assist them in managing their portfolios are likely to try them out before relying on them blindly and thus make them self-fulfilling. It seems improbable, in fact, that a misclassification error rate of 20 per cent for non-failing listed companies would be regarded as acceptable, even allowing for the significantly higher costs of incorrectly identifying a failing company as sound when compared to those of misclassifying a non-failing company as prima facie bankrupt.[7]

Consequently it would seem that, where analysts refer to the models, they do so primarily as a shorthand procedure for summarising data about a company. Given the diversity of the businesses from which the underlying financial ratio data are derived, it appears that the classifications really just reflect the fact that listed companies reporting losses or low profits and which are burdened with debt are more at risk than otherwise similar companies which are recording reasonable profit figures and which have lower gearing ratios. But the plight of a financially distressed company should be fairly obvious anyway, and it would not be unreasonable to expect relative share prices to reflect the outward condition of a business, regardless of whether it eventually fails or not. Interestingly, the evidence reported later in this book seems to support this argument.

The implication therefore seems to be that there are few, if any, unambiguous early warning signals of impending bankruptcy. On the other hand, there is a significant proportion of listed companies – around a fifth, perhaps – which in any one year might be regarded as 'at risk', but the vast majority of which are turned round and/or do not fail.

Some basic issues

It is against this background that various basic issues will be reviewed in this chapter. This will set the scene for the argument and review of the evidence which follows.

In particular, it is first necessary to identify what exactly is meant by the term 'failure'. At one extreme it can obviously mean liquidation; but at the other it could just mean reporting a profit figure below that expected. In between are various possible definitions of what precisely is meant by the term.

It is also appropriate to consider what is meant by 'prediction'. In fact, it can mean being able to discriminate *after* an event or *before* an event – although it is only the latter which is really of interest to a decision maker.

In addition, there are a number of other fundamental methodological issues which need to be considered. For instance, before engaging in empirical research it is highly desirable to try to establish by deductive reasoning what factors might reasonably be expected to bring about the failure of a business. Such theories are known as 'normative theories'. By contrast, it is possible to develop 'positive theories' (explaining *what is* rather than *what ought to be*) through empirical observation.

In fact, most of the research into corporate bankruptcy seems to be driven by empirical evidence, and it is therefore appropriate to try to identify potential difficulties in developing research designs which might bias the results in a particular direction. One obvious distorting factor is the matched pairing technique generally adopted, when the annual incidence of failure is far less than 50 per cent of a population of companies in a given year. Interestingly there are ways this can be allowed for, and when suitable adjustments are made it appears that the discriminatory power of so called prediction models is substantially reduced. Another problem concerns the way in which the accuracy of predictions is calculated, either pairwise or across complete samples of failed and non-failed companies.

It is also necessary at the outset to say something briefly about informational market efficiency and various 'user needs'. In the case of the former, the argument presented above suggesting that a successful prediction procedure should immediately be applied by analysts implies informational market efficiency. Brief reference will therefore be made to the very substantial body of empirical evidence which suggests that, after allowance is made for search costs and rewarding special skills, financial markets do appear to be very close to being informationally efficient.[8]

As for user needs, it is also appropriate to focus more closely on what might be of interest to investors, creditors, company directors, auditors and other third parties. In particular, it may well be that some investors or creditors will be prepared to put their money into specific high risk companies, knowing that some will fail but that overall the rewards will more than offset the losses they will make on some of the investments or loans in their portfolios.

As for directors and auditors, they have always had to decide whether or not a company is a 'going concern', since that will determine whether a company's accounts will be approved by the board and the nature of the audit certificate attached to them. However, recent changes in legislation affecting the rights and duties of directors, as well as to rules concerning the audit of company accounts, have focused attention more closely on the matter. It is therefore necessary to consider the significance for both directors and auditors of being able to assess the likelihood of a company failing in the foreseeable future.

The meaning of 'failure'

'Corporate failure' fairly obviously encompasses 'bankruptcy', which for a company effectively means a creditors' liquidation or the appointment of a receiver. However, the net can be drawn more widely to embrace situations where there is evidence of 'financial distress'. It may therefore be useful to list a spectrum of potential indicators of such distress, beginning with situations where there is general agreement on what constitutes failure and working down to other circumstances which are more indicative of a company's possible financial difficulties, e.g.

(1) creditors' or voluntary liquidation, appointment of a receiver;
(2) suspension of Stock Exchange listing;
(3) going concern qualification by the auditors;[9]
(4) composition with the creditors;
(5) protection sought from creditors (e.g. under Chapter 11 of the US Bankruptcy Code);
(6) breach of debt covenants, fall in bond or credit rating, new charges taken over the assets of the company or its directors;
(7) company reconstruction;
(8) resignation of directors, appointment of a company doctor, etc;
(9) take-over (although not all take-overs are witness to financial distress, of course);
(10) closure or sale of part of the business;
(11) a cut in dividends or the reporting of losses; or
(12) the reporting of profits below a forecast or acceptable level; and/or the fall in a company's relative share price.

Generally corporate failure studies concentrate on the first few items in the above list, although some of the others may be taken as indicators of impending difficulties. There is also an extensive literature on changes in corporate bond and credit ratings[10] and on corporate turnarounds.[11]

The meaning of 'prediction'

Many studies on corporate failure specifically refer to *predicting* bankruptcy. It is therefore necessary to deal with another semantic issue which is all too rarely addressed in the literature – namely, what exactly is meant by 'prediction'.

In fact, 'prediction' has two distinct meanings, and it is important to distinguish between them.

(1) Prediction can mean 'identification' – i.e. in a narrow statistical sense it should be possible historically (or 'ex post') for a given population of companies to predict (identify) which businesses went bankrupt and which did not. Such an *autopsy* can be useful as a way of enhancing understanding of the phenomena which characterise corporate failure.

(2) Prediction can mean 'forecast' – i.e. it implies that it should somehow be possible to distinguish in advance (or 'ex ante') those firms which, within a given time span, will fail and those which will not.

For decision makers it is essentially the second of these which is of interest, especially if there is a procedure which would enable them to increase returns (reduce losses) on their investment portfolios. However, in a highly competitive market analysts would be expected to use *any* procedure which would enable them to distinguish between 'winners' and 'losers'. In other words, just as in betting markets, it is difficult to conceive of an 'unfair game' situation existing for any length of time. Consequently, although it is possible that a new innovatory form of analysis might give its creator a momentary advantage, this would quickly be eroded as other 'players' mimic the procedure.

But, just as the alchemists of old sought to find the mystical substance, phlogiston, that would turn base metals into gold, so investors seek to find ways of 'beating the market'. Clearly there are situations where such opportunities exist: e.g. the whole notion of project appraisal, where positive net present values (NPVs) are identified, implies that there are situations where there is short term disequilibrium in markets. However, these are likely to arise where peculiar factors exist which limit competition: e.g. where the nexus of skills and resources which exist within a company gives it a competitive advantage over its rivals; or where there are other 'barriers to entry'. Where there are no such impediments, any risk-adjusted 'excess returns' can normally be viewed as a reward for 'search activity', and if individuals act rationally in a competitive economic environment such excess returns should be eliminated after search costs have been taken into account.

Essentially what is being argued is that, in a competitive market environment, it would be surprising to discover a way of successfully discriminating between failed and non-failed firms. If there were a means of identifying companies which are likely to collapse, the diagnosis should immediately be reflected in market judgements. As a result, as soon as a business is forecast as being highly likely to fail, presumably bankers and suppliers would starve it of credit,[12] auditors would enter going concern qualifications, and equity holders themselves would attempt to bale out to minimise their losses. Consequently if a failure 'prediction' model is successful, not only would it become self-fulfilling, but it would lose its ability to forecast as its predictions would immediately be impounded by the market. Further, there are incentives which would help to ensure such an eventuality. Thus an analyst who concludes that a listed company's shares are overvalued will benefit most by selling short and then disclosing his privileged information. This should push down prices so that he can close his speculative position and take his profit.[13]

Of course, in practice there may be institutional barriers which prevent analysts making such easy money – e.g. it is not always possible to implement a short selling strategy even with a listed company, and virtually impossible with a private business.[14] In the circumstances, the analyst has to pursue an alternative policy: for example, sell his innovatory diagnostic model to bankers or brokers who hope – for a moment at least – to steal a march on the market.[15]

In short, one might well expect that – for listed companies, at any rate – so called 'failure prediction models' will not enable investors to outperform the market significantly. Instead they will merely tend to mimic analysts' diagnoses, which will already (or simultaneously) be reflected in relative share prices.

Methodological issues

Normative and positive theories[16]

A basic distinction is drawn in the social sciences between 'normative' and 'positive' statements. The former are assertions of 'what ought to be', and consequently they require the application not only of value judgements, but also of deductive reasoning. By contrast, positive statements assert 'what is', and they can therefore be tested using inductive reasoning against empirical evidence.

Theories are essentially constructed to try to identify and explain cause-and-effect relationships. All are to a greater or lesser extent abstractions from reality, and they are therefore based on a number of simplifying assumptions. (Indeed, if they were not stylised 'models' of the world, but merely duplicated it, they would add very little to our understanding.) Consequently the fact that the assumptions underlying a theory at first sight appear to be unrealistic should not necessarily be a matter of great concern, particularly if the effects of relaxing them are later examined closely. Further, there are 'instrumentalists' who argue that the realism of a model is relatively unimportant so long as operationally it seems to explain observed phenomena.

From the point of view of studying company failure, it is first of all desirable to try to develop well defined normative theories which might explain corporate collapse. Yet, as will be argued in the first part of this book, this has usually not been the case. Rather, researchers have chosen to gather empirical evidence and rationalise inductively what phenomena may have led to bankruptcy. But it is also necessary to devise appropriate methodologies for testing hypotheses relating to positive statements. Regrettably there are various problems in undertaking research studies into the subject which seem likely to have led to 'inference errors'.

Research methodologies

In the pure sciences (such as chemistry, physics and medicine) it is usually possible to conduct tightly controlled experiments in laboratories. Thus, for instance, when testing a new drug it is common practice to treat a representative cross section of patients with a pharmaceutical compound; while another similar control group of patients is given a placebo. The results can then be compared to see whether there is any evidence that the new drug has healing properties. Of course, when undertaking such experiments it is vital to try to ensure that the two populations of patients are to all intents and purposes identical, otherwise it is quite possible to draw incorrect inferences. Equally, it is necessary to try to identify all relevant outcomes. Sadly, from time to time methodological errors come to light (e.g. with thalidomide).

In the social sciences it is often impossible to construct realistic experiments within a tightly controlled laboratory environment (although there are some examples in economics and accounting: see below, p. 185 et seq.). More commonly researchers have to collect statistical evidence from real world events and use the data to test various hypotheses.

It is this latter approach which is usually employed in bankruptcy studies. Moreover, there are various reasons why the vast majority of failure prediction studies relate to listed rather than private companies. One obvious factor is that it is easier to access data for the former. However, there are other reasons why researchers prefer to use data for quoted companies: e.g. fewer than half the small companies registered in the UK publish their profit figures,[17] and even when they do such figures are unreliable because of the somewhat arbitrary nature of directors' emoluments; the picture of small companies' indebtedness is often incomplete because of the widespread use of guarantees by directors and others; as indicated previously, some 30-40 per cent of small companies fail within five years of incorporation, so it is not only difficult to build up a track record of performance, but the data are effectively censored, giving rise to a 'survivorship bias'; and the accounts of small companies are often filed 8 or 9 months after the financial year end, which – given the high attrition rate – frequently makes reference to the figures irrelevant.

Another consideration is the fact that the use of data relating to listed companies makes it possible to compare the results of discriminatory models against those previously devised for such companies in the US and UK. But perhaps the most important factor is that, by examining share price behaviour, it is possible to assess how analysts in a 'multi-person' decision setting appear to use failure prediction information. Clearly this is less feasible for a sample of small companies, where – for equity investors, at least – there is effectively a 'single person' decision environment. This means that it is necessary somehow to identify an individual decision maker's attitude to risk, rather than let the market mechanism take account of the differences in attitude through trading.[18] More generally, however, it would not seem unreasonable to infer that bankers and creditors might use failure prediction information in a broadly similar way to investment analysts operating in a stock market environment.

At a wider level, as with all research experiments, it is especially important that the methodologies used should be carefully worked out to try to ensure that no unjustified inferences are drawn. In fact, researchers into the possibility of successfully discriminating between failing and non-failing businesses have spent a great deal of time and effort in critically examining their models. Frequently their concerns have been about the technical statistical requirements of the various models they have used to study corporate bankruptcy. Unfortunately, however, the research methodology generally used seems to be seriously flawed in two key respects, and while these are usually now acknowledged, the full implications are rarely spelt out in detail.

The first difficulty is that for studies involving listed companies – the vast majority, it will be recalled – the number of failures is relatively low (e.g. between 0.5 per cent and 2 per cent per annum across all such businesses). This means that it has been necessary to 'pool' the data over time to produce reasonably sized experimental samples, the implicit (and unwarranted) assumption being that the underlying economic circumstances are the same each year. This procedure also produces a discriminator that applies to all companies in a sample. This means, for instance, that a critical value for (say) the gearing ratio, return on capital employed or a combination of the two should apply regardless of industry membership. The implication must therefore be that where there is a degree of heterogeneity in the sample – as there usually will be, since (as already indicated) to make the sample large enough observations have not only to be taken from different industry categories but at different points in time – the ratios which represent a 'lowest common denominator' will be the best discriminators. These are likely to be indicators of profitability and indebtedness, since companies in crisis are almost certainly going to find themselves with low or negative profits and with increased borrowings. In other words, the best discriminator ratios are likely to be *symptoms* of financial distress rather than the *causes*. In the circumstances, it seems intuitively likely that in most cases a failure identification model is probably not telling analysts much they don't already know.

A second problem relates to the difficulty of collecting data, even for listed companies. This can be particularly awkward in the case of defunct businesses, since by their very circumstances it will often be well nigh impossible to obtain their records. This introduces another example of a 'survivorship bias'.[19]

A further difficulty is that in constructing a control sample of businesses it has been commonplace to use a 'matched pairing technique'. This approach is understandable, if only because it is desirable to try to isolate key factors which distinguish otherwise similar firms. On the other hand, the procedure does have a number of drawbacks. Thus if the pairing criteria are years, industry membership and size, these three possible explanations of failure are automatically excluded from consideration. This may be regarded as unfortunate, since size and industry membership certainly seem to be key factors in determining a firm's vulnerability to collapse.

But far more important is the problem already referred to, namely that there is a *sampling bias* in that the matched pairing procedure produces an experimental sample of companies which is totally different from that in the real world. Thus it is not uncommon for bankruptcy studies to be based on populations of (say) 60 failed and 60 paired non-failed companies. This *state* (or *choice*) *based sampling* approach assumes that there is an equal 50:50 per cent probability of any firm selected from the wider population of companies being a potential failure. Clearly this is untrue: as indicated previously, the prior probability of a randomly selected listed company failing in any one year is really between 0.5 per cent and 2 per cent. As has been demonstrated by Zmijewski (1984) and Palepu (1986), the effect of this has been greatly to exaggerate the discriminating power of failure identification models (see below, pp. 114, 126, 138 and 140-1).

Despite these problems, a number of writers have been all too ready to point out with the 20:20 vision of hindsight that companies which have failed were correctly identified as being at risk by one or other of the previously devised bankruptcy prediction models. But just as it would be a logical fallacy to say that because all failing companies earn low profits, therefore all companies which earn low profits must fail, it is equally unjustifiable to infer that because a failed company would have been correctly identified by a model, therefore the model accurately discriminates between failed and non-failed companies.

Further, as previously argued, it is evident that the market can hardly believe the models, since if it did it would automatically and immediately bankrupt all companies strongly signalled as likely to fail. Consequently just what the models are telling the reader is unclear.

Finally, there is another problem that arises with the matched pairing procedure which is rarely referred to. This relates to the basis of comparison between pairs. The usual procedure is to apply the discriminatory function to *all* companies in a sample. However, an alternative approach is to rank each pair. This can give rather different results and makes some allowance for differences in the industry membership, size and time of failure which are captured in the pairing process.

In fact, it is easy to demonstrate the impact this may have if, for instance, a single ratio (such as return on capital employed) is the discriminator and is applied to a population of matched pairs of failed and non-failed companies.

Example:

Suppose that there are four pairs of companies with the following ratios (failed first, non-failed second): *Pair I* .02, .03; *Pair II* .03, .04; *Pair III* .09, .10; *Pair IV* .10, .11. In terms of a pairwise comparison of basically similar companies, the failed companies always have lower ratios: i.e. there is a 100 per cent correct discrimination. However, if the data are pooled and an average of the ratios is calculated, a cut-off point of .065 results.[20] According to this, for the first two pairs of observations both failed and non-failed companies are classified as failures; and the last two pairs as non-failures. Consequently at best there is only a 50 per cent success rate in categorising the companies.

The main problem which arises when a strict matched pairing criterion is adopted is that it is difficult to devise suitable hold out tests to assess the predictive ability of a model. All that can be done is to see whether a procedure correctly ranks *pairs* of companies out-of-sample (e.g. in terms of the greater risk of their going bankrupt). Consequently the best that can be achieved is to see whether a model derived in this way still correctly discriminates between pairs of bankrupt and surviving firms in subsequent periods. It will not generally be feasible to apply the procedure to the general population of companies and so identify those which appear to be prima facie failures.

Informational market efficiency

In a competitive environment investors, creditors and other interested parties would be expected to search actively for any clue which gives them an inkling of which companies are going to be winners and which losers. But one result of such activity ought to be that it will become increasingly difficult to steal a march on rival investors and creditors. Any informational advantage, though real enough and valuable, will be slender and short-lived.

This concept of informational market efficiency has been the subject of intense study by financial economists, and a substantial body of empirical evidence indeed suggests that in active markets there are few opportunities to make easy money. The only obvious occasions where investors can earn profits are where they have particular skills or advantages – or they are the result of pure chance, just as there will always be a punter who wins the national lottery or the football pools.

As an entrepreneur, an investor's advantage can be the result of a particular skill or combination of skills; or, alternatively, of having a degree of monopoly power. As explained earlier, such advantages give rise to the net present values which are the focus of attention in investment project analysis. By contrast, in highly competitive securities markets the opportunities to secure such a superiority might be expected to be much more limited. Certainly there should be many potential investors engaging in search activity in the hope of securing the equivalent of 'inside information' and thus gaining a momentary advantage. Moreover, it is true that someone has to be first with the news. But if any method of obtaining superior information is discovered, it is only to be expected that the technique will be copied.

It follows from this that, after extracting trends (e.g. for inflation or other systematic factors), prices of actively traded assets, such as securities, might be expected to follow a random walk. In fact, there is an impressive amount of empirical evidence which suggests that this is indeed the case, the phenomenon being referred to as 'weak form' market efficiency.

The reason why prices move, of course, is that unanticipated news items impact the market in a random fashion. It is therefore appropriate to see whether prices move in response to the release of news items (such as profit announcements, some of the content of which will not have been correctly anticipated by the market). Again, a substantial body of empirical research has been undertaken to test so called 'semi-strong' form market efficiency. In undertaking such studies, great care has to be taken to avoid the risk of making unjustified inferences. Thus it is necessary to allow for the market's expectations prior to a news announcement so that the incremental piece of news can be isolated. But it is equally necessary to eliminate systematic factors which are driving asset prices so that only the *residual element* responding to a specific news announcement is identified. Again, a substantial body of empirical research gathered over the years is impressive in suggesting that financial markets are 'semi-strong' informationally efficient.[21]

At the extreme it has been suggested that at any moment in time *all* analysts might be able correctly to perceive future outcomes and hence the true value of assets. This is the equivalent of suggesting that somehow they possess inside information. In practice it does not seem likely that this 'strong form' of informational market efficiency holds, and indeed empirical evidence suggests that it does not.

User needs

Information relevant to failure identification can be viewed from a number of angles, but it is probably most appropriate to concentrate on the user demand perspective.

The most obvious *decision makers* who might use accounting information would seem to be equity investors, creditors and employees. Yet typically they will hold portfolios of assets, and it follows that they will have to adjust their holdings to maintain the desired risk/return balance if one of the assets suffers a sharp fall in value. This is especially important to those users whose asset holdings are not widely diversified as their exposure to risk will be that much greater. This is likely to be a particular problem for employees, for some types of