Journal of Real Estate Finance and Economics, 17:2, 163–178 (1998)

Commercial Mortgage Default: A Comparison of Logit with Radial Basis Function Networks

ATHANASIOS EPISCOPOS
National Bank of Greece, Eolou 86, Athens, Greece 102-32

ANDREAS PERICLI
Asset/Liability Management, Federal Home Loan Mortgage Corporation, 8200 Jones Branch Drive, McLean, VA 22102

JIANXUN HU
Advanta Mortgage Corporation, 850 Ridgeview Dr., Horsham, PA 19044

Abstract

This article explores the use of artificial neural networks in the modeling of foreclosure of commercial mortgages. The study employs a large set of individual loan histories previously used in the literature of proportional hazard models on loan default. Radial basis function networks are trained (estimated) using the same input variables as those used in the logistic. The objective is to demonstrate the use of networks in forecasting mortgage default and to compare their performance with that of the logistic benchmark in terms of prediction accuracy. Neural networks are shown to be superior to the logistic in terms of discriminating between ''good'' and ''bad'' loans. The study performs sensitivity analysis on the average loan and offers suggestions on further improving prediction of defaulting loans.

Key Words: commercial mortgages, default, prediction, artificial neural networks, radial basis functions, logistic function

1. Introduction

Mortgage default (foreclosure) results when the borrower decides to exercise the default put option embedded in the mortgage contract (Foster and Van Order, 1984, 1985; Vandell, 1992; Quigley and Van Order, 1994). In essence, the defaulter exercises the put option and ''sells'' the property back to the lender in exchange for eliminating the mortgage obligation.

The American Council of Life Insurance (ACLI) has recently reported that foreclosures on commercial mortgages in 1994 amounted to $4.4 billion or a proportion of 2.25% of the commercial mortgage loans. The average loss from foreclosures was 29.04% or $1.28 billion (ACLI, 1995). While foreclosure rates were the lowest since 1990, when they averaged 1.66% of the mortgage loans, in the middle Atlantic and New England regions foreclosure rates averaged 5.15 and 4.09%, respectively, in the four quarter of 1994 (ACLI, 1995).

The considerable default risk and a recent increase in commercial mortgage securitization have renewed interest in improving commercial mortgage default predictions. Commercial mortgage lenders credit rating agencies, brokerage houses, and commercial mortgage-backed securities investors are now eager to explore new techniques that could improve their credit risk predictions.

Empirical research on commercial mortgage default risk has been very limited. A small set of empirical papers (Titman and Torrus, 1989; Kau, Keenan, Muller, and Epperson, 1990; Vandell, 1992; and Vandell, Barnes, Hartzell, Kraft, and Wendt, 1993) examine commercial mortgage credit risk. The lack of empirical research arises largely because of the lack of available data. Vandell et al. (1993) is an excellent source that summarizes the shortcomings of previous research. These shortcomings are mainly attributable to the use of aggregate rather than disaggregate data.

If prediction of commercial mortgage default risk is the item of current concern, it is desirable to use the best available data and the most powerful classification techniques to improve default risk predictions. This is indeed the aim of this article and the main reason for using artificial neural network (ANN) techniques for the modeling of commercial mortgage default risk. The study uses a data set of 2,899 loan histories provided by a major multiline insurance company. The same set has been employed by Vandell et al. (1993) to model default risk in a proportional hazard framework. The Vandell et al. (1993) study is used as the starting point for determining the explanatory variables (inputs) of the ANN and logistic approach.

ANNs have been recently used extensively for solving several highly nonlinear financial problems (e.g., Refenes, 1995). ANNs are learning algorithms that are capable of recognizing patterns in data. Function approximation has been generally their major strength (see White, Gallant, Hornik, Stinchcombe, and Wooldridge, 1992, for the theory, and Hutchinson, Lo, and Poggio, 1994, for an application to stock option pricing). In particular, financial applications that require classification of inputs into two or more groups are excellent candidates for this methodology. A special case of this problem is binary classification in which the groups are limited to two (for instance, default or nondefault). Other related financial applications of neural nets include predicting the rating of corporate bonds (Dutta and Schekar, 1988), bank failure prediction (Tam and Kiang, 1992), risk assessment of single-family mortgage applications (Reilly, Collins, Scofield, and Ghosh, 1991), and emulation of mortgage underwriting judgment (Collins, Ghosh, and Scofield, 1988).

This article contributes to the ongoing interest regarding the forecasting accuracy of ANNs for discriminating between good and bad loans in at least four important ways. First, to the best of our knowledge, there have been no other studies in the literature that employ neural networks in predicting the default risk of commercial mortgages. Second, the article uses radial basis function networks (RBFN) to model default probabilities of mortgages. RBFN are a comprehensive and important class of networks different from, and in many ways superior to, the ubiquitous back-propagation networks. Third, it performs forecasting accuracy comparisons on training and holdout samples of RBFN with the logistic benchmark. Fourth, the article employs an extensive, proprietary, and unique commercial mortgage database that has been used in a proportional hazard model.

Fifth, the article proposes a way of reducing the severity of censored samples, a problem often overlooked by practitioners.

The findings suggest that radial basis functions predict commercial mortgage defaults with higher forecasting accuracy than the logistic functions do. In particular, in the holdout sample the forecasting accuracy improves by 4%. Given the $1.8 billion losses due to defaults in 1994, the 4% improvement in forecasting accuracy implies a reduction in losses of about $72 million.

The rest of the article is organized as follows: In section 2 the sample and the explanatory variables used in the empirical analysis are presented. Section 3 provides a brief description of the logistic and the radial basis function network. Section 4 shows the empirical results and compares the classification accuracy of the two methods. Section 5 performs a sensitivity analysis that offers additional insights on the behavior of the estimated classification function. Section 6 demonstrates how default prediction can be improved. Section 7 summarizes the principal findings.

2. Sample and Explanatory Variables

The database used to compare the forecasting accuracy of default probabilities between RBFNs and the logistic approach is provided by a large multiline insurance company. The database contains 2,899 first-lien mortgage loans, beginning with originations in 1962 and concluding with originations in the third quarter of 1989, among which 175 satisfy the definition of default proposed by Vandell et al. (1993). Default loans are defined either as those that have been foreclosed by the lender or as those for which the payment is made in full after the loan has gone through the foreclosure process.

Observations are highly censored, and thus the logit model suffers from length-based sampling bias since it treats both the foreclosed and current loans as completed payment records. To reduce the severity of this problem we have considered only the loans that have been originated before December 1986. The filtered sample contains 1,619 loans and includes all 175 default ones. Table 1 displays summary statistics on the loan characteristics that are available for all the set of 1,619 loans.

As in Vandell et al. (1993), several dummies were created in order to investigate the differential impact of regions, borrower types, property types, and loan types on default risk. Specifically, there are eight regional dummies,[1] six property-type dummies, four borrower-type dummies, and four loan-type dummies.

To put the two methods on equal footing, the same inputs will be used by both the logistic model and the radial basis function network. The data set is broken down to a training and holdout set, in a proportion of about 80 to 20%, respectively, and the records were randomly selected so as to avoid localization of the testing results.

Table 2 shows the structure of the training and holdout samples. As it can be seen, the training set includes 150 default and 1,162 nondefault loans, while the holdout sample includes 25 default and 282 nondefault loans.

Table 1. Summary statistics for the 1,619 loans.

Definition	Variable	Mean	Minimum	Maximum
Market loan-to-value (%)	MDMLTV	51.66	0.30	161.80
Interest rate (%)	INTRTE	10.0	6.0	18.0
Debt coverage	ORIGDSC	1.29	0.75	2.17
Loan term (months)	LNTERM	176.53	19.00	499.00
Loan amount ($)	LNAMOUNT	4,329,497.36	1,000.00	201,431,818.00
Loan date (years)	LNDATE	1978	1962	1985
Maturity date (years)	MATDT	1993	1985	2013
Apartment	PROPAPT	0.32		
Hotel	PROPHOTL	0.02		
Office	PROPOFF	0.32		
Retail	PROPRET	0.14		
Industrial	PROPINDU	0.16		
Others	PROPOTH	0.04		
Borrower type				
Individual	BRWRTY01	0.14		
Partnership	BRWRTY02	0.51		
Corporation	BRWRTY04	0.09		
Other	BRWRTY03	0.26		
Apartment	PROPAPT	0.32		
Hotel	PROPHOTL	0.02		
Office	PROPOFF	0.32		
Retail	PROPRET	0.14		
Industrial	PROPINDU	0.16		
Others	PROPOTH	0.04		
Location by region				
East North Central	ENC	0.16		
Mideast	ME	0.08		
Mountain	MTN	0.10		
New England	NE	0.09		
Pacific	PAC	0.22		
Southeast	SE	0.11		
Southwest	SW	0.17		
West North Central	WNC	0.07		
Cash flow	GPMDUMMY	0.18		
Step rate	GPMDUMY2	0.04		
Accrual	ACRDUMMY	0.03		
Amortization	AMZDUMY2	0.85		
Default	KDVDEF	0.07		
Borrower type				
Individual	BRWRTY01	0.14		
Partnership	BRWRTY02	0.51		
Corporation	BRWRTY04	0.09		
Other	BRWRTY03	0.26		

Notes: The summary statistics are performed on all the loans that have been originated between January 1962 and December 1986 and for which completed records were found. This sample includes 1,619 loans and includes all 175 default ones. ACRDUMMY is a dummy variable that takes the value of one if the loan was originally an accrual one and zero otherwise. GPMDUMMY is a dummy variable indicating if there were changes in the cash flows of the loan. GPMDUMY2 is a dummy variable indicating if the loan was a true graduate loan.

Table 2. Training and holdout sets.

	Training Set	Holdout Set	Total
Default	150	25	175
Nondefault	1,162	282	1,444
Total	1,312	307	1,619

Note: Figures show the number of loan records included in each category.

3. Model Construction

This section presents a brief description of the logistic model and the RBFN.

3.1. The Logistic Approach

The logit model is based on the cumulative logistic probability function and is specified as

$$P(Y_i = 1 \mid X_i) = F(X_i\beta),$$

where the dependent variable Y takes either the value of 0 or 1, X is the set of explanatory variables, and

$$F(z) = \exp(z)/(1 + \exp(z))$$

is the logit model.

AMZDUMY2 is a dummy variable indicating if complete or partial amortization was scheduled. KDEVDEF is a dummy variable indicating if the loan has been defaulted. LNTERM are the number of months from the first full payment date through the maturity date. LNAMOUNT is the original principal amount. LNDATE is the closing date for the loan—that is, the date cash is disbursed or the documents are signed, whichever occurs later. MATDT is the date that the loan matures as set forth by the note. The regional variables are dummy variables that take the value of one if the property is in the specified region and zero otherwise. For example, MTN is a dummy variable indicating if the property is in the mountain region. The property type variables are dummy variables taking the value of one if the property is of certain type and zero otherwise. For example, PROPAPT is a dummy indicating if the property is an apartment. Borrower type variable are dummy variables that take the value of one if the borrower is of certain type and zero otherwise. For example, BRWRTY01 is a dummy indicating if the borrower is an individual. INTRTE is the origination coupon rate as set forth on the note. ORIGDSC is the net operating income to annual debt service at the time of origination. MDMLTV is the market loan-to-value and is computed as follows. MDMLTV = (LOANTOVL * LINDEX)/PINDEX, where LOANTOVL is the loan amount divided by the appraisal amount at the time of origination, LINDEX is the present value of the remaining cash flows of the loan discounted at the prevailing market rate and divided by the present value of the remaining cash flows discounted at the original contract rate, and the PINDEX is the value of a property index based on geographic region as of the date of default or the date of maturity divided by the value of the property index at the time of origination. The right-censoring point is December 30, 1989.

Estimates of the parameters are obtained by maximizing the sample log-likelihood function

$$\log L(\beta \mid X) = \sum_{i=1}^{n} \log P(Y_i) + \sum_{i=n+1}^{N} \log\,(1 - P(Y_i)),$$

where $Y_i = 1$ for the first n observations, and 0 for the last N-n observations. The maximization of the log L(.) is based on the Berndt, Hall, Hall, and Hausman (1974) algorithm. A loan will be accurately classified if the estimated probabilities for the specific binary choice exceed 0.5—that is, when $P(Y_i = 1) > 0.5$ or $P(Y_i = 0) > 0.5$.

3.2. Radial Basis Function Networks (RBFN)

A neural network is generally a nonlinear mapping from an input to an output space. A network can be represented via its *architecture*—that is, a schematic showing the information flows through each component of the mapping. *Nodes*, *processing elements*, or *neurons* are points of (usually nonlinear) transformation of values. The functions commonly used are *sigmoids* such as the logistic. However, the RBFN use ''bell-shaped'' or *radial basis* functions. The set of nodes between the input and the output nodes from the *hidden layer*, and more than one hidden layers are sometimes encountered. Arrows indicate the direction of the flow of values.

The objective of *training* is to present or feed into the network the input records and compare the output with the actual data in order to *learn* (estimate) the parameters of the model. Due to the often severe nonlinearity involved, estimation techniques such as maximum likelihood fail, in contrast to learning algorithms, which attempt to minimize sequentially an objective function such as the error of the network.

Figure 1 shows the architecture of the RBFN used in this study. The input vector includes all the loan characteristics. The output vector takes the value of (1,0) if the loan is in default and (0,1) if it is not. The architecture and training follows Moody and Darken (1989). The network response function is given by $F(x) = Wf(x)$, where $f(x) = (f_1(x), f_2(x), \ldots, f_N(x))^T$ and W is a 2 by N matrix of weights. Each transfer function is given by

$$f_k(x) = R(\|x - c_k\|/s_k)$$

for $k = 1, 2, \ldots, N$ where N is the number of pattern units, R is the Gaussian function, $R(z) = \exp(-0.5z^2)$, c_k is a vector in the input space called center of the kth unit, s_k is the width of the unit, and $\|\cdot\|$ is the Euclidian norm.

The parameters to be estimated are c_k and s_k for all k, and the weights forming the matrix W. Obviously, if we eliminate the pattern layer and one of the outputs, the network reduces to a simple linear discriminant model.

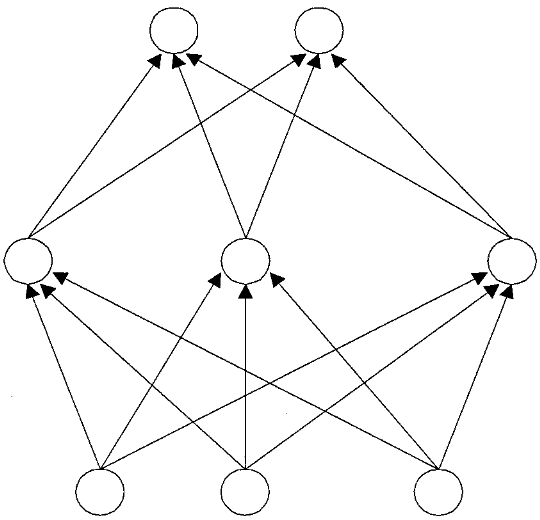

Figure 1. Radial basis function network.

Training takes place in two stages. In the first stage the centers and widths are estimated. The adaptive k-means clustering algorithm starts with randomly chosen input records (exemplars) as initial centers. Then in each step, a new random exemplar is chosen, and the nearest cluster center is moved proportionally to its vector difference from the new exemplar. Thus, the input space is effectively covered with as many regions (clusters) as neurons in the pattern layer. Each width parameter is computed as the root mean square distance of each center from the P nearest neighbor centers:

$$s_k = \left[(1/P) \sum_{p=1}^{P} \|x_k - x_{kp}\|^2 \right]^{1/2}.$$

In the second stage the upper half of the network in figure 1 is trained by a standard gradient descent algorithm, and the weight matrix W is obtained. The learning algorithm starts with a random W and adjusts it to minimize the mean square error

$$E = \left(\sum_t \|a_t - d_t\|^2 \right) / 2,$$

where a_t is the actual output from the network in each step t, and d_t is the desired output, ((1,0), (0,1)). For an intuitive explanation of the process, imagine the error function plotted against two weights only, for instance, as a U-shaped surface, and suppose we are at a point on the surface with coordinates some initial weights. The gradient descent technique adjusts the weights in a direction that makes the point move opposite to the gradient of the surface—that is, downward. The factor of adjustment is a number much less than one called the *learning rate*. Given specific values for centers and widths, the global minimum of the error function can thus be reached to a desirable degree of proximity.

Radial basis functions have certain advantages over other networks such as the frequently used back-propagation network (BPN). RBFN train faster than BPN, avoid the problem of multiple local minima, and, more important, can classify better than BPN, as shown vividly in Leonard and Kramer (1991). The ability to classify better is due to the network architecture and the bell shape of the radial basis functions. Take, for instance, a specific node in the pattern layer. If an input record is located far from the cluster center the node corresponds to, the value of the Gaussian function (at the tails) will be negligible, and the particular record will have insignificant influence during the second stage of the training. In other words, its *signal* will be weak. In contrast, a record that is close to the center of that node will be strongly represented. The upshot is that the network features neurons each of which is dedicated to a specific area of the input space, allowing it to be in a sense ''specialized'' in the prediction of that particular section.

It is to be noted that the general neural network approach is a legitimate and useful scientific technique and by no means a ''black box.'' First, the training process can be replicated exactly if the same random numbers are used and if the same steps are taken. Second, once the network is trained, it can be used just like any other model. For instance, because the radial basis functions are differentiable, one can take partial derivatives to perform comparative statics in some economics applications or to construct hedging portfolios in some financial option applications. In our case, the model is used to generate predictions and to conduct sensitivity analysis. Thus, the model becomes an explicit analytical tool after training.

4. Empirical Results

As mentioned earlier, this article seeks to compare the forecasting accuracy between the RBFN and the logistic model that can be regarded by practitioners as the standard tool to model default risk.[2]

The selection of explanatory variables is an important task for the development of accurate forecasts. As a starting point we have chosen those variables that Vandell et al. (1993) have found to be statistically significant in their proportional hazards model. Then we added variables that are known to explain the economic behavior of default risk, such as regional dummies, in order to capture the differential impact of regions on the foreclosure rates. Table 3 shows the logistic model results when explanatory variables are the ones adopted by Vandell et al. (1993), and table 4 shows the logistic model with additional variables included. We refer to these as model 1 and model 2, respectively. The regional variables are statistically significant and add to the explanatory power of the model. Furthermore, the likelihood ratio statistic *LR* (given in the bottom of table 4) suggests that the inclusion of the constant and the regional variables add to the explanatory power of the models. It should be pointed out, however, that the inclusion of the regional

Table 3. Model 1: logistic model using the Vandell et al. (1993) explanatory variables.

Variable	Parameter Estimate	*T*-Statistic	*P*-Value
MDMLTV90	0.043	7.288	0.000
INTRTE	– 14.234	– 2.536	0.011
ORIGDSC	– 2.676	– 5.136	0.000
PROPAPT	1.576	3.933	0.000
PROPHOTL	0.704	0.873	0.383
PROPOFF	0.917	2.299	0.022
PROPINDU	0.411	0.843	0.399
PROPOTH	1.907	2.280	0.023
BRWRTY01	0.312	0.691	0.490
BRWRTY02	– 0.025	– 0.069	0.945
BRWRTY04	0.016	0.041	0.967
GPMDUMMY	0.713	2.548	0.011
GPMDUMY2	– 2.007	– 3.668	0.000
ACRDUMMY	0.787	1.761	0.078
AMZDUMY2	– 1.517	– 5.669	0.000
LOGL	– 337.69		

Note: See the notes of table 1 for variable definitions and other explanations where LOGL is the loglikelihood function evaluated at the maximum.

dummies is known to have weaknesses. However, this being an empirical study, and because the sample period covers several business cycles, we have undertaken this experiment. The results without the regional dummy variables are also presented. A more accurate approach is to use regional demographic and economic variables to explain the behavior of defaults.

Other variables were tried, such as loan amount and loan term. However, they were colinear with other variables and were dropped.[3] During training several architectures were tried for the RBFN.[4] The search went through a variety of models with very similar performance to the logistic. In most cases the RBFN performed significantly better within sample, as expected. The final RBFN and the logit are applied to the data, and the predictions are reported in tables 5 and 6.

In table 5, using the variables of model 1, the results show that in the training set out of 150 defaulting mortgages logit classified forty of them correctly for a correct classification rate of 27%, whereas RBFN classified fifty-six of them correctly for a correct classification of 37%. Out of the 1,162 nondefaulting loans the logit model predicted 98% correctly, while the RBFN predicted only 97% of them correctly. In the holdout set, out of twenty-five defaulting loans the RBFN classified seven of them correctly for a correct classification rate of 28%, whereas RBFN classified twelve correct for a correct classification rate of 48%. Out of the 282 nondefaulting loans the logit model predicted 98% correctly, whereas the RBFN predicted 96% correctly.

In table 6, the empirical findings of model 2 show that out of the 150 defaulting loans the logit classified 43% correctly, whereas the RBFN classified 65% correctly. Out of 1,162 loans nondefaulting the logit classified 97% correctly, whereas the RBFN predicted 95%

Table 4. Model 2: logistic model by considering additional variables.

Variable	Parameter Estimate	*T*-Statistic	*P*-Value
CONSTANT	− 3.749	− 1.904	0.057
MDMLTV90	0.015	2.129	0.033
INTRTE	19.456	2.139	0.032
ORIGDSC	− 0.663	− 0.662	0.508
PROPAPT	1.983	4.334	0.000
PROPHOTL	1.111	1.193	0.233
PROPOFF	1.534	3.392	0.001
PROPINDU	1.028	1.861	0.063
PROPOTH	1.968	1.989	0.047
BRWRTY01	0.472	0.895	0.371
BRWRTY02	0.297	0.704	0.481
BRWRTY04	0.257	0.557	0.577
GPMDUMMY	1.092	3.420	0.001
GPMDUMY2	− 2.220	− 3.588	0.000
ACRDUMMY	0.596	1.158	0.247
AMZDUMY2	− 1.997	− 6.003	0.000
NEC	− 2.555	− 4.669	0.000
ME	− 3.734	− 4.660	0.000
MTN	− 3.174	− 4.208	0.000
NE	− 0.606	− 1.829	0.067
SE	− 1.919	− 4.635	0.000
LOGL	− 286.93		
LR	101.52		

Note: See the notes of table 1 for variable definitions and other explanations. Where LOGL is the loglikelihood function evaluated at the maximum, and *LR* is the likelihood ration statistic for H_{01}: CONSTANT = slope coefficients for the five regional dummies = 0. It is distributed as χ^2_6 and has a 95% critical value of 12.59.

correctly. In the holdout set, the logit model predicted correctly 48% of the defaulting loans, whereas the RBFN had a correct prediction rate of 52%. Out of the 282 nondefaulting loans the logit model predicted 95% correctly, whereas the RBFN predicted 92% correctly. Further analysis of the incorrect predictions of the RBFN revealed that for both model 1 and model 2 the defaulting loans incorrectly classified as nondefault were also misclassified by the logit model.

5. Sensitivity Analysis

The selection of input variables is another issue of interest. There is no statistical methodology (such as *t*-ratios, for instance) to help us discriminate between good and bad input variables in nonlinear models. This is a drawback of neural network techniques. However, instead of statistics for each variable, one often uses sensitivity analysis. All inputs are set to their average level—that is, they describe the average loan in the data set. One of the inputs is varied, and the set is fed through the network. Plotting the output (the

Table 5. Model 1: logit and RBFN.

	Correct Classifications		
	Actual	Logit	RBFN
Training set			
Default	150	40 (27%)	56 (37%)
Nondefault	1,162	1,141 (98%)	1,131 (97%)
Total	1,312	1,181 (90%)	1,187 (90%)
Holdout set			
Default	25	7 (28%)	12 (48%)
Nondefault	282	275 (98%)	270 (96%)
Total	307	282 (92%)	282 (92%)

column of the output designated for defaulted loans) offers insights into the responsiveness of prediction to changes in that input.

Figure 2 plots the probability of default for the average loan as the market-to-loan value is varied for the average loan. As it can be seen from the graph, the RBFN shows higher likelihood of default for the average loan than the logistic function does. Figures 3 and 4 show the sensitivity of default for the average loan over different values of the coupon rate and the debt coverage ratio, respectively. It becomes evident that one can explore other sections of the estimated mapping. For instance, one may want to see how contracts in the Northeast are affected by changes in the borrower type or changes in the loan term, or one may want to focus on a particular type of loan and see for example how the average defaulted loan behaves.

A related method of analysis often used is that of geometrically representing the regions in the input space that lead to default. Although graphing all twenty variables is impossible, by using two or three variables it is possible to see which combinations of

Table 6. Model 2 logit and RBFN.

	Correct Classifications		
	Actual	Logit	RBFN
Training set			
Default	150	65 (43%)	98 (65%)
Nondefault	1,162	1,132 (97%)	1,108 (95%)
Total	1,312	1,197 (91%)	1,206 (92%)
Holdout set			
Default	25	12 (48%)	13 (52%)
Nondefault	282	268 (95%)	259 (92%)
Total	307	280 (91%)	272 (89%)

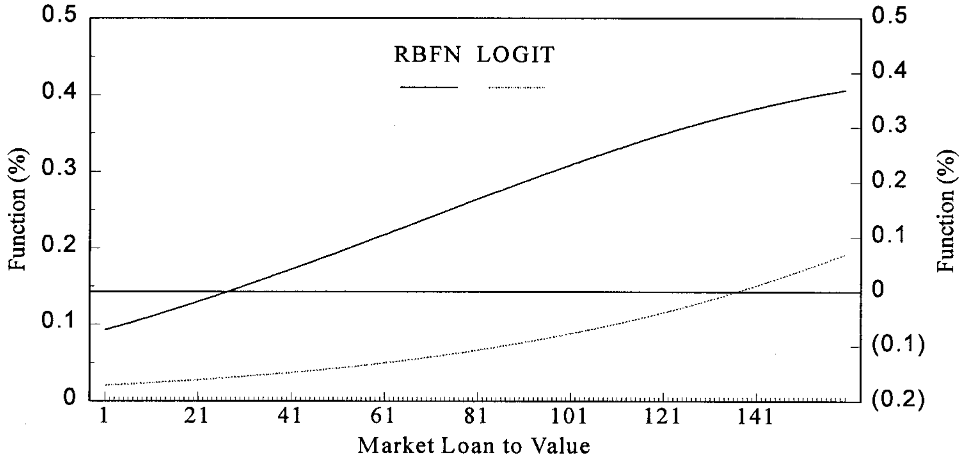

Figure 2. Probability to default for changes in market loan to value.

input levels are associated with foreclosure. This idea is useful in constructing intuitive rules of thumb for deciding on a new application for a loan.

6. Improving Prediction Performance

There are cases where practitioners are interested exclusively in predicting default rates and where the risk of incorrect prediction of a nondefault loan is not important. This section sketches how one can improve the success rate on foreclosed loans.

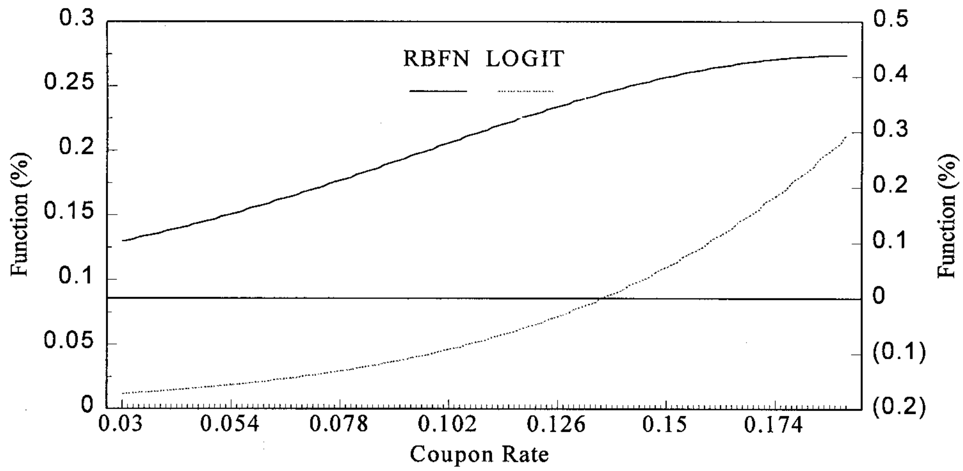

Figure 3. Probability to default for changes to the origination coupon rate.

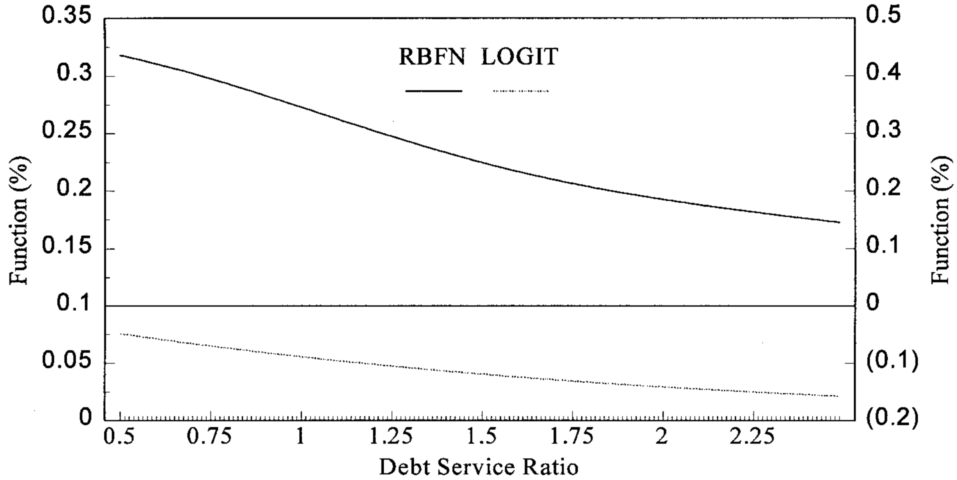

Figure 4. Probability to default for changes in debt service ratio.

As the previous sections have shown, it is not easy achieving a good prediction rate. This is a result of the existence of a small percentage of defaulted loans in the sample, along with the censoring problem. Given the data set and variables, a practitioner can do one of the following: (1) increase the weight of the defaults in the likelihood function or (2) replicate the defaulted loan records and insert them back in the original sample. Either of these methods enhances the presence of the defaulted loans and brings out the characteristics of these firms. Thus the estimation techniques can perform better. In particular, the neural net can learn the features of these records better and can classify them more accurately.

This practice has the unwelcome side effect of reducing the success rate of predicting ''good'' loans. In fact, as the replication of the defaults is repeated, the type II error gets higher. However, one can construct a loss function and optimize it to choose the right degree of replication of the defaulted records. To demonstrate the effect of the method, we have doubled the defaulted loan records in both the training and the holdout set, arriving at 163 and 186 records, respectively. This doubling is performed as follows. Each default record is copied once (and only) and inserted back into the sample. In other words, if the old sample contained just two loan records, $\{n, d\}$, denoting nondefault and default, the new sample would contain $\{n, d, d\}$. Table 7 shows the estimated logit. An RBFN was trained on the new data. Table 8 shows the results from the application of the logit and radial basis functions on the new data sets (model 3). The success rate has increased for the default loans, although it has decreased for the nondefault ones for both neural net and logit.

7. Concluding Comments

Default on commercial mortgages is a serious problem for government agencies and brokerage houses. This article has used an extensive database of individual loan histories to demonstrate the use of radial basis function networks in the prediction of mortgage default and to compare the network's prediction performance with that of the logit benchmark. It is shown that artificial neural networks are superior to logit in terms of success rate. The superiority of RBFNs is due to their ability to learn and capture nonlinearities in the data. In other words, RBFNs can seek and approximate the unknown mapping from the space of features (loan characteristics) to the space of categorical variables signifying default and nondefault. Logit parameterizes the unknown function with as many coefficients as independent variables. In contrast, an RBFN with N pattern nodes estimates as many as $6N$ parameters, leading to a more complex structure. If additional variables that are loan-specific are included as inputs, the results could be further improved.

Artificial neural networks are a new methodology in finance. Financial data contain noise, and the relationships among variables are often not stable over time or across sections. ANNs should be able to offer an edge in picking up intricate patterns. However, before neural networks enter the mainstream, we have to rely on established techniques to provide benchmarks for forecasting purposes. Statistical tests is also a much needed

Table 7. Model 3: logit estimated on the enhanced sample.

Variable	Parameter Estimate	*T*-Statistic	*P*-Value
CONSTANT	– 3.214	– 2.016	0.044
MDMLTV90	0.018	3.133	0.002
INTRTE	16.876	2.281	0.023
ORIGDSC	– 0.464	– 0.569	0.569
PROPAPT	1.970	5.548	0.000
PROPHOTL	0.914	1.233	0.218
PROPOFF	1.508	4.304	0.000
PROPINDU	1.187	2.760	0.006
PROPOTH	1.866	2.406	0.016
BRWRTY01	0.457	1.105	0.269
BRWRTY02	0.282	0.854	0.393
BRWRTY04	0.202	0.562	0.574
GPMDUMMY	1.033	4.141	0.000
GPMDUMY2	– 2.159	– 4.340	0.000
ACRDUMMY	0.476	1.072	0.283
AMZDUMY2	– 1.967	– 7.249	0.000
NEC	– 2.476	– 6.197	0.000
ME	– 3.519	– 6.057	0.000
MTN	– 3.046	– 5.558	0.000
NE	– 0.701	– 2.648	0.008
SE	– 1.852	– 5.845	0.000

Note: See the notes of table 2 for variable definitions and other explanations.

Table 8. Model 3: logit and RBFN predictions for enhanced sample.

	Correct Classifications		
	Actual	Logit	RBFN
Training set			
Default	300	178 (59%)	226 (75.33%)
Nondefault	1,162	1,103 (95%)	1,080 (92.94%)
Total	1,462	1,281 (88%)	1,306 (89%)
Holdout set			
Default	50	32 (64%)	34 (68%)
Nondefault	282	263 (93%)	248 (88%)
Total	332	295 (89%)	282 (85%)

component to most nonlinear models, including neural networks, and the literature awaits such developments. Of course, the ultimate test is the quality of a technique's results, and that is the bottom line in practice.

Acknowledgments

The authors wish to thank the Travellers Realty Investment Company for providing the data for the project; Smith Barney, Inc., National Bank of Greece, Michael Youngblood, and Ranzjit Fernando for support on this project; and Professor Ronald Anderson of the Universite Catholique de Louvain and two anonymous reviewers for their comments and helpful suggestions. The paper was presented at the Eastern Finance Association (April 1996, Charlotte, NC), the annual conference of the American Real Estate Society (March 1996, Lake Tahoe, CA), and the Third International Conference on Neural Networks in the Capital Markets (October 1995, London).-

Notes

1. The states included in each region are given below:
East North Central: Illinois, Indiana, Michigan, Ohio, and Wisconsin;
Mideast: Delaware, District of Columbia, Maryland, New Jersey, New York, Pennsylvania;
Mountain: Colorado, Idaho, Montana, Nevada, Utah, Wyoming;
New England: Connecticut, Main Massachusetts, New Hampshire, Rhode Island, Vermont;
Pacific: Alaska, California, Hawaii, Oregon, Washington;
Southeast: Arkansas, Florida, Georgia, Louisiana, Mississippi, North Carolina, South Carolina, Tennessee, Virginia, West Virginia;
Southwest: Arizona, New Mexico, Oklahoma, Texas;
West North Central: Iowa, Kansas, Minnesota, Missouri, Nebraska, North Dakota, South Dakota.
2. We have chosen to compare the forecasting accuracy of RBFN versus the logistic approach for two reasons: (1) the logistic approach is still the most widely used approach that is used by practitioners for classification purposes; (2) both approaches model unconditional default probabilities. Comparison of a neural network

with a model that employs conditional probabilities, such as the proportional hazard model, is outside the scope of this article.

3. An important variable that would explain the incidence of default is the volatility of the property price. Actually, this variable would affect both the default option as well as the prepayment option. However, volatility data are not available to us.
4. The optimal number of pattern units is a purely empirical matter. A method similar to *cascade correlation algorithm* can be used to sequentially add nodes to the point where no improvements are made. Genetic algorithms can also be used. Unfortunately, these methods impose a huge computational burden with no apparent gain except economizing somewhat on the coding of the network. Instead, a search alternative was pursued in this article by trying different architectures and assessing their prediction performance.

References

American Council of Life Insurance. (1995). ''Quarterly Survey of Mortgage Loan Delinquencies and Foreclosures,'' *Investment Bulletin*, 1289 (March 31).

Berndt, E. K., B. H. Hall, R. E. Hall, and J. A. Hausman. (1974). ''Estimation and Inference in Nonlinear Structural Models,'' *Annals of Economic and Social Measurement* 3/4, 653–665.

Collins, E., S. Ghosh, and C. Scofield. (1988). ''An Application of Multiple Neural Network Learning System to Emulation of Mortgage Underwriting Judgements,'' *Proceedings of the IEEE International Conference on Neural Networks* 2, 495–566.

Dutta, S., and S. Shekhar. (1988). ''Bond Rating: A Non-Conservative Application of Neural Networks,'' *Proceedings of the IEEE International Conference on Neural Networks* 2, 443–450.

Foster, C., and R. Van Order. (1984). ''An Option-Based Model of Mortgage Default,'' *Housing Finance Review* 13, 351–372.

Foster, C., and R. Van Order. (1985). ''FHA Terminations: A Prelude to Mortgage Pricing,'' *Journal of the American Real Estate and Urban Economics Association* 13(3), 292–316.

Hutchinson, J., A. Lo, and T. Poggio. (1994). ''A Nonparametric Approach to Pricing and Hedging Derivative Securities Via Learning Networks,'' *Journal of Finance* 49, 851–889.

Kau, J. B., D. C. Keenan, W. J. Muller III, and J. T. Epperson. (1990). ''Pricing Commercial Mortgages and Their Mortgage-Backed Securities,'' *Journal of Real Estate Finance and Economics* 2, 333–356.

Leonard, J. A., and M. A. Kramer. (1991). ''Radial Basis Functions for Classifying Process Faults,'' *IEEE Control Systems* (April).

Moody, J., and C. J. Darken. (1989). ''Fast Learning of Locally Tuned Processing Units,'' *Neural Computation* 1, 281–294.

Quigley, J., and R. Van Order. (1994). ''Explicit Tests of Contingent Claims Models of Mortgage Default,'' Working paper, University of California, Berkeley.

Refenes, A. (ed.). (1995). *Neural Networks in the Capital Markets*. New York: Wiley.

Reilly, D. L., E. Collins, C. Scofield, and S. Ghosh. (1991). ''Risk Assessment of Mortgage Applications with a Neural Network System: An Update as the Test Portfolio Ages,'' *Proceedings of the IEEE International Conference of Neural Networks* 2, 479–482.

Tam, K. Y., and Y. M. Kiang. (1992). ''Managerial Applications of Neural Networks: The Case of Bank Failures,'' *Institute of Management Sciences* 38(7), 926–947.

Titman, S., and W. Torrus. (1989). ''Valuing Commercial Mortgages: An Empirical Investigation of the Contingent-Claims Approach to Pricing Risky Debt,'' *Journal of Finance* 44(2), 345–373.

Vandell, Kerry D. (1992). ''Predicting Commercial Mortgage Foreclosure Experience,'' *Journal of the American Real Estate and Urban Economics Association* 20(1), 345–373.

Vandell, K. D., W. Barnes, D. Hartzell, D. Kraft, and W. Wendt. (1993). ''Commercial Mortgage Defaults: Proportional Hazard Estimation Using Individual Loan Histories,'' *Journal of the American Real Estate and Urban Economics Association* 21(4), 451–480.

White, H., A. Gallant, K. Hornik, M. Stinchcombe, and J. Wooldridge. (1992). *Artificial Neural Networks: Approximation and Learning Theory.* Oxford: Blackwell.