Canada Research Chair in Risk Management
Working Paper 06-05

Estimation of the Default Risk of Publicly Traded Canadian Companies

Georges Dionne, Sadok Laajimi, Sofiane Mejri, and Madalina Petrescu*
Canada Research Chair in Risk Management, CIRPÉE and CREF
HEC Montréal

30 November 2005
Revised 21 March 2006

Abstract

Two models of default risk are prominent in the financial literature: Merton's structural model and Altman's reduced-form model. The former has the benefit of being responsive, since the probabilities of default can continually be updated with the evolution of firms' asset values. Its main flaw lies in the fact that it may over- or underestimate the probabilities of default, since asset values are unobservable and must be extrapolated from the share prices. The latter, on the other hand, is more precise, since it uses firms' accounting data—but it is less flexible. In this paper, we investigate the hybrid contingent claim approach with publicly traded Canadian companies listed on the Toronto Stock Exchange. Our goal is to assess how combining their continuous valuation by the market (structural model) with the value given in their financial statements (reduced-form model) improves our ability to predict their probability of default. Our results indicate that the predicted structural probabilities of default (PDs from the structural model) contribute significantly to explaining default probabilities when PDs are included alongside the retained accounting variables. We also show that quarterly updates to the PDs add a large amount of dynamic information to explain the probabilities of default over the course of a year. This flexibility would not be possible with a reduced-form model. We also conducted a preliminary analysis of correlations between structural probabilities of default for the firms in our database. Our results indicate that there are substantial correlations in the studied data.

Keywords: Default risk, public firm, structural model, reduced form model, hybrid model, probit model, Toronto Stock Exchange, correlations between default probabilities.

JEL classification: G21, G24, G28, G33.

* Financial support by Bank of Canada is acknowledged. We would like to thank Pierre St-Amant, Toni Gravelle, and Céline Gauthier for their valuable comments and Translation Services of the Bank of Canada for their very efficient contribution.

Résumé

Deux modèles sont généralement présentés dans la littérature financière : le modèle structurel de Merton et le modèle non structurel (forme réduite) de Altman. Le premier a l'avantage d'être souple, car les probabilités de défaut peuvent être continuellement mises à jour selon l'évolution de la valeur des actifs des entreprises. Sa principale lacune réside dans le fait qu'il peut surestimer ou sous-estimer les probabilités de défaut, parce que les valeurs des actifs ne sont pas observables et doivent être estimées par les valeurs des actions en bourse. Le second a l'avantage d'être plus précis, en utilisant les données comptables des entreprises, mais est par ailleurs moins souple. Dans cette recherche, nous analysons l'approche contingente hybride à l'aide de données sur des entreprises privées transigées à la Bourse de Toronto. Notre objectif est de montrer comment la combinaison des valeurs continues de marché des entreprises (modèle structurel) avec celles des états financiers (modèle à forme réduite) améliore les possibilités de prédiction des probabilités de défaut. Nos résultats indiquent que les probabilités de défaut structurelles prédites (PDs du modèle structurel) contribuent significativement à expliquer le risque de défaut lorsqu'elles sont incluses dans un modèle probit avec des données comptables. Nous montrons aussi que des mises à jour trimestrielles des PDs ajoutent beaucoup d'information dynamique pour expliquer le risque de défaut au cours d'une année. Cette flexibilité serait impossible avec l'utilisation exclusive du modèle à forme réduite. Nous avons aussi effectué une analyse préliminaire des corrélations entre les probabilités de défaut du modèle structurel. Nos résultats indiquent qu'il existe des corrélations importantes dans les données étudiées.

Mots clés : Risque de défaut, entreprise publique, modèle structurel, modèle à forme réduite, modèle hybride, modèle probit, Bourse de Toronto, corrélations entre les probabilités de défaut.

Classification JEL : G21, G24, G28, G33.

1. Introduction

In this paper, we investigate the hybrid contingent claim approach with publicly traded Canadian companies listed on the Toronto Stock Exchange. Our goal is to assess how combining their continuous valuation by the market with the value given in their financial statements improves our ability to predict their probability of default. The reference model is a reduced-form specification that uses only accounting information from financial statements available at the beginning of the fiscal year.

Two models are prominent in the financial literature: Merton's structural model and Altman's reduced-form model. The former has the benefit of being responsive, since the probabilities of default can continually be updated with the evolution of firms' asset values. Its main flaw lies in the fact that it may over- or underestimate the probabilities of default, since asset values are unobservable and must be extrapolated from the share prices. The latter, on the other hand, is more precise, since it uses firms' accounting data—but it is less flexible. Because the periodicity of the information is usually annual, the probabilities of default cannot be updated during the fiscal year. Technically, quarterly financial statements can be found, but their use is not widespread and they are not always audited by an external accounting firm. A final criticism of the second approach is that the predictions yielded by accounting data are not as reliable as those from market data.

To address these issues, Moody's has proposed a hybrid model. During a first phase, the probabilities of default are estimated using both methods, and subsequently, the probabilities of default from the structural model are integrated into the reduced-form model at each point in time as an additional explanatory variable. The appeal of the hybrid model is that it allows the probabilities of default to be continually updated by incorporating market information via the probabilities of default from the structural model.

The Bank of England estimated the hybrid model with data from British firms and obtained some interesting results. In this report, we apply the hybrid model to Canadian firms listed on the Toronto Stock Exchange. We also propose a preliminary analysis of the correlations between the estimated probabilities of default. This additional information is useful both for banks with portfolios including these firms and for eventual aggregate analyses by industry or by geographical region.

This paper is divided as follows. Section 2 reviews the principal models in the literature. Section 3 describes the database used. Section 4 presents the estimation of the structural model and section 5 that of the hybrid model. Section 6 provides a preliminary assessment of the issue of correlation between the estimated probabilities of default, and the final section summarizes the main results and concludes the study by proposing some extensions.

2. Review of the Principal Models for Evaluating Default Risk

In this section, we will focus exclusively on scoring models. We will not examine models for rating securities, such as bonds or credit risk portfolio management models, although we shall analyze correlations in the last section. Readers interested in portfolio models can consult the excellent literature reviews in Crouhy, Galai, and Mark (2000) and Gordy (2000), or delve further into the subject by studying the CreditMetrics model of Gupton, Finger, and Bhatia (1997). We shall also ignore the distribution of losses conditional on default. This distribution is often approximated by the product of the loss given default (LGD) and the risk exposure at the time of default (EAD). This risk exposure can be evaluated by the residual face value or the market value of the debt at the time of default (Dionne et al. 2005).

2.1 *Reduced-form (or non-structural) models*

The first scoring model for firms was developed by Altman (1968). Known as Z-score, it used five financial ratios to attribute a credit score to firms. These ratios, obtained from a discriminant analysis model, are weighted differently. The five ratios are: working capital/total assets, retained earnings/total assets, earnings before taxes and interest/total assets, market value of equity/book value of total liabilities, and sales/total assets.

An extension to this approach has been the use of linear or non-linear regression models to directly estimate the probabilities of default. These models allow several ratios and assorted financial data to be considered simultaneously and provide descriptive statistics for the estimated parameters. Furthermore, they can explicitly model non-linearities between the financial variables and the score and, finally, directly compute the probability of default. Logit and probit models are often used. Typically, the greatest variations in the probabilities of default come from ratios capturing firms' profitability, level of indebtedness, and liquidity. These models can be estimated on cross-section or panel data.

Several banks use this method for privately owned and publicly traded firms, either by buying a model or its extension, such as Moody's RiskCalc, or by programming their own estimation method. A problem they frequently encounter involves building an adequate database. Very often, credit files are not computerized or contain no historical data.

The main benefit of reduced-form models is their precision in estimating probabilities of default. Furthermore, they are easy to use for financial institutions equipped with strong database management systems. Beaulieu (2003) demonstrates how data from a Canadian bank can yield very precise probabilities of default. On the other hand, these models are not flexible, since they require information from financial statements. Thus, it proves very difficult to update probabilities of default over the course of a year. Some institutions may demand financial statements on a quarterly basis, but these are rarely audited. Another criticism is the absence in accounting data of anticipations regarding the future. They reflect the past well, but tell us nothing of the future. Market data are more relevant to forecasting probabilities of default.

2.2 *Structural models*

To respond to these criticisms, Merton (1973, 1974) proposed a structural model for calculating probabilities of default from market data. This model is a direct application of the Black-Scholes model (1973) for valuing European options. Stockholders own call options on the firm's assets, the strike price of which is the debt level. At the horizon date, they exercise the option if the value of the assets exceeds that of the debt, then reimburse the debt and share the surplus. Otherwise, the firm is in default and stockholders do not exercise their option. Their loss is then equal to the initial investment. Thus, the probability of default is the probability that the option is not exercised. To evaluate this probability, we need to assign a value to the option. After having computed the mean value of the asset and its standard deviation by iteration, we can find the distance to default (DD), which is equal to the gap between the mean asset value and the value of the debt, normalized by the standard deviation of the asset value. The shorter this distance, the greater the probability of default (PD).

To improve the basic Merton model, several extensions have been suggested in the literature. The one most relevant to our project is from Brockman and Turtle (2003). The main criticism levelled at Merton's model is that it does not account for the possibility that the firm may default before the debt matures. Also, only stockholders are involved in exercising the option. In general, firms default before this horizon date and lenders (banks and other creditors) possess options (debt convenants) allowing them to pull the plug on firms if they observe that the latter are in breach of their debt obligations or are simply unable to pay.

To formally account for these two dimensions, Brockman and Turtle (2003) propose using barrier options, which were introduced into the literature by Brennan and Schwartz (1978), Leland (1994), and Briys and de Varenne (1997). They use the down-and-out option, but other types can be applied. Thus, rather than stockholders who wait for the debt to mature before exercising a standard European call option, we have a down-and-out option on the assets in which lenders hold a portfolio of risk-free debt and a short put option combined with a long down-and-out call option on the firm's assets. The last part gives them the right (but not the obligation) to place the company into bankruptcy when they anticipate that its financial health can only deteriorate.

This option makes it possible to place the firm into bankruptcy as soon as the value of its assets reaches the barrier, at any time before, or at, the debt's maturity. The appeal of this option is that it can be adjusted to bankruptcy laws throughout the world, including in Canada. One simply adjusts the parameters of the model. It can also account for the various restrictions imposed by creditors on the borrowing firms, such as maintaining a low debt-to-asset ratio, limiting dividend payments, curtailing merger activity, and not issuing further debt without permission.

The authors demonstrate that Merton's standard call option model is a special case of the barrier option model, and test their model on U.S. data. They empirically verify that the barriers are statistically different from zero, thus rejecting the standard European call option for all years, capital structures, and industries studied. Finally, they show that their model, with a barrier option, dominates Altman's (1968) Z-score reduced-form model. It is important to emphasize that they may not have used the most advanced version of the reduced-form model.

Other versions of the structural model have been suggested in the literature, including Moody's KMV. We will not discuss them, since we do not use them. We refer the reader to the paper by Crosbie and Bohn (2003).

Duan, Gauthier, and Simonato (2004) demonstrate that estimating the parameters of the Brockman and Turtle (2003) model by maximum likelihood yields results that resemble those from the iterative estimation method used in this literature when the theoretical model is Merton's or when the capital structure is fixed. The appeal of the maximum-likelihood method is that it allows for statistical inference or, more specifically, calculating descriptive statistics for the estimated parameters, such as the value of the firm. Another important aspect in the contribution by Duan, Gauthier, and Simonato (2004) is that the correspondence between the two estimation methods is not necessarily perfect when we insert an additional parameter into the structural model to account for the capital structure, as when Brockman and Turtle (2003) estimate three parameters (the value of the firm, its standard deviation, and a parameter for the capital structure owing to the barrier option) instead of two. In this particular instance, the maximum-likelihood method dominates, since it yields unbiased estimates of the parameters. Wong and Choi (2004) developed a maximum-likelihood model with endogeneity of the capital structure. In our study, we use the maximum-likelihood method with two parameters, as in the Bank of England study. However, we conduct various sensitivity analyses by shifting the barrier, which is equivalent to a sensitivity analysis of the capital structure.

The structural approach has been criticized for overestimating the probabilities of default (Duan and Fulop 2005). The presence of trading noises on the exchange introduces randomness into the correlation between unobservable asset values and stock prices, thus annulling the one-to-one relationship between these two values. This relationship is very important, however, in applications of the maximum-likelihood method to unobservable data on assets, as in the structural model. The authors demonstrate that the presence of trading noises can affect the standard deviation in the Merton model. On the basis of their sample of securities, they find an average increase of 7.64 per cent in the standard deviation, with a maximum of 25 per cent, which has an effect on the projected probabilities of default.

Moody's developed a hybrid model that combines the benefits of the structural and reduced-form approaches: Estimates of the probabilities of default are both flexible and precise. It is an extension to their reduced-form RiskCalc model. Moody's version uses the structural model, adding the possibility that default may occur before maturity of the debt (Sobehart, Keenan, and Stein 2000). Sodehart, Keenan, and Stein (2000) do not use the structural approach to directly compute the probabilities of default. First, they find the value of the firm and its volatility, then they estimate the distance to default and use this measure in a logistic estimation of the probability of default. They show that integrating information from the structural model significantly improves the calculation of the probabilities of default in the reduced-form model.

The hybrid model allows supplementary information to be integrated. For example, the structural model does not account for liquidity risk, though this risk is generally significant in multivariate structural analyses. The hybrid model also allows the profitability of firms to be incorporated more directly than the structural approach, at least in the case of firms whose stock is not very liquid in responding to good news, for example. Finally, the hybrid model allows macroeconomic factors to be included when the estimation period is sufficiently long.

Tudela and Young (2003) present an application of the hybrid model. This application uses barrier options with a down-and-out call option. The researchers estimate various models on data from non-financial English firms for the period 1990–2001. One interesting particularity of their application is in the tests they propose. For their estimates of probabilities of default in the structural model, they use data on firms that did, and did not, default. Thus, they first verify whether the two firm types represent different predicted probabilities of default (type I and II errors). Second, they compare their hybrid model with other reduced-form models to see whether the added probability of default (PD) variable is significant for explaining probabilities of default. Third, they measure the performance of their model with power curve and accuracy ratio type instruments.
The authors establish that, over a one-year interval, the mean probability of default for the non-defaulting firms is 5.44 per cent, while that percentage rises to 47.33 per cent for those that did default. The results of the error tests are satisfactory. They tested the model for probabilities of default over a two-year interval. They also performed dynamic analysis and found that the probabilities of default rise as the date of default nears.

They further confirm that the PD variable is significant in their probit model, increasing the estimated likelihood significantly. Finally, the hybrid model outperforms other models when different tests (power curve and accuracy ratios) are used.

3. Database

In this section, we present the raw data and their sources, and explain how we constructed the database for the probability calculations.

Our initial database contained 3,712 financial securities representing 1,339 firms that did not default and 130 firms that did, for a total of 1,469 publicly traded Canadian firms. A number of firms issue several different equities, which explains the higher number of securities. The study period for the probabilities of default is from January 1988 to December 2004. The methodology we use to compute the probabilities of default with the structural model requires that our data window extend 24 months prior to the estimation period for the predicted probabilities of default in order to ensure statistical reliability.

Thus, stock exchange and accounting data were gathered for the period from January 1986 to December 2004.

3.1 *Defaults*

Firms that have defaulted are catalogued in Financial Post Predecessors & Defunct, Cancorp Financials (Corporate Retriever), and Stock Guide. Market data, to wit the market capitalization or market value (MV), are extracted daily from DATASTREAM's DEAD.LLT series, while the frequency of the accounting data, from Stock Guide and CanCorp Financials, is annual. Between 1990 and 2004, 130 firms were identified as being in default: 112 were bankrupt and 18 were (or are) undergoing reorganization.

Table 1: Distribution of defaults and reorganizations

Year	Non-default	Bankrupt	Reorganization	Total defaults
1990	379	9	4	13
1991	401	10	2	12
1992	427	10	1	11
1993	482	8	1	9
1994	522	7	0	7
1995	561	8	0	8
1996	627	4	0	4
1997	704	5	1	6
1998	788	5	0	5
1999	857	9	0	9
2000	939	6	3	9
2001	1,012	7	3	10
2002	1,104	8	2	10
2003	1,182	10	1	11
2004	1,328	6	0	6
Total	11,313	112	18	130

To illustrate, column 2 of Table 1 shows the total number of observations on firms listed on the Toronto Stock Exchange (firm-year) that recur in our database, for dynamic analysis of probabilities of default. These are firms that did not default. The data are from DATASTREAM. We point out that DATASTREAM did not gather data on all firms in the early 1990s, which is why there are fewer observations. These data appear as reliable as those for later years, however, though they are less complete in their coverage of firms on the stock exchange. Figure 1 represents the evolution of defaults over time.

Figure 1: Distribution of defaults and reorganizations by year

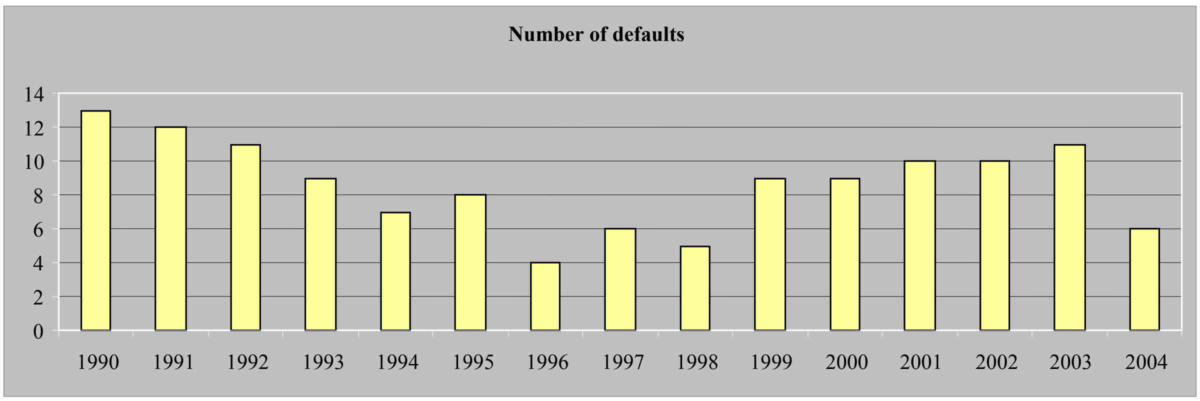

For the 130 firms that defaulted, we have 436 dynamic observations (firm-year) on accounting data, of which 378 are for firms that went bankrupt and 58 for those undergoing reorganization.

After merging the accounting data with the daily market data, 108 firms remained in the intermediary database of defaults, i.e., for the first stage of our study during which we compute the probabilities of default using the structural model. This attrition is mostly attributable to missing market data for some firms, and the fact that we had only one year of accounting data for others—rendering the data unusable for our study. In fact, to be able to apply the structural model, we require at least one year for the estimation and another year for computing the probabilities for each firm.

Two variables are vital for the first stage of calculating the probabilities of default with the structural model: market value and liabilities. We have 60,331 daily observations on market value and 69,822 observations on liabilities. This gap is attributable to missing market value data.

Table 2 presents statistics on market value and liabilities for 108 firms having defaulted, after cleaning the data and merging the accounting and market data.

Table 2: Descriptive statistics for the 108 defaults
(in millions of Canadian dollars)

Statistic	Market value	Liabilities
Mean	335.04	185.25
Median	14.57	15.33
Mode	22.00	0.32
Standard deviation	1,301.36	948.32
Skewness	6.00	10.60
Kurtosis	40.06	126.58
Number of observations	60,331	69,822

During the second phase of the study, i.e., during the probit regressions, only 57 of the defaulting firms remained in the final database. In several cases, the financial statements were insufficiently complete to create the variables required for a multivariate analysis, while others had not produced financial statements during the 18 months preceding the default. We believe that going back more than 18 months would not provide a representative picture of the firm's real situation at the time of default.

The data were processed, cleaned, and merged using SAS, version 9.1.

3.2 *Firms listed on the Toronto Stock Exchange that did not default*

The data on the firms listed on the Toronto Stock Exchange that did not default are from DATASTREAM's FTORO.LLT series. The frequency of market value (MV) data is daily for the period from 1 January 1988 to 31 December 2004. Accounting data are from Stock Guide. Annual accounting data for the new fiscal year are only used four months after publication, since this is the average delay before investors have access to this information.

In total, we have 3,109,201 daily observations for the market value variable. The mean, over all firms, is Can$854.93 million. The standard deviation is Can$4,758.10 million, owing to the existence of very high market capitalization values for some firms—the maximum being Can$366,399.75 million.

In Table 3, we present the descriptive statistics for the daily market value variable for all firms listed in Toronto that did not default.

Table 3: Descriptive statistics for all firms listed on the Toronto Stock Exchange that did not default (in millions of Canadian dollars)

Statistic	Market value
Mean	854.93
Median	59.03
Mode	285.38
Standard deviation	4,758.10
Skewness	26.12
Kurtosis	1,188.95
Range	366,400.00
Interquartile range	255.77
Number of observations	3,109,201

3.3 *Various statistics on the firms retained for the study*

To begin our estimations, we merged the accounting database from STOCKGUIDE with the market database from DATASTREAM. However, we removed all financial companies from the two databases, considering that they did not belong in the study since the structure of their financial statements differs from those of non-financial firms.

Our final database included 684 publicly traded non-financial Canadian firms, 627 of which did not default and 57 of which did.

After merging and cleaning the data, we were left with 1,885,707 daily market value observations. The mean over all firms is Can$747.12 million, while the median is Can$48 million. This large difference is attributable to the very high value of market capitalization in the case of some firms.

Compared to the entries in Tables 2 and 3, Table 4 provides some statistics for the 627 publicly traded non-financial Canadian firms that did not default.

We notice that, in terms of the mean, the median, and the standard deviation, there is little difference between the initial sample values in Table 3 and those in Table 4 for the final sample.

Table 4: Descriptive statistics for the firms retained for the analysis (in millions of Canadian dollars)

Statistic	Market cap
Mean	747.12
Median	47.99
Mode	2.09
Standard deviation	4,247.97
Skewness	38.79
Kurtosis	2,380.02
Range	366,400.00
Interquartile range	269.35
Number of observations	1,885,707
Number of firms	684

We also looked at the lags separating the default dates from the last financial statements of some firms. Many firms do not publish financial statements during their final years prior to bankruptcy. We felt obliged to withdraw from the database those for which these lags exceeded 18 months in duration. For the others, i.e., those that had defaulted between 12 and 18 months after their final financial statement, we moved the date of the default up to reconcile it with the last observable accounting year.

For example, in the case of a firm that defaulted in 2000 and whose fiscal year ended on 31 December we need accounting data from 1999 to estimate the probabilities of default in 2000. To be able to estimate models using accounting variables, the time elapsed between the date of the publication of the final financial statement and the date of default should not exceed 12 months. Unfortunately, this time is longer for many of the defaults in our database. Knowing that many defaulting firms do not publish financial statements during the year leading up to official bankruptcy or reorganization, we dropped this condition and included those whose defaults occurred from the thirteenth through the eighteenth month after publication of their final financial statement. This, in turn, forced us to move up the default dates to make them correspond to the last year for which we had valid accounting data. This explains why we have defaults for the years 1988 and 1989, despite the fact that defaults began in 1990 in our initial sample.

Table 5 provides an overview of the annual frequency of the data used in the final model for calculating the probabilities of default. We observe that eight defaults were moved up in 1988 and 1989.

Table 5: Annual frequency of the data used for the final model

Year	0: Non-default 1: Default	Total number	Percentage
1988	0 1	99 3	97.06% 2.94%
1989	0 1	111 5	95.69% 4.31%
1990	0 1	132 1	99.25% 0.75%
1991	0 1	149 2	98.68% 1.33%
1992	0 1	150 3	98.04% 1.96%
1993	0 1	177 3	98.33% 1.67%
1994	0 1	185 5	97.37% 2.63%
1995	0 1	229 2	99.13% 0.87%
1996	0 1	280 0	100.00% 0.00%
1997	0 1	328 5	98.50% 1.50%
1998	0 1	381 3	99.22% 0.78%
1999	0 1	427 3	99.30% 0.70%
2000	0 1	462 6	98.72% 1.28%
2001	0 1	505 8	98.44% 1.56%
2002	0 1	563 5	99.12% 0.88%
2003	0 1	588 3	99.49% 0.51%
2004	0 1	66 0	100.00% 0.00%
Total	0 1	4,766 57	99.49% 0.51%

4. Estimation of the probabilities of default with the structural model

The literature on structural models distinguishes between at least two techniques for estimating the parameters. One is based on plain vanilla (standard) options, and the other is based on barrier options. In the second model type, the payoff is a function of a threshold price for the underlying asset. In our case, we use a down-and-out barrier call option; i.e., the option vanishes when the underlying asset reaches the barrier.

We assume that the firm's capital structure consists exclusively of debt plus equity. The level of the debt is denoted B, while $(T-t)$ is the time to maturity. The firm's value is A_t and the value, at time *t*, of debt maturing at time T is $V(A,T,t)$. The value of equity at time *t* is $v(A,t)$. Consequently, the total value of the firm at time *t* is:
$A_t = V(A,T,t) + v(A,t)$.

To compute the probability of default for this firm, we assume that the value of its assets follows the following Brownian motion process:

$$dA = \mu_A A dt + \sigma_A A dz, \tag{1}$$

where μ_A is the mean of the value of the firm and σ_A its standard deviation.

Let $dz = \varepsilon\sqrt{dt}$ with $\varepsilon \sim N(0,1)$, i.e., the distribution of ε is normal with mean zero and unit standard deviation.

As to the liabilities, assume, on one hand, that the firm's liabilities L are the sum of short-term liabilities plus one-half of long-term liabilities. This assumption, which is used by Moody's KMV for North American firms, ensures that the firm's liabilities are not overstated. We perform sensitivity analyses on this assumption. On the other hand, we assume that L follows a deterministic process:

$$dL = \mu_L L dt. \tag{2}$$

Default occurs when the value of the firm's assets falls below that of its liabilities. The "barrier" is the default point $\breve{k}$. Consequently, we only need to monitor the ratio:

$$k = \frac{A}{L} \tag{3}$$

throughout the evolution of the firm's debt and check whether it reaches $\breve{k}$, the default point.

From equations (1), (2), and (3), we can compute the path of k:

$$dk = (\mu_k)\,kdt + \sigma_k kdz$$

with $\mu_k = (\mu_A - \mu_L)$ and $\sigma_k = \sigma_A$.

Estimation of μ_k and σ_k is based on the probability density function of k or, more specifically, the probability density function of $\left(\frac{k_T}{k_t}\right)$. The probability that the barrier will not be reached and that the value will be $\left(\frac{k_T}{k_t}\right)$ at time T, equals:

$$h\left[\ln\left(\frac{k_T}{k_t}\right)\right] = \frac{1}{\sqrt{2\pi\sigma_k^2(T-t)}}\left\{\exp\left[-\frac{\left(\ln\left(\frac{k_T}{k_t}\right)-\left(\mu_k-\frac{\sigma_k^2}{2}\right)(T-t)\right)^2}{2\sigma_k^2(T-t)}\right]\right.$$
$$\left.-\exp\left[-\frac{2\ln\left(\frac{\breve{k}}{k_t}\right)\left(\mu_k-\frac{\sigma_k^2}{2}\right)}{\sigma_k^2}-\frac{\left(\ln\left(\frac{k_T}{k_t}\right)-2\ln\left(\frac{\breve{k}}{k_t}\right)-\left(\mu_k-\frac{\sigma_k^2}{2}\right)(T-t)\right)^2}{2\sigma_k^2(T-t)}\right]\right\}.$$

This equation will be used to construct the likelihood function that we will maximize to estimate μ_k and σ_k for a given capital structure. The Bank of England set $\breve{k}$ equal to 1. We shall adopt this normalization.

The conditional probability, given that the firm has not defaulted by time T, is:

$$PD = 1 - \left[\left[1-\Phi(u_1)\right] - \varpi\left[1-\Phi(u_2)\right]\right],$$

where

$$u_1 = \frac{\breve{K} - \left(\mu_k - \frac{\sigma_k^2}{2}\right)(T-t)}{\sigma_k \sqrt{T-t}}, u_2 = \frac{-\breve{K} - \left(\mu_k - \frac{\sigma_k^2}{2}\right)(T-t)}{\sigma_k \sqrt{T-t}}, \varpi = \exp\left[\frac{2\breve{K}\left(\mu_k - \frac{\sigma_k^2}{2}\right)}{\sigma_k^2}\right],$$

and Φ is the cumulative density function of the normal distribution.

Let $\breve{K} = \ln\left(\breve{k}/k_t\right)$. In the case of a European call option, the probability of default equals $\Phi(u_1)$. However, for a barrier option, we see that the term $\varpi\left[1-\Phi(u_2)\right]$ adjusts the probability of default to account for the fact that the firm may default before the horizon date *T*.

For the estimation, we use $y = \frac{MV}{L}$, i.e., the ratio of market value (MV) to debt L, as a proxy for the ratio $k = \frac{A}{L}$. We use Matlab to estimate μ_k and σ_k with the maximum-likelihood method. Subsequently, we compute the probabilities of default. The parameters μ_k and σ_k are estimated daily, on the basis of a 261-day window for firms having defaulted and a 261×2, or 522-day window for those that did not. We needed to proceed thus in order not to lose too many observations for defaulting firms. A sensitivity analysis that consisted of imposing the same windows on both types of firms did not alter the results. The mean computed probability for firms that did not default is 8.08 per cent with the 522-day window and 8.8 per cent with the 261-day window. We pursued our work with the longer window to have more information and thus greater statistical reliability, since the default probabilities of several firms are very low. In our study, we set the default barrier at 1 ($y = 1$), i.e., the firm defaults when its market value (MV) equals its debt (L), which, in turn, initially equals short-term liabilities plus 50 per cent of long-term liabilities.

Table 6 reproduces the probabilities of default, computed one year prior to the period of risk exposure, for firms that did, and did not, default. The mean of the probabilities for defaulting firms is 35.46 per cent, while that for non-defaulting firms is 8.08 per cent. These annual means are 35.36 per cent and 7.51 per cent, respectively, when the percentage of long-term debt is 25 per cent, and 35.44 per cent and 8.63 per cent when it is 75 per cent.

Table 6: The probability of default for all firms, computed one year prior to risk exposure

Year	Probability of default for firms that did not default	Probability of default for firms that did default
1988	11.28%	40.05%
1989	10.65%	46.62%
1990	12.82%	44.93%
1991	11.36%	53.53%
1992	4.76%	36.09%
1993	3.15%	31.56%
1994	3.07%	29.78%
1995	7.09%	48.96%
1996	3.72%	33.91%
1997	5.03%	24.24%
1998	10.17%	30.27%
1999	12.29%	34.29%
2000	7.73%	20.56%
2001	13.71%	35.69%
2002	10.14%	41.07%
2003	7.29%	15.94%
2004	3.12%	39.12%
Mean	8.08%	35.68%
Number of firms	627	57

The next two figures show the evolution of the probabilities of default over time for several firms. The ones in Figure 2 defaulted, while those in Figure 3 did not.

Figure 2: Daily default probabilities (1 year) of defaulting firms

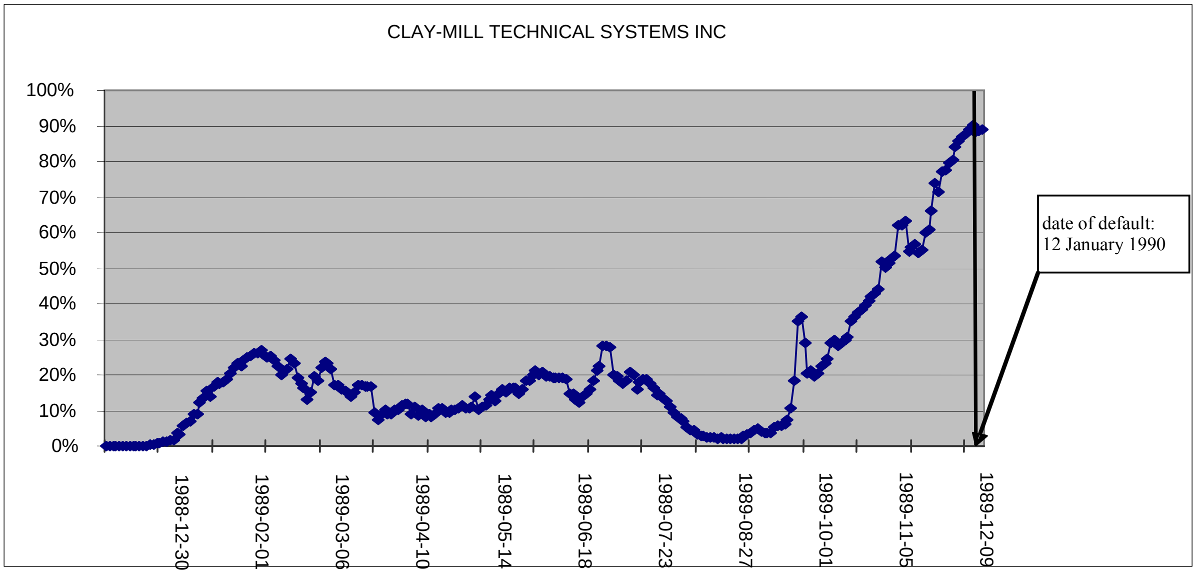

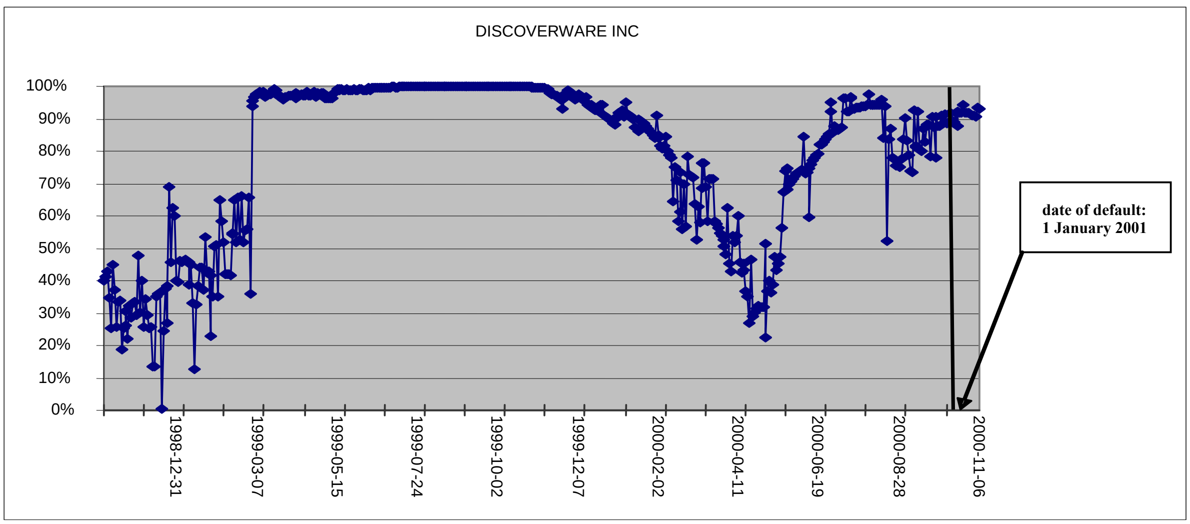

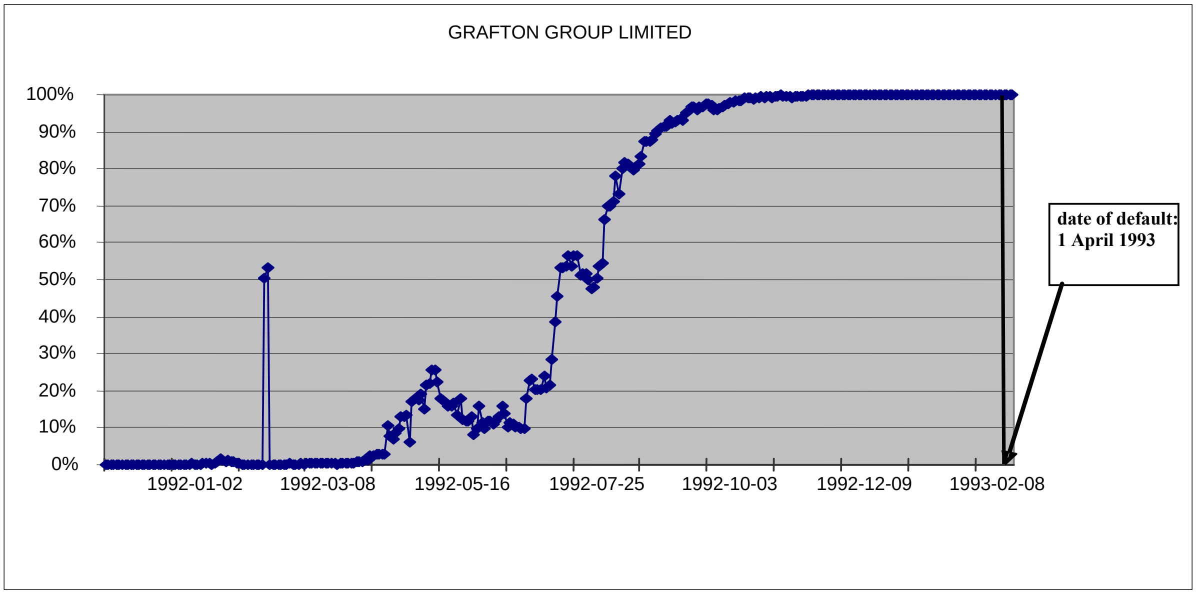

The results in Figure 2 are consistent with the model, while those in Figure 3 are somewhat surprising. They represent two extreme cases and reflect our comments on the overstatement of the probabilities of default in the structural model, especially when there are shocks to the market. It is of interest to note, for example, that the value of Asbestos stock held between $15 and $20 during the period from early 2000 to mid-2001, then fell to below $5 from the middle of 2001 until the beginning of 2004, then rose again. The model appears very sensitive to significant fluctuations in the values of this firm's stocks.

The explanation is different in the case of Agricore. This firm resulted from a merger in 2001. In 2002, the firm announced a change in its fiscal year, and consequently its financial statements covered a 15-month period. When computing the probabilities of default with the structural model, we use the ratio market value/liabilities from the annual reports. In the case of Agricore, liabilities covered 15 months in 2002, resulting in an increase in the default barrier (short-term liabilities plus 50 per cent of long-term liabilities) in the calculation of the probabilities of default for the following period. Consequently, the probabilities of default are very high in the years following 2002, as we see in Figure 3. The liabilities must be adjusted here to yield a more accurate picture of the firm's true situation. This adjustment is far from straightforward, since it reflects not only a change of period but also a merger, and the required data are not available. This illustrates the structural model's extreme sensitivity to variations in the inputs, providing the rationale for using the hybrid model, which contains more information for conditioning the estimates of the probabilities of default.

As Figure 8 reveals very clearly, variations in the conditional probabilities of the hybrid model for Asbestos and Agricore Corporation are much more modest, though they remain sensitive to fluctuations in the inputs operating over variations in the probability of default. Three smoother examples are featured in Figure 4. It is of interest to note that the probabilities of default of the three firms move in tandem with Moody's default cycles.

Figure 3: Daily default probabilities (1 year) of non-defaulting firms

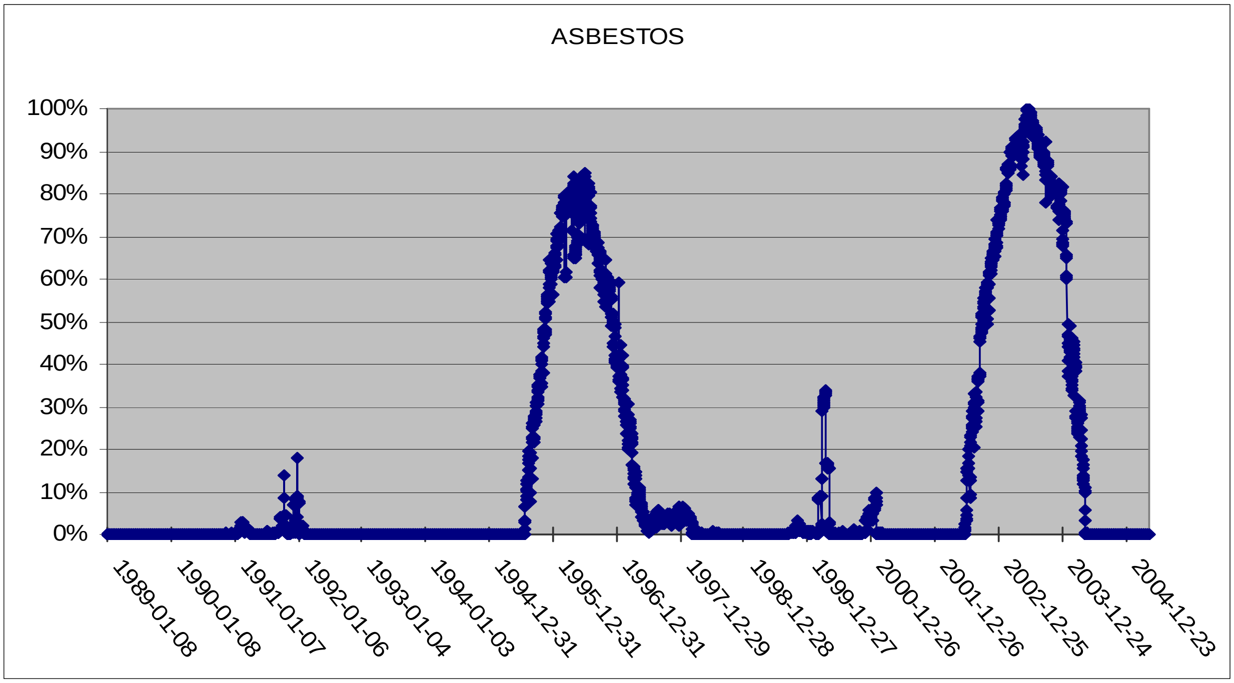

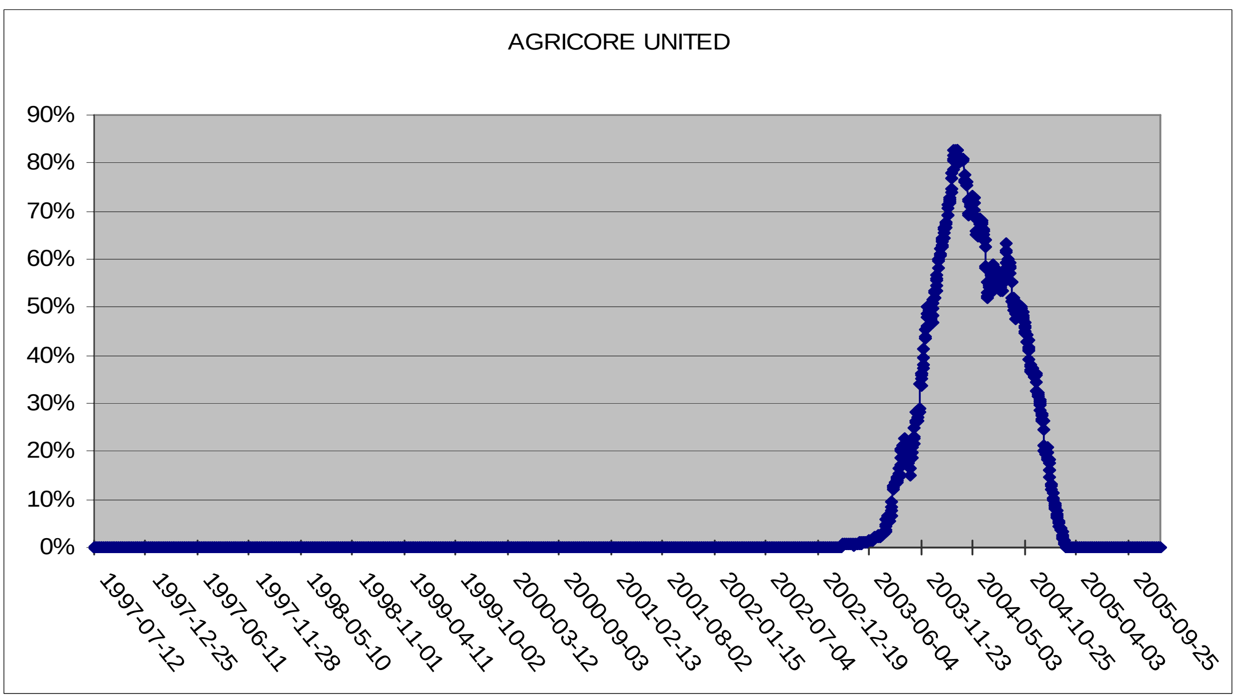

Figure 4: Other PDs (1 year) of non-defaulting firms

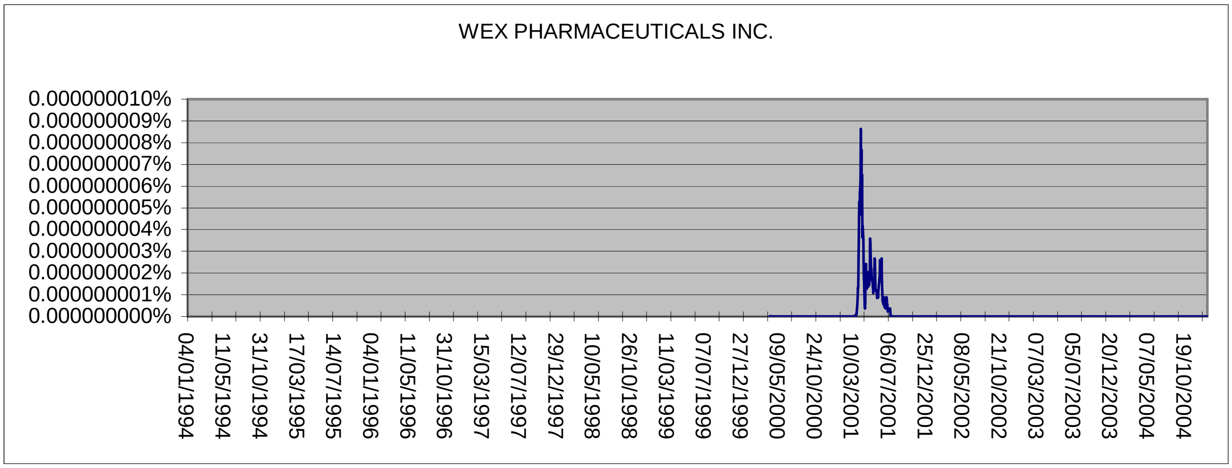

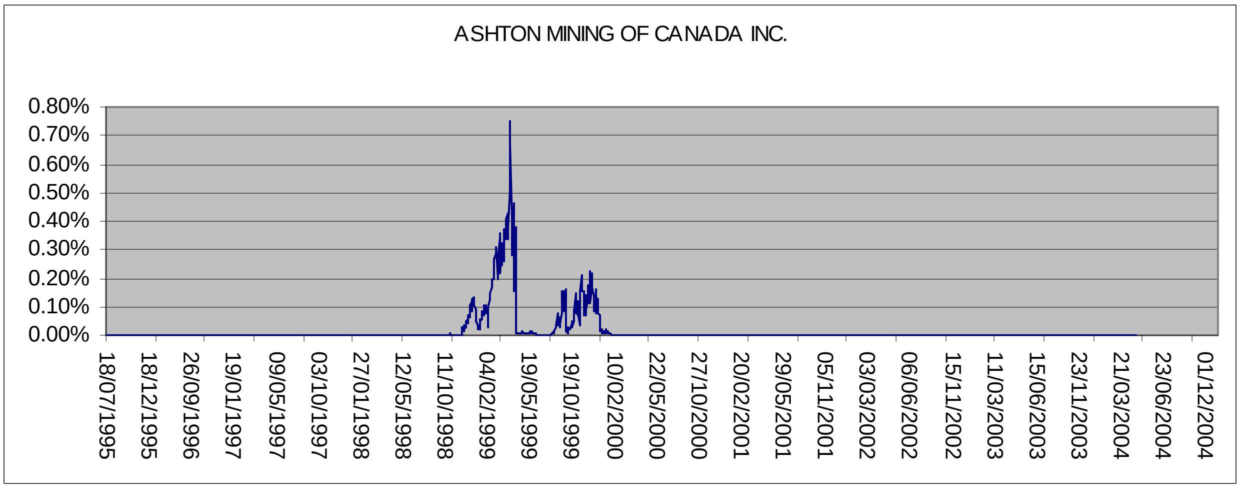

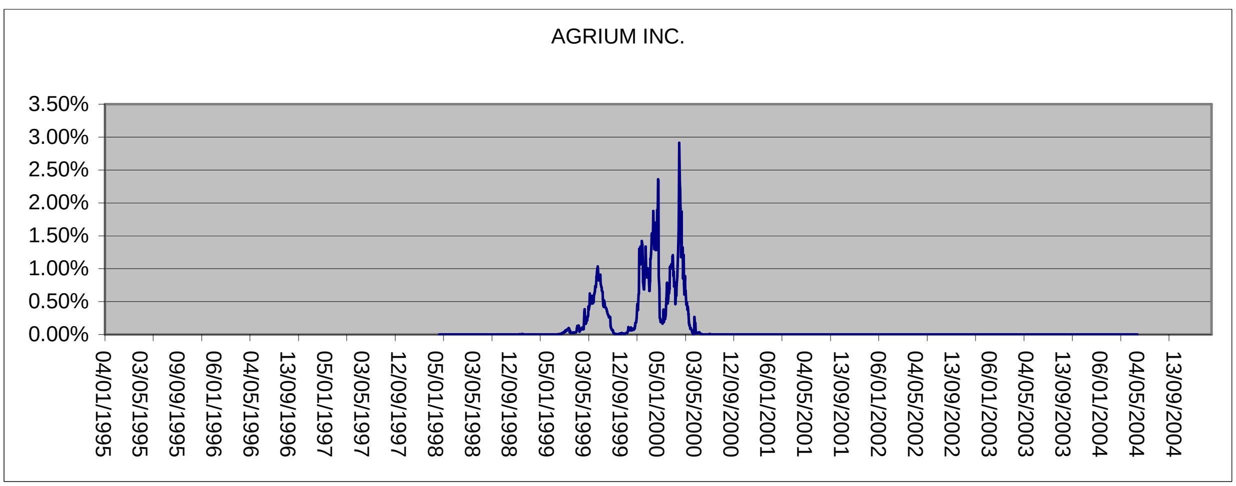

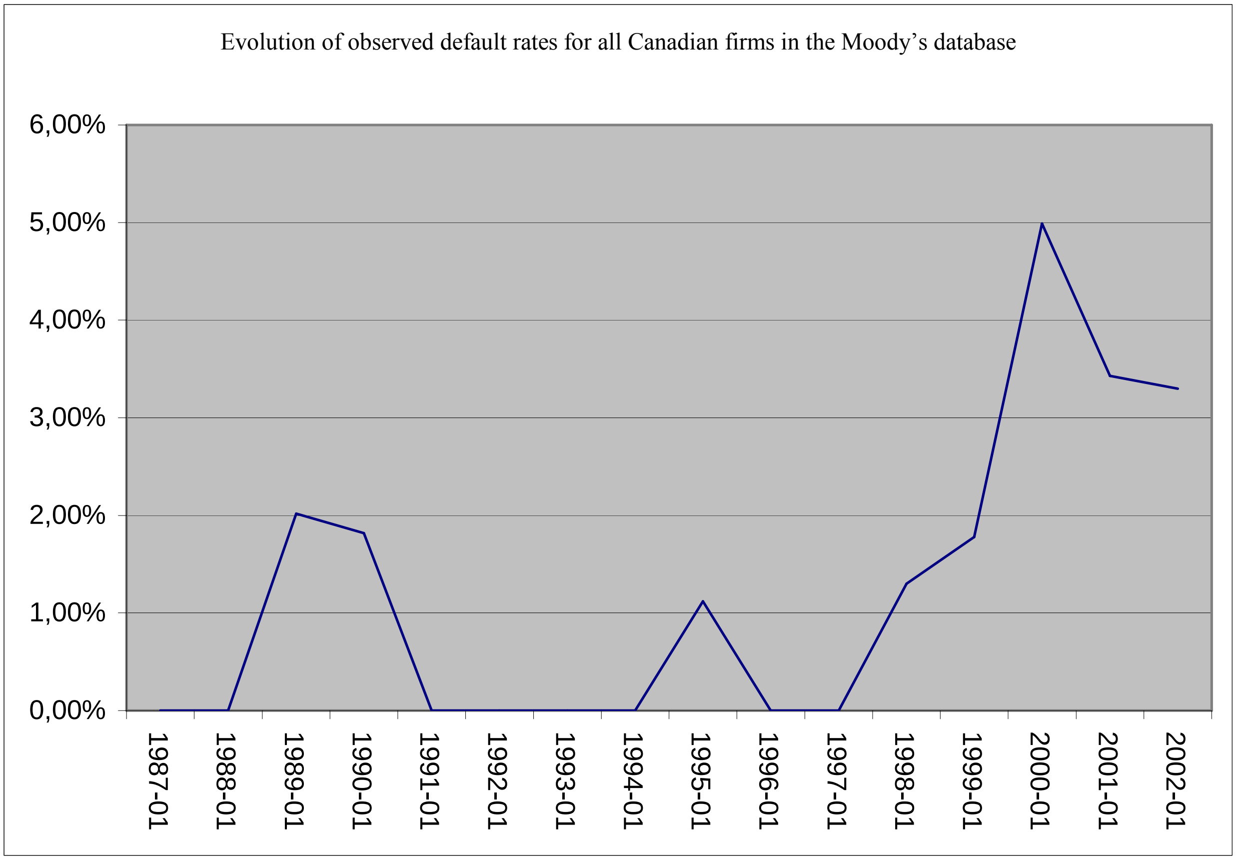

5. Estimation of the hybrid model

5.1 *Methodology*

We did not estimate the model with a simple linear regression, since we know that it must reflect non-linear behaviour of the explanatory variables for defaults. In addition, it is well documented that simple linear models are inappropriate when the dependent variable is a probability.

5.1.1 The probit model with cross-section data

In the probit model, the dependent variable y_i is a dichotomous variable assuming the value 1 if an event occurs, and 0 otherwise. In our case, the variable y_i assumes the following values:

$y_i = 1$ if firm *i* defaults, and
$y_i = 0$ otherwise.

The vector of explanatory variables (financial ratios and other financial or business cycle variables) for firm *i* is denoted x_i, and β represents the vector of parameters to be estimated. The premise of the probit model is that there exists a qualitative response variable (y_i^*) that is defined by the following relationship:

$$y_i^* = \beta' x_i + \varepsilon_i . \tag{4}$$

In practice, however, y_i^* is an unobservable latent variable. The variable we actually observe is the dichotomous y_i, such that:

$$\begin{aligned} y_i &= 1 \text{ if } y_i^* > 0; \\ y_i &= 0 \text{ otherwise.} \end{aligned} \tag{5}$$

In this formulation, $\beta' x_i$ is not $E\left(y_i / x_i\right)$, as in the simple linear model, but rather $E\left(y_i^* / x_i\right)$. From equations (4) and (5), we have:

$$\text{Prob}\,(y_i = 1) = \text{Prob}\,(\varepsilon_i > -\beta' x_i) = 1 - F(-\beta' x_i), \tag{6}$$

where F is the cumulative distribution function of ε_i.

The observed values of y are simply realizations of a binomial process with probabilities that are given by (6) and vary from one observation to the next (with x_i). The likelihood function can thus be written:

$$\ell = \prod_{y_i=0} F(-\beta' x_i) \prod_{y_i=1} \left(1 - F(-\beta' x_i)\right), \tag{7}$$

and the parameter estimates β are those that maximize ℓ.

The functional form of F in equation (7) depends on the retained assumptions regarding the distribution of the residual errors (ε_i) in equation (4). The probit model is based on the assumption that these errors are independently and identically distributed (i.i.d.) and follow a standard normal distribution N(0,1). The functional form can thus be written:

$$F(-\beta' x_i) = \int_{-\infty}^{-\beta' x_i} \frac{1}{(2\pi)^{1/2}} \exp\left[-\frac{t^2}{2}\right] dt . \tag{8}$$

5.1.2 The probit model with panel data

This model accounts for potential correlations between different observations on the same firm at different points in time (different financial statements). It is defined by the following regression:

$$y_{it}^{*} = \beta' x_{it} + \varepsilon_{it} \tag{9}$$

and the observed dichotomous variable is such that:

$$\begin{aligned} y_{it} &= 1 \text{ if } y_{it}^{*} > 0; \\ y_{it} &= 0 \text{ otherwise,} \end{aligned} \tag{10}$$

where *i* represents the firm and *t* the time of firm *i*'s financial statement. To account for intertemporal correlation using a random-effects model, the error must be decomposed into two terms:

$$\varepsilon_{it} = v_{it} + \vartheta_i ,$$

where $v_{it} \sim N(0,1)$ is the stochastic-error component and $\vartheta_i \sim N(0, \sigma_u^2)$ is the part of the error correlated with *i*, so that the two error components (v_{it} and ϑ_i) are normally distributed with mean 0 and are independent of each other. The variance of the error term ε_{it} can now be represented by:

$$\text{var}(\varepsilon_{it}) = \sigma_v^2 + \sigma_\vartheta^2 = 1 + \sigma_\vartheta^2$$

and the correlation is:

$$\text{corr}(\varepsilon_{it}, \varepsilon_{is}) = \rho = \frac{\sigma_\vartheta^2}{1 + \sigma_\vartheta^2}.$$

Thus, the new free parameter is:

$$\sigma_\vartheta^2 = \frac{\rho}{1 - \rho}.$$

This is the parameter that will make it possible to measure the existence of a correlation between the different observations (financial statements) of a single company over time.

5.2 *Data and variable selection*

The principal objective of this study is to verify whether combining the reduced-form and the structural model into a hybrid model yields a better measure of the default risk than is obtained from traditional econometric reduced-form and structural models estimated separately. To accomplish this, we seek to explain defaults by estimating a probit model in which the explanatory variables are the estimated probabilities of default from the structural model, financial ratios, and other accounting and cyclical data. The dependent variable is a binary variable assuming the value of 1 if the firm defaults and 0 otherwise. Using this same methodology, we also estimate a model with only accounting data as explanatory variables (reduced-form model) and a third probit model in which the only exogenous variable is the probability of default (PD) from the structural model (the model that contains only structural information).

Thus, we test the predictive power of the PD variable for explaining corporate bankruptcy by including it in the reduced-form model. If, after controlling for the effect of the firm's accounting data, we find that the estimated coefficient of the PD variable is statistically different from zero, the probabilities of default yielded by the structural approach will be shown to contain information that is supplementary to that in the accounting data, and we will be able to use its coefficient to update the probability of default when the PD changes.

As to the selection of accounting variables and financial ratios used in the reduced-form and hybrid models, we retained a wide array of variables and financial ratios liable to have an impact on the quality of the firm's credit and for which we were able to obtain satisfactory data. This choice of variables was based on both empirical studies addressing the determinants of default in Canadian privately owned or publicly traded firms (Petrescu, 2005; Beaulieu, 2002; RiskCalc; Z-score) and on studies conducted in other countries (Bank of England).

Since missing accounting items are quite frequent in the defaults database, we were faced with a dilemma: retain more accounting variables, and thus reduce the number of defaults in our database and, by extension, the statistical significance of our results; or eliminate variables at the cost of undermining the model specification. We found a compromise that seemed best to us.

To make a sound selection, we started by estimating the probit model on each accounting variable separately. This allowed us to retain the most significant ones and thus reduce the number of missing observations in our final estimation. Starting from an original sample with 5,556 observations including 60 defaults, corresponding to observations for which were able to obtain the PD on an annual basis, we arrived at a sample of 4,889 observations with 57 defaults. Descriptive statistics for the retained variables are presented in Table 7. The analysis covers the period 1988 to 2004.

Table 7: Descriptive statistics for the explanatory variables

Variable	Mean	Median	Standard deviation	Minimum	Maximum
Mean annual PD (1 year)	0.09	0	0.21	0	1
Cash	67.94	3.18	323.29	0	9,390
Short-term assets	361.68	37.43	1,364	0	36,811
Short-term liabilities	261.63	21.32	1,006	0	23,330
Retained earnings	53.71	2.58	3,873	-156,950	15,426
Net value	490.55	42.13	4,222	-1,726	163,016
Total liabilities	792.94	41.81	2,875	0	53,466
Total assets	1,361	92.39	4,688	0.02	93,931
Profitability < 0[(1)]	0.25	0	0.43	0	1
0 % < profitability < 6%[(2)]	0.12	0	0.32	0	1
Net value/total liabilities	6.76	0.88	35.21	-0.95	1,214
Retained earnings/total liabilities	-3.3	0.09	26.74	-1,052	72.1
Total liabilities/total assets	0.67	0.49	9.43	0	658.1
Cash/total liabilities	0.73	0.08	3.01	0	68
Cash/total assets	0.10	0.04	0.16	0	1.92
EBITDA/short-term liabilities	0.11	0.41	4	-127.66	55.71
GDP growth	0.03	0.03	0.02	-0.02	0.06

The accounting variables are in millions of Canadian dollars.

(1) A dummy variable assuming the value 1 if the margin of profit (EBITDA/sales) is negative, 0 otherwise.

(2) A dummy variable assuming the value 1 if the margin of profit (EBITDA/sales) is between 0 and 6%, 0 otherwise.

5.3 *Analysis of the results*

5.3.1 Estimation of the probit model with different specifications

In this section, we analyze the characteristics and performances of three models: the hybrid model, the reduced-form model, and the model containing only structural information. We summarize the results of these estimations in Table 8.

In Model 1, we only use the information from the structural model, which is tantamount to treating the annual mean of the structural PD as an explanatory variable. Notice that the probabilities of default used here are computed from data for the year prior to the estimation year. The coefficient of PD is 1.07 per cent, and it has the expected sign. It is a very significant predictor of the probabilities of default, with a *p*-value of less than 1 per cent. However, the corrected pseudo-R^2 is low (6.3 per cent).

In Model 2, we estimate the reduced-form model on all of the retained accounting variables and financial ratios, adding the rate of growth of the Canadian GDP to account for the business cycle. Examination of the estimated coefficients reveals that these variables have the expected sign, aside from total liabilities, total assets, and their ratio. For example, the dummy variable for negative profitability has a positive coefficient, indicating that a negative profitability increases the probability of default compared to profitability exceeding 6 per cent. We also notice that 10 of the 14 accounting variables considered are significant at the 95 per cent confidence level, and that the effect of GDP growth on the probability of default is not statistically different from zero at any of the usual confidence levels. Moreover, examination of Model 2 reveals that the reduced-form specification largely outperforms the one using only information from the structural model (Model 1) in terms of its ability to explain corporate bankruptcy. The likelihood ratio is 246.98 for the reduced-form model, versus 36.93 for the structural model with only PD as an exogenous variable (the corresponding values of R^2 are 20 per cent and 6.3 per cent). At first glance, the PDs from the structural approach appear unable to generate adequate predictions of defaults of publicly traded Canadian firms, compared with accounting data.

In Model 3, we estimate the hybrid model by adding the probabilities of default computed from the structural model to the explanatory variables from Model 2. An analysis of the results reveals that the PD variable is statistically significant at the 5 per cent confidence level. This suggests that the probabilities of default from the structural approach have an additional predictive power for corporate defaults than the firms' financial statements. In addition, we observe that the estimated parameters from our reduced-form model are robust to the introduction of probabilities of default through the structural approach. We see neither major changes to the estimated coefficients of the accounting variables nor any loss of significance for some of these variables. Furthermore, to acquire a better understanding of the contribution of Model 3 relative to that of Model 2, we test the null hypothesis from Model 2 against the alternative hypothesis from Model 3. To accomplish this, we compute the likelihood ratio (LR) as follows:

$$LR = -2\,[\log(L2) - \log(L3)] = -2\,(-186.9488 + 185.0268) = 3.844,$$

where log(L2) and log(L3) indicate the log likelihood of Models 2 and 3, respectively. The distribution of the resulting statistic (LR) is chi-square with 1 degree of freedom. The critical value of this distribution at the 95 per cent confidence level is 3.8414. Thus, we can reject the Model 2 specification in favour of that of Model 3 at the 95 per cent confidence level.

We now repeat this analysis in Models 4 and 5, but this time only retaining the significant variables from Models 2 and 3. In Model 5, the PD variable is significant at the 10 per cent confidence level, confirming the information contributed by the probabilities of default from the structural approach. Here again, the explanatory power of the model, as measured by the adjusted R^2 and the likelihood ratio, increases only marginally with the introduction of the PD. The likelihood ratio rises from 240.81 to 243.65. The likelihood-ratio test allows us to reject Model 4 in favour of Model 5 at the 90 per cent confidence level. In fact, the value of the LR test is 2.83, compared to a critical value of 2.79 for the chi-square distribution function at the 90 per cent confidence level.

We proceed with a further test to assess the informational contribution of the PD variable, using a specification based solely on the significant accounting values. When we use SAS 9.1 to perform a stepwise selection of variables to retain in the model, we observe that the PD variable is always kept along with the same four accounting variables from Models 6 and 7. Thus, the accounting variables that are most relevant for predicting defaults in publicly traded Canadian firms are: profit margins, or, more specifically, the dummy variables for profit margins that are negative or less than 6 per cent, the ratio net value/total liabilities, and the ratio retained earnings/total liabilities. With this latter specification, the PD is significant at the 1 per cent confidence level. Finally, when we re-ran the regressions using a panel probit model with random effects, the preliminary results remained essentially unchanged compared to those from the simple probit estimation. Details are available.

Also, we re-estimated Model 5 with dummy variables for each year of the observations, with 1988 serving as the reference year. The only temporal variable that proved significant corresponded to the year 1999 (significant at the 90 per cent confidence level). However, when the macroeconomic variable for GDP growth was included, none of the year dummies remained significant.

Table 8: Analysis of the maximum-likelihood estimators

The estimated coefficients are on top, and the corresponding *p*-values below.

Parameters	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
Constant	-2.4419 <0.0001	-1.9541 <0.0001	-2.0344 <0.0001	-1.9663 <0.0001	-2.0411 <0.0001	-2.2734 <0.0001	-2.3473 <0.0001
Annual mean PD (1 year)	1.0694 <0.0001		0.4568 0.0446		0.3791 0.0863		0.5432 0.0080
Cash		0.00752 0.0113	0.00746 0.0104	0.00586 0.0824	0.00584 0.0580		

Parameters	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
Short-term assets		-0.00919 0.0002	-0.00912 0.0002	-0.00857 0.0004	-0.00836 0.0003		
Short-term liabilities		0.00346 0.0030	0.00345 0.0029	0.00308 0.0069	0.00299 0.0073		
Retained earnings		-0.00400 0.0027	-0.00379 0.0036	-0.00434 0.0012	-0.00416 0.0014		
Net value		-0.00640 0.0013	-0.00613 0.0014	-0.00668 0.0009	-0.00647 0.0008		
Total liabilities		-0.00317 0.0083	-0.00301 0.0096	-0.00333 0.0074	-0.00325 0.0060		
Total assets		0.00372 0.0035	0.00356 0.0038	0.00395 0.0023	0.00386 0.0018		
Profitability < 0		0.9915 <0.0001	0.9619 <0.0001	0.9603 <0.0001	0.9113 <0.0001	1.1999 <0.0001	1.1012 <0.0001
0 % < Profitability < 6 %		0.4098 0.0460	0.4088 0.0500	0.4079 0.0411	0.3996 0.0476	0.4510 0.0097	0.4261 0.0166
Net value/total liabilities		-0.5401 <0.0001	-0.5353 <0.0001	-0.5462 <0.0001	-0.5249 <0.0001	-0.6444 <0.0001	-0.6174 <0.0001
Retained earnings/total liabilities		-0.0707 0.0026	-0.0708 0.0022	-0.0680 0.0037	-0.0669 0.0037	-0.0909 0.0003	-0.0885 0.0003
Total liabilities/total assets		-0.0345 0.5903	-0.0588 0.3842				
Cash/total liabilities		-0.2422 0.8399	-0.0497 0.9660				
Cash/total assets		-1.6476 0.4245	-2.0148 0.3223				
EBITDA/short-term liabilities		-0.0146 0.8130	-0.00553 0.9332				
GDP growth		2.9519 0.4346	3.3481 0.3773				
Number of observations	4,889	4,889	4,889	4,889	4,889	4,889	4,889
Number of defaults	57	57	57	57	57	57	57

Parameters	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6	Model 7
AUC	0.740	0.946	0.950	0.944	0.942	0.905	0.920
Likelihood ratio	36.9380 <0.0001	246.9282 <0.0001	250.7721 <0.0001	240.8134 <0.0001	243.6483 <0.0001	183.2470 <0.0001	189.9129 <0.0001
Log likelihood	-292.0380	-186.9488	-185.0268	-190.0061	-188.5887	-218.7893	-215.4564
McFadden's Pseudo R^2	0.063	0.40	0.40	0.39	0.39	0.30	0.31

* The likelihood ratio measures the explanatory power of the independent variables when the model is compared to a model with only a constant.

5.3.2 Various tests

In this section, we assess the capacity of the retained models to adequately rank defaulting and surviving firms. We do this using the accuracy ratios of the different models estimated, ROC curves, and gain charts. Sobehart, Keenan, and Stein (2000) provide a detailed description of methodologies for validating quantitative credit-risk models.

In Table 9 and Figure 5, we reproduce the accuracy ratios of the seven models estimated so far. The accuracy ratio of the model containing only structural information is 48 per cent. This ratio is maximized at 90 per cent for the hybrid model with all variables (Model 3). The same model, but without the probabilities of default from the structural approach, comes in at 89.2 per cent. This confirms the results from the preceding section. In fact, though the proportion of accurate rankings by the model with only structural information is 48 per cent of a perfect ranking, moving from the accounting model (Model 2) to the hybrid model (Model 3) only translates into a negligible improvement in the accuracy ratio.

Figure 5: Accuracy ratio

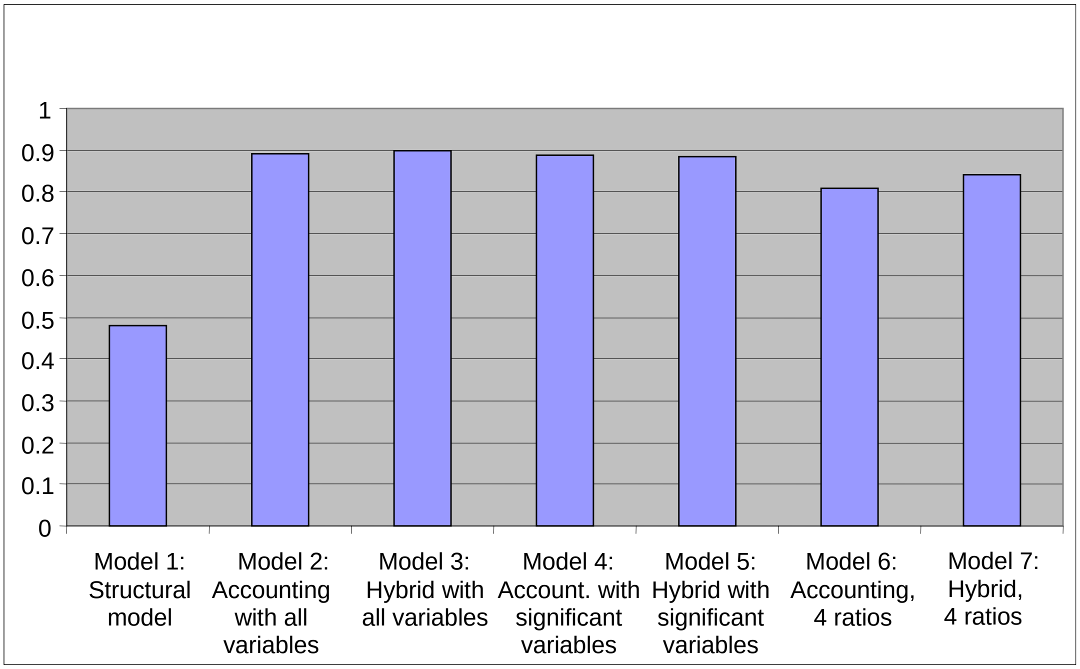

To better understand the performance of the probit model, we point out that the accuracy ratios from similar studies, to wit those by Sobehart, Keenan, and Stein (2000) of Moody's KMV and Tudela and Young (2003) of the Bank of England, are 76.7 per cent and 69 per cent, respectively, when only information from the structural model is included (versus 48 per cent in our study) and 77.09 per cent and 76 per cent for the hybrid model (versus 90 per cent for our model). Thus, it is clear that, despite the poor performance of our model with structural information, in comparison to those mentioned above, the hybrid model is able to correctly predict a greater proportion of the defaults in our sample.

Table 9: Accuracy ratio

Model	Accuracy ratio
Model 1: with only structural information	48.00%
Model 2: accounting, with all variables	89.20%
Model 3: hybrid, with all variables	90.00%
Model 4: accounting, with significant variables	88.80%
Model 5: hybrid, with significant variables	88.40%
Model 6: accounting, 4 ratios	81.00%
Model 7: hybrid, 4 ratios	84.00%

For the ROC curve and gain-chart analysis of model performance, the model with structural information is Model 1. As to the rest, we retain the models with significant variables, to wit Model 4 for the reduced-form model and Model 5 for the hybrid model, so as to avoid errors in the predicted probabilities of default calculated from non-significant coefficients.

Table 10 presents an analysis of Type I and II errors for the three retained models. We use the predicted probabilities of default computed from the probit coefficient estimates. Then, we rank the observations according to these probability predictions. If their values exceed a certain threshold (in this case, the mean of the probabilities), the firm is considered to be in default. Conversely, if the predicted probability is below the threshold, the observation is considered not to be in default. Table 10 juxtaposes this ranking with the actual occurrence of defaults. For the chosen thresholds, a comparison of the performances of the three models reveals that the Type I and II errors for the hybrid and accounting models are identical and relatively small compared to those from the model with only structural information. Once again, we see the superiority of the accounting model compared to the structural model for predicting corporate bankruptcy.

Table 10: Performance in predicting defaults

Model prediction	Actual defaults	Actual non-defaults
Hybrid model (Model 5) *Threshold: Mean of the predicted defaults = 11.46%*		
Defaults	53 92.98%	763 15.79%**
Non-defaults	4 7.02%*	4,069 84.21%
Total	57	4,832
Accounting model (Model 4) *Threshold: Mean of the predicted defaults = 11.45%*		
Defaults	53 92.98%	763 15.79%**
Non-defaults	4 7.02%*	4,069 84.21%
Total	57	4,832
Information from the structural model only (Model 1) *Threshold: Mean of the predicted defaults = 11.27%*		
Defaults	30 52.63%	728 15.07%**
Non-defaults	25 47.37%*	4,104 84.93%
Total	57	4,832

* Type I error in percentage; ** Type II error in percentage.

The downside of the preceding analysis is that the choice of threshold is arbitrary, and the percentages of Type I and II errors depend upon this choice. The ROC curve provides a correction for this limitation. This curve compares the proportion of defaults that were correctly predicted (top, left-hand cell in Table 10) with the proportion of firms that were incorrectly predicted as having defaulted (the "false alarms," top right-hand cell in Table 10) for all thresholds of the ranking. In Figure 6, we reproduce the ROC curves for Models 1, 4, and 5. This graph clearly shows that the reduced-form accounting model dominates the model with only structural information. Despite performing adequately in terms of predicting the probabilities of default, the PD variable does not make a material contribution to the performance of the hybrid model. However, it should be noted that the structural information is not only used to improve the performance of calculations of the probabilities of default. It is also used for quarterly, or even weekly, updates to the probabilities of default. We shall return to this point below. A brief description of the ROC function is presented in the appendix.

Figure 6: Receiver operating characteristic function

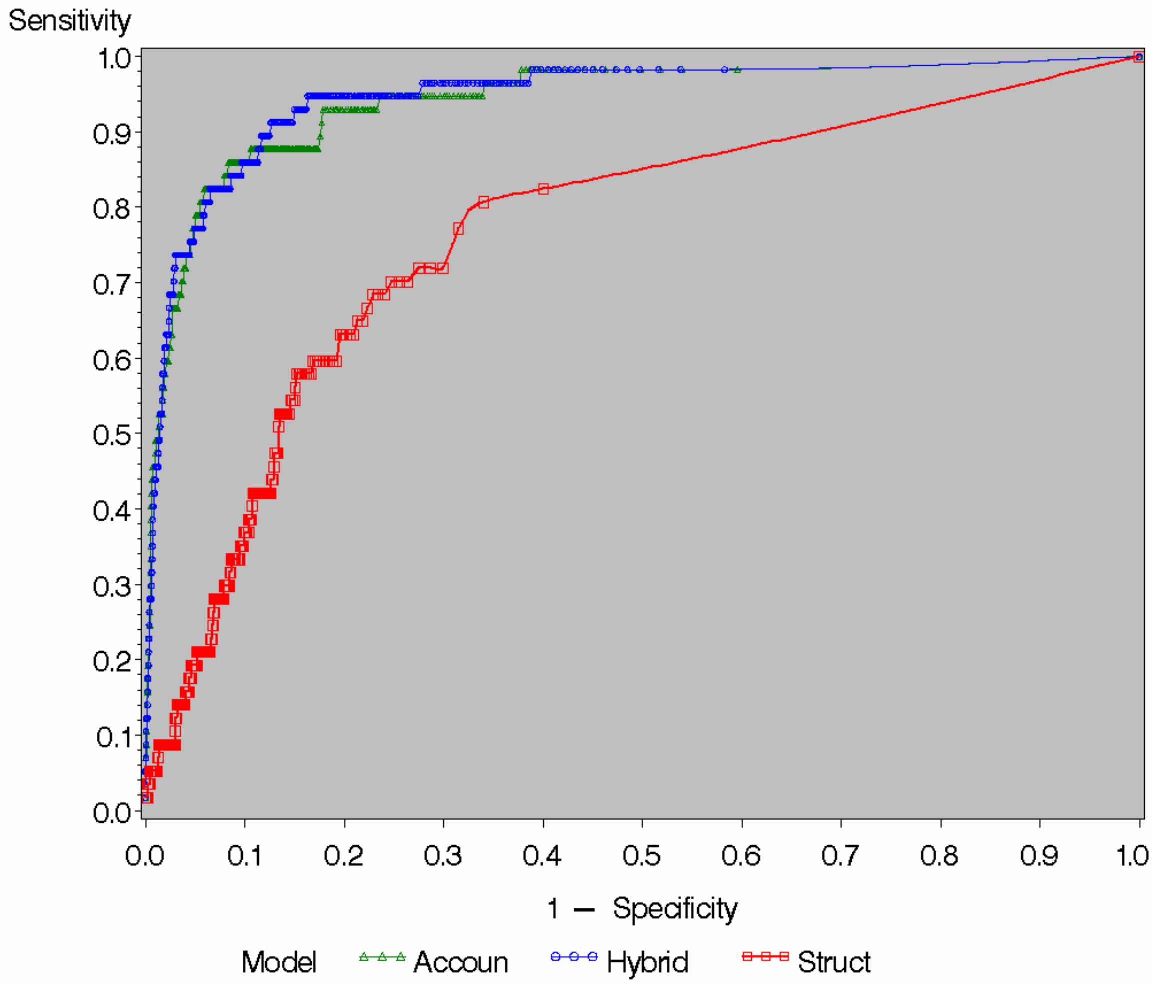

In the case of the gain chart, we proceed differently. We begin by ranking observations in decreasing order of their predicted probability of default. The ranking percentiles are on the abscissa and the percentage of defaults captured by the model is on the ordinate. A purely random ranking model would yield a 45° line from the origin, while a perfect model would capture every default within the first percentile. The gain chart for our hybrid model illustrates its capacity to detect a large percentage of corporate bankruptcies. In fact, for the hybrid model, Figure 7 reveals that 95 per cent of bankruptcies occur in the bottom 20 percentiles of predicted probabilities of default. The hybrid model does not owe this excellent performance to the structural model, but rather to the accounting model.

Figure 7: Gain chart

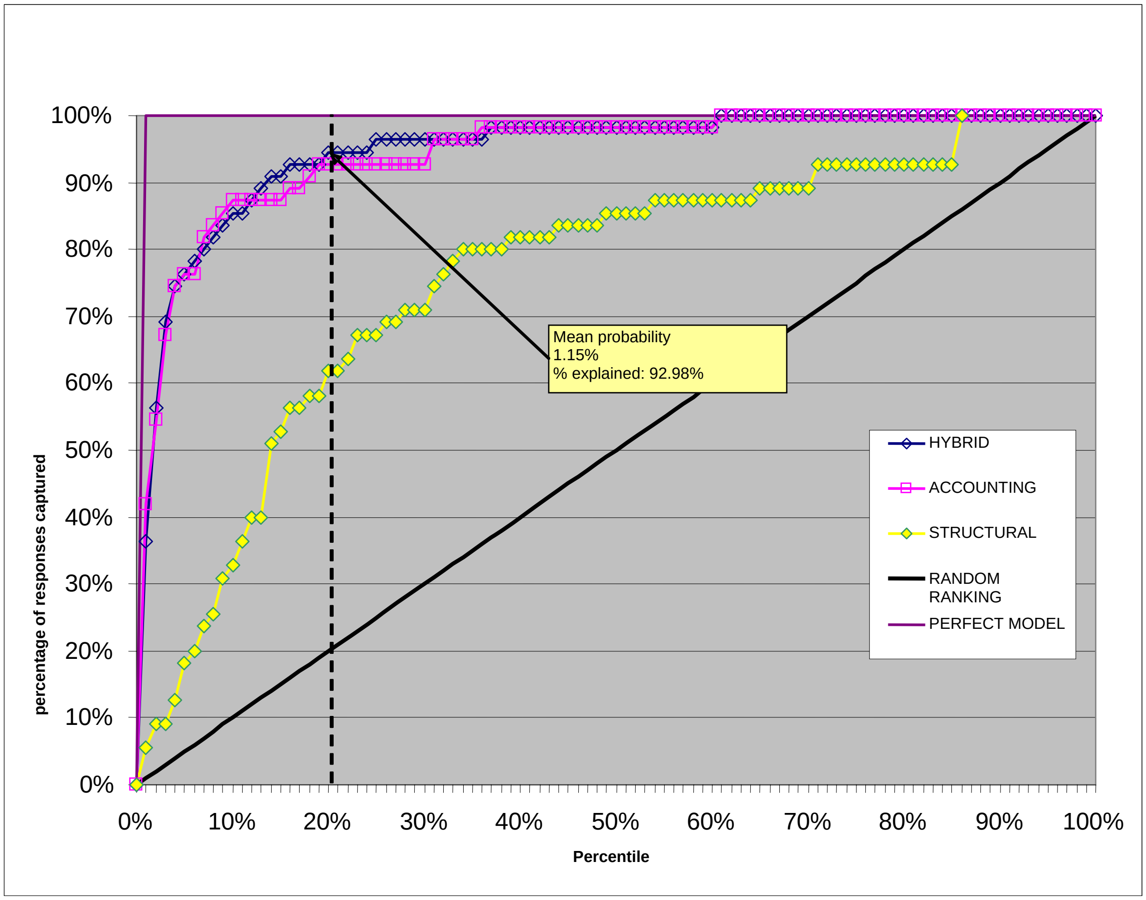

5.4 *Update of the predicted probabilities of default*

Despite the fact that the structural model proves unable to make a significant contribution to the accounting information in our model, in terms of its ability to adequately predict defaults of publicly owned non-financial companies in Canada, it does feature one undeniable advantage. Measures of probabilities of default can be obtained at a much higher frequency from the structural approach than from accounting data. It is also possible to update the probabilities of default predicted by the hybrid model by incorporating PD variables computed on a quarterly, monthly, or even daily basis—and we note that the PD variable is significant in all of the specifications we have examined. We conducted this exercise for some of the firms in our sample. We updated the probabilities of default predicted by the hybrid model by including a quarterly PD. Figure 8 shows that the probability of default can increase dramatically in as much as a year. The same analysis can be applied at a greater level of aggregation, for example, to a given group of firms or sector. Finally, we provide an example with annual variations in the PD and financial statements. We observe substantial differences between the variations in the

PD of the structural model (Figure 3) and those in the probabilities from the hybrid model. These differences can be explained as follows.

The only input used by the structural model is the ratio of the firm's asset-market values to its liabilities. Estimates of the probabilities of default yielded by this model are very sensitive to the level and evolution of this ratio. However, results of the probit estimation of the hybrid model reveal that there are other factors that explain the occurrence of default in publicly owned Canadian firms. Moreover, it is known in the literature that structural models may overstate the probabilities of default, which is consistent with their very high values for Agricore and Asbestos in the structural model. Conversely, the probabilities predicted by the hybrid model incorporate variables other than the probabilities of default from the structural model, allowing a better estimation and calibration of those probabilities. Examination of Figure 8 reveals how the hybrid model allows errors in the estimates of the probabilities of default from the structural model to be corrected. Indeed, it appears clear that, despite the large increases in the structural probabilities of default, forecasts of the probabilities in the hybrid model vary more moderately.

Figure 8: Quarterly and annual probabilities of default

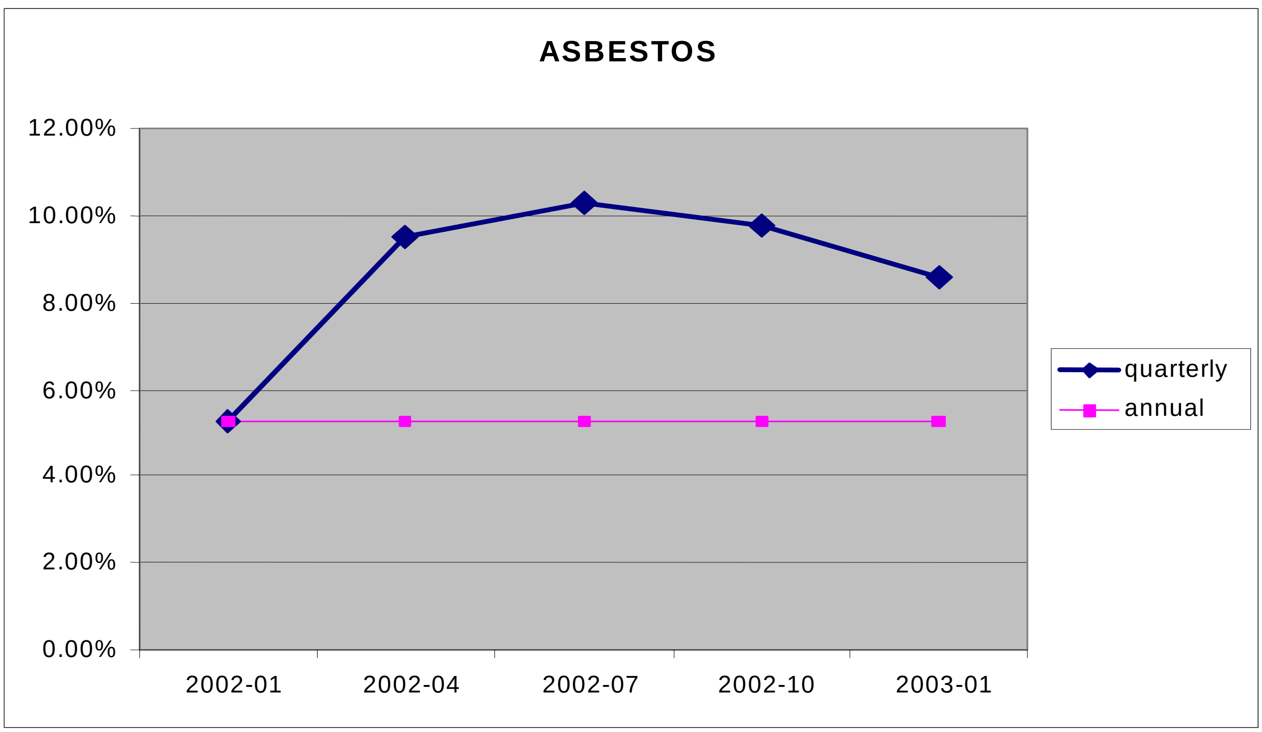

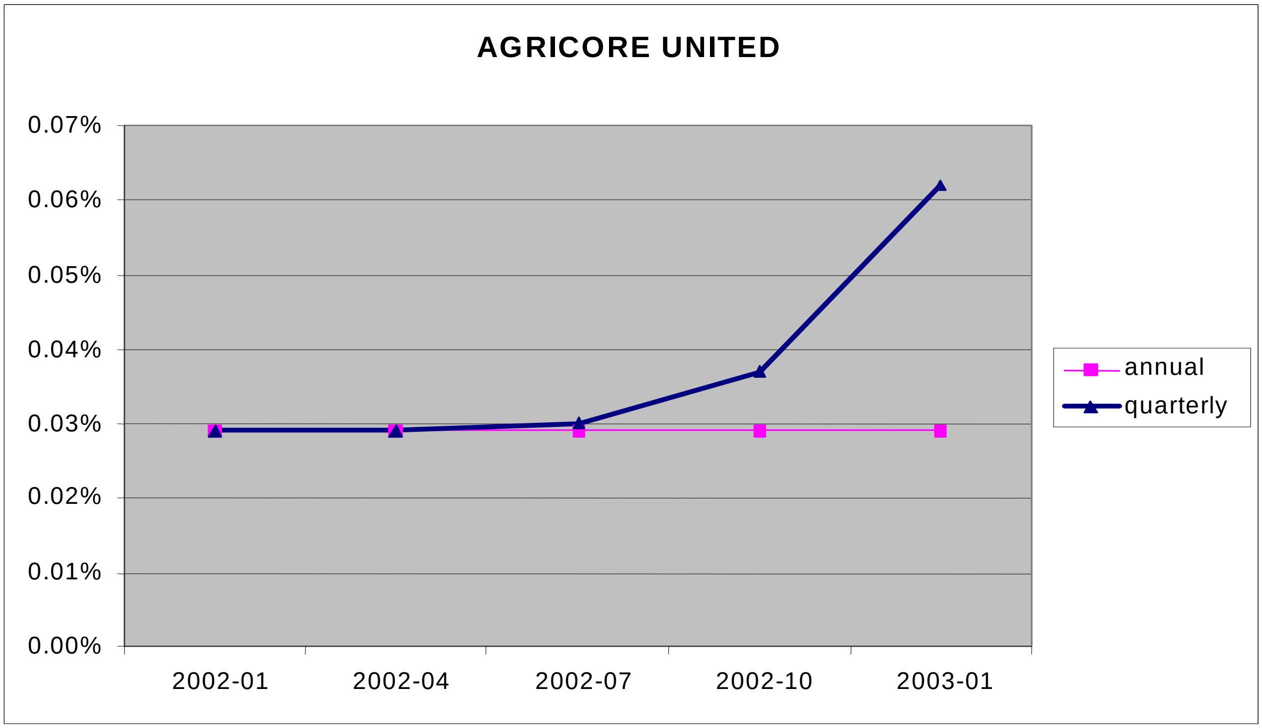
AGRICORE UNITED
0.07%
0.06%
0.05%
0.04%
0.03%
0.02%
0.01%
0.00%
2002-01
2002-04
2002-07
2002-10
2003-01
annual
quarterly

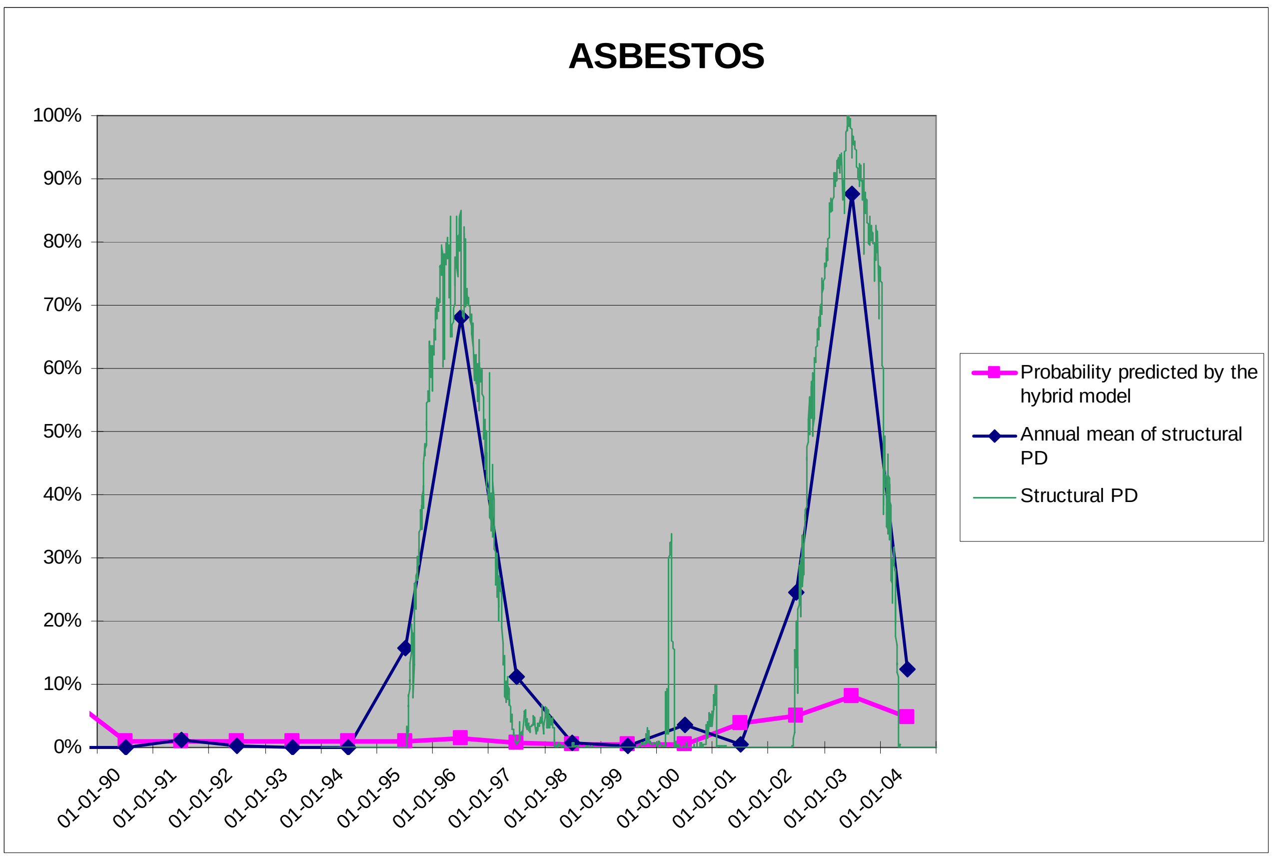
ASBESTOS
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
01-01-90
01-01-91
01-01-92
01-01-93
01-01-94
01-01-95
01-01-96
01-01-97
01-01-98
01-01-99
01-01-00
01-01-01
01-01-02
01-01-03
01-01-04
Probability predicted by the hybrid model
Annual mean of structural PD
Structural PD

5.5 *Comparison of our results with those in Tudela and Young (2003)*

In the Tudela and Young (2003) paper, the structural model explains a greater proportion of defaults compared to our other research. In fact, the accuracy ratio of the structural model for British firms is 69 per cent, compared to 48 per cent for the Canadian firms in our sample.

This difference is not necessarily exclusively attributable to differences in the samples. Our methodology differs from that of the Bank of England in several points of application of the hybrid model. These may explain the difference in the performance of the two categories of models despite the fact that we use the same structural model.

1. To compute firms' liabilities, Tudela and Young (2003) use cubic splines to smooth the series of liabilities into a continuous curve and thus avoid stepwise asset/liability ratios with abrupt jumps at the dates of publication of the financial statements, as we discussed in section 4. Use of the cubic splines method, or any interpolation method, is justified by the elimination of these undesired discontinuities that are detrimental to the estimation of the parameters and by the more accurate reflection of the firm's true situation. However, from a practical perspective it must be borne in mind that implementation of such methods assumes that financial information is available for the following year, which is not at all true for the investor—as a result, the explanatory power of the structural model is overstated. Since we must assume the perspective of the investor, we do not apply this type of interpolation for calculating the inputs of the structural model, which may partly explain why the structural model performed less impressively in our model.

2. Furthermore, following Vassalou and Xing (2003), we use only accounting information on firms' liabilities that is also available to investors. In fact, accounting information is only published four months after the end of the fiscal year. Since we have the dates on which the fiscal years end for the firms in our sample, we lag the financial data four months to be sure that we are only using information available to investors at the time of estimation. Therefore, if the fiscal year ends on 31 December 1999, we use only the 1999 accounting data as of 1 May 2000. This allows us to avoid overstating the structural model's power of discrimination. Conversely, it should be noted that the database manager corrects the data of the preceding years when this proves necessary and, of course, possible. The firm must, however, provide a detailed justification for the alterations before they are made.

3. Also, as to the use of the probability of default from the structural model as an explanatory variable in the probit, Tudela and Young (2003) use the mean of the probabilities of default from the structural model for the 12 months preceding the default. This assumes advance knowledge of the date of default, allowing this information to be integrated into the aggregate weekly probabilities of default. This

type of procedure creates a type of endogeneity that can bias the explanatory power of their model upwards. We, on the other hand, use the mean probability of default for the calendar year regardless of whether or not the firm defaults.

4. Finally, Tudela and Young (2003) select the accounting variables to include in the hybrid model on the basis of preliminary research on British firms (Geroski and Gregg 1997). We select these variables from a large array of accounting variables and financial ratios based on their appropriateness for the sample studies and their potential for explaining firm defaults. In other words, we invested considerable effort into identifying the best possible specification for the reduced-form model before introducing the information from the structural model. This may also explain the better performance of the reduced-form model relative to the structural model in our study, in contrast to the results in Tudela and Young (2003).

6. Correlations between the probabilities of default

In this section, we will propose a portfolio approach to credit risk. This means that we will seek to account for correlations between the risk of default for the firms within a portfolio in order to achieve better estimates of the global default risk for the entire portfolio. One rationale for developing such an approach to credit risk management lies in what we may call "concentration risk." Concentration risk refers to an incremental portfolio risk resulting from increased exposure to a single debtor or a correlated group of debtors (for example, same bank, same industry, or same geographical zone).

Another important reason for the portfolio approach to default risk is to more effectively and rationally account for diversification. Indeed, a bank's decision to increase its exposure to a debtor will result in increasing the marginal risk. Conversely, an equivalent increase in exposure involving a debtor of the same quality, but that was not in the initial portfolio, will substantially lower the marginal risk. Thus, some positions, while individually risky, may only represent a small increase to the global risk of the portfolio, owing to the benefit of diversification. In the two following sections, we present a preliminary analysis of the correlations between the probabilities of default.

6.1 Testing for the presence of correlation

Before proposing a methodology to allow us to account for correlations between the defaults in a given portfolio of firms, we should test for the presence of this type of correlation among the probabilities of default of the firms in our sample. With respect to this, we propose a measure to detect the existence of such correlations between the probabilities of default. For this exercise, we will use the probabilities of default from the structural model, since these are available in time series that are better adapted to capturing the presence of correlations between them. It is reasonable to extrapolate that, if such correlations are found in the probabilities yielded by the structural model, they should also be present in those computed from the hybrid model.

We construct an index of monthly PDs, which is nothing other than the series of individual PDs for the 824 firms for which we have the data required by the structural model. We notice, for example, that this index may cover only a group of firms or a given sector. Thus, for each month in the period covered by our sample, we compute the mean PD for the firms as follows:

$$I_t = \frac{1}{n_t}\sum_{i=1}^{n_t} PD_{it},$$

where n_t is the number of firms for which we have the probability of default for month *t* and PD_{it} is the probability of default of firm *i* during month *t* from the structural model. We obtain a monthly series for our index covering the 204 months (17 years) of our sample (1988–2004).

Subsequently, we compute the correlation between each firm's monthly PD and the index as follows:

$$corr(PD_i, I) = \frac{1}{T}\sum_{t=1}^{T}(PD_{it} - \overline{PD}_i)(I_t - \overline{I}),$$

where

$$\overline{PD} = \frac{1}{T}\sum_{t=1}^{T} PD_{it}, \; \overline{I} = \frac{1}{T}\sum_{t=1}^{T} I_t,$$

and *T* is the number of months for which we have valid observations on PD.

This yields a correlation coefficient for each firm in our sample. Next, we simply calculate the mean of the correlation coefficients to obtain our measure of correlation, $\overline{corr}$.

$$\overline{corr} = \frac{1}{N}\sum_{i=1}^{N} corr(PD_i, I),$$

where *N* is the number of firms for which we have valid monthly observations on the probabilities of default from the structural model. We thus obtain a correlation index of $\overline{corr}$ = 19.68 per cent.

We check the validity of this measure by dividing the sample of firms into two equally sized subsamples and repeating the procedure. The mean correlation for each subsample is very near to the correlation coefficient for the entire sample. As a further test of robustness, we divided the observations into two subperiods of the same length, the first extending from January 1988 to June 1996, and the second from July 1996 to December 2004. We obtain a correlation coefficient of 17.71 per cent for the first subperiod and 19.38 per cent for the second. This demonstrates that the presence of correlation is not specific to the period or the sample, but is rather characteristic of the probabilities of default. Our results underscore the importance of accounting for correlations for adequately estimating joint probabilities of default for a set of firms.

6.2 *The portfolio approach to credit risk, an example*

CreditMetrics proposes an interesting approach to incorporating dependencies between default risks in a portfolio model of credit risk. We draw on it to derive an equivalent method that is better adapted to the data available to us. In the following, we select only two firms and use a concrete example for purposes of illustration, as we compute the credit risk for a portfolio of firms.

The approach proposed here is based on Merton's (1974) theoretical framework, in which the firm's debts are treated as stockholders' call options on the firm's assets. In this context, the value of the firm is a stochastic variable with a certain distribution. If the value of the firm falls below the value of its debts, which we term the default threshold, it becomes impossible for the firm to honour its obligations and it is in default (Figure 9). It is a simple matter to extend this analysis to include changes on the credit side, as defined by rating agencies (e.g., Aa, B). This generalization requires assuming that, in addition to the default threshold, there are thresholds for the value of the firm below which its rating will deteriorate, as illustrated in Figure 10. The value of the firm's assets relative to these thresholds thus determines its future ranking, allowing us to establish a link between the firm's value and its credit rating. In the final analysis, if we know the distribution of the asset values, the default and credit rating transition thresholds, and the correlation between firms' asset values, we can compute the joint probability of default for said firms (for example, the probability that two firms default simultaneously).

Figure 9: Structural model of the firm

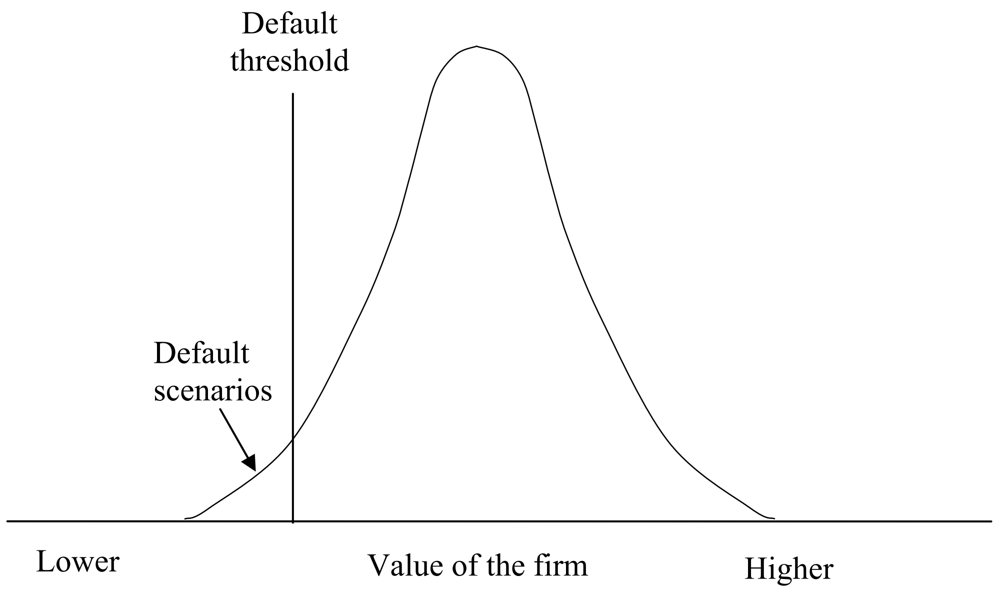

Source: CreditMetrics™ – Technical document.

Figure 10: Structural model of the firm account for changes in credit ratings

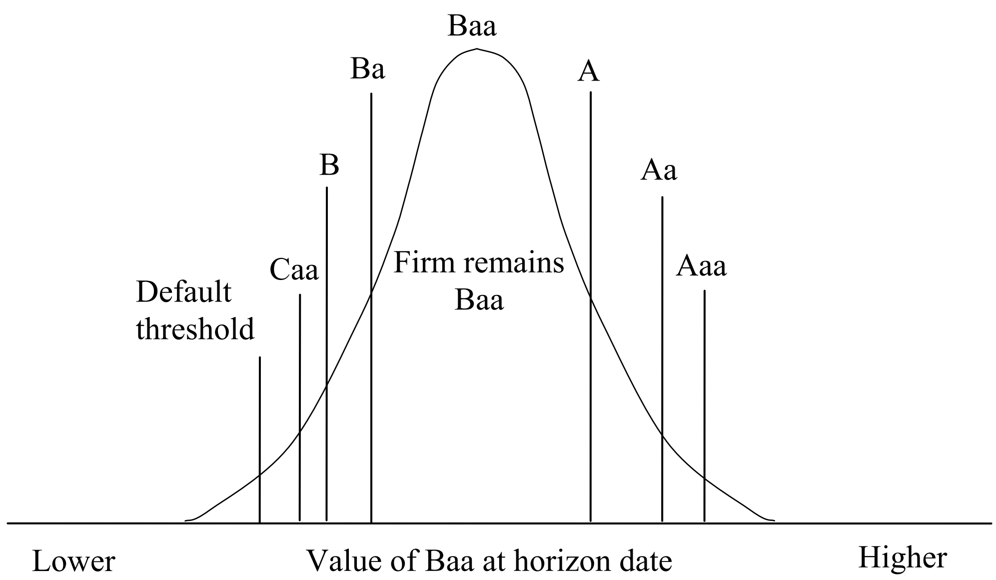

Source: CreditMetrics™ – Technical document.

If we assume that we know the thresholds for the firm's asset values, it remains for us to model the changes to these asset values in order to be able to describe the evolution of its credit. For this, we make the assumption that the distribution of the return to the assets (which we denote R) is normal with known mean and standard deviation.

It is now a simple matter to express the default thresholds in terms of the returns. Thus, we assume that for each class of risk, as defined by the rating agencies, there exists a floor below which the return to the firm's assets will be insufficient to pay its debt.

The default and risk migration thresholds are obtained from the transition matrices of the rating agencies. For example, according to Moody's transition matrix (Table 11), the probability that a firm that was initially rated Baa will be rated Ba the following year is 4.76 per cent, while the probability that it remains unaltered is 88.48 per cent.

Table 11: Annual migration matrix

Initial Rating	Final rating							
	Aaa	**Aa**	**A**	**Baa**	**Ba**	**B**	**Caa**	**Default**
Aaa	93.40	5.94	0.64	0	0.02	0	0	0
Aa	1.61	90.55	7.46	0.26	0.09	0.01	0	0.02
A	0.07	2.28	92.44	4.63	0.45	0.12	0.01	0
Baa	0.05	0.26	5.51	88.48	4.76	0.71	0.08	0.15
Ba	0.02	0.05	0.42	5.16	86.91	5.91	0.24	1.29
B	0	0.04	0.13	0.54	6.35	84.22	1.91	6.81
Caa	0	0	0	0.62	2.05	4.08	69.20	24.06

In our case, we do not have the risk rating of the firms in our sample, so we use the probabilities predicted by our structural model to generate our own. We rank them according to their probabilities of default from the structural model and divide them into 10 classes (10 deciles). Class 0 represents the firms with the poorest credit quality and the highest default risk. These firms correspond to the last decile of the structural PDs. Firms in class 9 are the least risky and have the lowest probabilities of default, corresponding to the first decile. We let the risk of default be measured by the probability that the firm will fall into the last category (class 0). Using this classification, we estimate the transition matrix from the history of risk class migrations for the firms in our sample. For example, we see in Table 12 that the estimated probability that a firm initially in class 3 will finish the year in the same class is 38.65 per cent. To obtain this estimate, we count the number of times a firm in class 3 remains in the same class the following year, then divide by the total number of observations starting in this class.

Moreover, since we assume that the distribution of the return to assets is normal, we know that there exist threshold returns, Z_i, where $i = 0, 1, ..., 9$, delimiting the transitions between risk classes. For example, for a firm in class 3, if the return to assets *R* falls below Z_0, the firm will be downgraded to class 0, and if $Z_0 < R < Z_1$, the firm will be downgraded to class 1. Furthermore, the assumption of normality allows us to compute the probabilities of these events:

$$\Pr\{\text{class}\,0\} = \Pr\{R \le Z_0\} = \Phi\left(\frac{Z_0 - \mu}{\sigma}\right)$$

$$\Pr\{\text{class}\,1\} = \Pr\{Z_0 \le R \le Z_1\} = \Phi\left(\frac{Z_1 - \mu}{\sigma}\right) - \Phi\left(\frac{Z_0 - \mu}{\sigma}\right),$$

etc. Here, μ and σ are the expectation and the standard deviation of the returns, and Φ is the distribution function of the normal distribution. The series of equations connecting the transition probabilities to the threshold values Z_i allows us to compute these latter from the transition matrix we have already estimated.

Working from the transition probabilities, we can estimate the values of Z_i where $i = 0, 1, ..., 9$ (see Table 9) for the two firms in our example, Adaptron Technologies and Agricore United.

Table 12: Calculation of the threshold values associated with risk class migrations

Class in 2002	Predicted class in 2003	Probability	Thresh-old	Value
Adaptron Technologies				
3	0	1.64%	Z_0	-0.61
3	1	5.93%	Z_1	-0.38
3	2	23.72%	Z_2	-0.08
3	3	38.65%	Z_3	0.24
3	4	12.68%	Z_4	0.37
3	5	6.54%	Z_5	0.46
3	6	3.27%	Z_6	0.53
3	7	2.45%	Z_7	0.59
3	8	2.45%	Z_8	0.68
3	9	2.66%		
Agricore United				
4	0	1.43%	Z_0	-1.14

Class in 2002	Predicted class in 2003	Probability	Thresh-old	Value
4	1	2.25%	Z_1	-0.97
4	2	5.93%	Z_2	-0.77
4	3	20.04%	Z_3	-0.44
4	4	35.58%	Z_4	-0.05
4	5	17.18%	Z_5	0.18
4	6	6.75%	Z_6	0.31
4	7	5.11%	Z_7	0.46
4	8	2.86%	Z_8	0.59
4	9	2.86%		

Thus far, we have not examined the transition probabilities of the two firms individually by linking them to the returns to their assets. To compute the joint transition and default probabilities, we must establish whether the returns to assets are correlated. It remains to calculate the correlation coefficient ρ between the assets of the two firms. The estimated correlation coefficient between the assets of these firms is 0.04. We now find the variance-covariance matrix for the multivariate normal distribution.

$$\Sigma = \begin{bmatrix} \sigma_1^2 & \rho\sigma_1\sigma_2 \\ \rho\sigma_1\sigma_2 & \sigma_2^2 \end{bmatrix}.$$

This allows us to calculate the joint probability of the two firms simultaneously being downgraded to 0.

$$\Pr\left\{R \le Z_0, R' \le Z_0'\right\} = \int_{-\infty}^{Z_0}\int_{-\infty}^{Z_0'} f(r,r',\Sigma)\,dr\,dr',$$

where $f(r,r',\Sigma)$ is the density function of the bivariate normal distribution with variance-covariance matrix Σ. (With Matlab, the normal bivariate distribution function requires downloading the function *mvncdf*.)

As an exercise, we use several values for the correlation coefficient in this paper. In the first instance, we assume that the returns of the two companies are independent, so that the joint probability of default is simply the product of the two marginal probabilities. In this case, it is 0.02 per cent. In the second instance, we assume perfect correlation between these returns (ρ = 1). In this case, the joint probability is 1.47 per cent, i.e., 62 times higher than when the returns are independent. Finally, we look at the intermediary case of a correlation coefficient equal to 0.5. This yields a joint default probability equal to 0.22 per cent, which is over ten times its value in the case of independence (Table 13).

This underscores the importance of accounting for correlations between the probabilities of default in order to be able to properly assess firms' default risks.

Table 13: Calculation of the joint default probabilities given three hypotheses regarding the correlations between asset values

	Ranking	Individual default probabilities (class 0)	Expected returns	Variance of returns	Joint default probability		
					ρ = 0	ρ = 1	ρ = 0.5
Adaptron Technologies	3	1.63%	7.10%	0.0038	0.02%	1.47%	0.22%
Agricore United	4	1.43%	-21.31%	0.0005			

Conclusion

The goal of this research was to determine how a continuous evaluation of the probabilities of default of publicly traded firms by the stock exchange might improve the prediction that a firm may default. One way of accomplishing this goal is to estimate a hybrid model in which the estimated probabilities of default from the structural model are introduced into the reduced-form model as explanatory variables.

We conducted this exercise for publicly traded Canadian companies whose shares are traded on the Toronto Stock Exchange. Our results indicate that the predicted probabilities of default (PDs) contribute significantly to explaining default probabilities when they are included alongside the retained accounting variables. We also show that quarterly updates to the PDs add a large amount of dynamic information to explain the probabilities of default over the course of a year. This flexibility would not be possible with reduced-form models unless audited accounting data were available on a quarterly basis.

We also conducted a preliminary analysis of correlations between structural probabilities of default for the firms in our database. Our results indicate that there are substantial correlations between these probabilities of default. If this information were to be borne out by a more detailed analysis of the data, that would suggest that holders of portfolios of corporate debt, such as banks, should account for these correlations when assessing their capital requirements. It would also indicate that these correlations should be accounted for when probabilities of default are aggregated across industries or regions for purposes of economic policy. Finally, we provided a cursory presentation of a model allowing correlations between the values of firms' assets to be used in calculations of joint default probabilities in the hybrid model.

There are several possible extensions to this initial analysis. First, a method could be developed for aggregating the analysis over industrial sectors or over financial institutions' portfolios. This aggregation should account for correlations between the probabilities of default of the firms included. Ultimately, this model could be used to construct more diversified loan portfolios by the banks.

A second extension pertains to the estimation of the PD by the structural model. Two issues raised in the literature review have been ignored thus far. The first consists of estimating the parameter of the capital structure simultaneously with the other parameters by using the maximum-likelihood method in the structural model. The second involves applying the data-filtering algorithm of Duan and Fulop (2005) in order to reduce bias in the estimates of standard deviations associated with trading noises on stock exchanges that has an impact on the one-to-one relationship between asset values and firm values.

Finally, it would be very useful to adapt this method to the purposes of economic policy. This requires finding the relevant aggregates and choosing the periods in which the aggregates must be continuously updated, so as to disseminate the information to the affected financial institutions.

References

Altman, E.I. 1968. "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy." *Journal of Finance* 23 (4): 589–609.

Altman, E.I. and M. Lavallée. 1981. "Business Failure Classification in Canada." *Canadian Journal of Business Administration* (Summer).

Beaulieu, S. 2003. "Modèle d'évaluation des probabilités de défaut un an des firmes privées canadiennes : approche probit non-structurelle." Master's thesis, HEC Montréal.

Black, F. and M. Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (3): 637–54.

Brennan, M.J. and E.S. Schwartz. 1978. "Corporate Income Taxes, Valuation, and the Problem of Optimal Capital Structure." *Journal of Business* 51 (1): 103–14.

Briys, E. and F. de Varenne. 1997. "Valuing Risky Fixed Rate Debt: An Extension." *Journal of Financial and Quantitative Analysis* 32 (2): 239–48.

Brockman, P. and H.J. Turtle. 2003. "A Barrier Option Framework for Corporate Security Valuation." *Journal of Financial Economics* 67 (3): 511–29.

Crosbie, P. and J. Bohn. 2003. "Modeling Default Risk." Mimeo, Moody's/KMV.

Crouhy, M., D. Galai, and R. Mark. 2000. "A Comparative Analysis of Current Credit Risk Models." *Journal of Banking and Finance* 24 (1–2): 59–117.

Dionne, G., G. Gauthier, K. Hammami, M. Maurice, and J.G. Simonato. 2005. "Default Risk in Corporate Yield Spreads." Working Paper No. 05-08, Canada Research Chair in Risk Management, HEC Montréal.

Duan, J.C. and A. Fulop. 2005. "Estimating the Structural Credit Risk Model When Equity Prices Are Contaminated by Trading Noises." Mimeo, University of Toronto.

Duan, J.C., G. Gauthier, and J.G. Simonato. 2004. "On the Equivalence of the KMV and Maximum Likelihood Methods for Structural Credit Risk Models." Mimeo, HEC Montréal.

Georgakopoulos, V. 2005. "Current Approaches to Credit Risk Measurement." Mimeo, National and Kapodistrian University of Athens.

Geroski, P.A. and G. Gregg. 1997. *Coping with Recession: UK Company Performance in Adversity*. Cambridge: Cambridge University Press.

Gordy, M.B. 2000. "A Comparative Anatomy of Credit Risk Models." *Journal of Banking and Finance*, 24 (1–2): 119–49.

Gupton, G.M., C.C. Finger, and M. Bhatia (1997). *CreditMetrics – Technical Document*. J.P. Morgan & Co. Incorporated. Available at <http://www.defaultrisk.com>.

Jarrow, R.A., D. Lando, and S.M. Turnbull. 1997. "A Markov Model for the Term Structure of Credit Risk Spreads." *Review of Financial Studies* 10 (2): 481–523.

Jarrow, R.A. and S.M. Turnbull. 1995. "Pricing Derivatives on Financial Securities Subject to Credit Risk." *Journal of Finance* 50 (1): 53–85.

Leland, H.E. 1994. "Corporate Debt Value, Bond Covenants, and Optimal Capital Structure." *Journal of Finance* 49 (4): 1213–52.

Leland, H.E. and K.B. Toft. 1996. "Optimal Capital Structure, Endogenous Bankruptcy, and the Term Structure of Credit Spreads." *Journal of Finance* 51 (3): 987–1019.

Merton, R.C. 1973. "Theory of Rational Option Pricing." *Bell Journal of Economics and Management Science* 4 (1): 141–83.

———. 1974. "On the Pricing of Corporate Debt: the Risk Structure of Interest Rates." *Journal of Finance* 29 (2): 449–70.

Sobehart, J.R., S.C. Keenan, and R.M. Stein. 2000. "Benchmarking Quantitative Default Risk Models: A Validation Methodology." New York: Moody's Investors Service.

Tudela, M. and G. Young. 2003. "A Merton-Model Approach to Assessing the Default Risk of UK Public Companies." Working Paper No. 194, Bank of England.

Vassalou, M. and Y. Xing. 2004. "Default Risk in Equity Returns." *Journal of Finance* 59 (2): 831–68.

Wong, H. and T. Choi. 2004. "The Impact of Default Barrier on the Market Value of Firm's Asset." Working Paper, Chinese University of Hong Kong.

Zhou, C. 1997. "A Jump-Diffusion Approach to Modeling Credit Risk and Valuing Defaultable Securities." Working Paper No. 1997-15, Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System.

———. 2001. "The Term Structure of Credit Spreads with Jump Risk." *Journal of Banking and Finance* 25 (11): 2015–40.

Appendix: Receiver operating characteristic (ROC) curves

Given a sample of *n* observations, let n_1 represent the number of defaults. We denote this group C_1 and the remaining non-defaulting group of $n_2 = n - n_1$ observations C_2. Risk factors are identified in the sample and a probit regression model is fit to the data. An estimated default probability, $\hat{\pi}_i$, is calculated for the *i*-th observation.

Now, assume that, for the *n* observations, we undertake to test a default prediction on the basis of the estimated probability of default. The highest values of the estimated probability are expected to correspond to default. A receiver operating characteristic (ROC) curve can be constructed by shifting the threshold at which the estimated probability is considered a predictor of default. For each threshold *z*, the following measures can be calculated, where *I* designates the index function:

number of defaults correctly predicted: $POS(z) = \sum_{i \in C_1} I(\hat{\pi}_i \geq z)$;

number of non-defaults correctly predicted: $NEG(z) = \sum_{i \in C2} I(\hat{\pi}_i < z)$;

number of non-defaults incorrectly predicted as defaults: $FALPOS(z) = \sum_{i \in C_2} I(\hat{\pi}_i \geq z)$;

number of defaults incorrectly predicted as non-defaults: $FALNEG(z) = \sum_{i \in C_1} I(\hat{\pi}_i < z)$.

We obtain the values for:

$$\text{Sensitivity} = \frac{POS(z)}{n_1} \text{ and } 1 - \text{specificity} = \frac{FALPOS(z)}{n_2}.$$