The Quantification of Model Risk According to the Principle of Relative Entropy with a Case Study

This is a summary of a research study available at https://www.michaeljacobsjr.com/wp-content/uploads/2023/07/Jacobs_ModelRiskQuant_20230701_V4.pdf.

Name: Michael Jacobs, Jr.
Company: PNC Financial Services Group

Bio: Michael Jacobs, Jr., Ph.D., CFA, Senior Vice-President, and Lead Modeling Expert, PNC Financial Services Group, Model Development Group, 340 Madison Avenue, New York, N.Y., 10173, 917-324-2098, michael.jacobsjr@pnc.com. The views expressed herein are those of the authors and do not necessarily represent a position taken either by PNC Financial Services group or of any affiliated firms.

1 Introduction

Risk measurement relies on modeling assumptions, the errors in which expose such models to model risk. This article introduces and applies a tool for quantifying model risk and making risk measurement robust to modeling errors. As simplifying assumptions are inherent to all modeling frameworks, the prime directive of model risk management is to assess vulnerabilities to and consequences of model errors. This article presents a study summary that is consistent with this objective in model risk measurement, focusing on calculating bounds on measures of loss, which can result in a range of model errors within a certain distance of a nominal model for a range of alternative models. To this end, it is proposed to quantify such changes according to the *principle of relative entropy*. Illustrating the application of this principle the measure for corporate *probability-of-default* (PD) is the *Aikakie Information Criterion* (AIC).

2 Discussion and Case Study

In the building of risk models professionals are subject to errors from model risk; one source being the violation of modeling assumptions. This can be addressed by applying a methodology for the quantification of model risk by using a tool in building models robust to such errors. A key objective of model risk management is to assess the likelihood, exposure, and severity of a model error in that all models rely upon simplifying assumptions. For this reason, a critical component of an effective model risk framework is the development of bounds upon a model error resulting from the violation of modeling assumptions. This measurement is based upon a reference nominal risk model and is capable of rank ordering the various model risks, as well as indicating which perturbation of the model has a maximal effect upon some risk measure.

In line with the objective of managing model risk in the context of measuring and managing risk in various contexts (e.g., credit or market risk, portfolio management), confidence bounds are calculated around some measure of risk (or loss) spanning model errors in a vicinity of a nominal or reference model defined by a set of alternative models. These bounds can be likened to confidence intervals that quantify sampling error in parameter estimation. However, these bounds are a measure of model robustness that instead measures model error due to the violation of modeling assumptions. In contrast, a standard error estimate conventionally employed in risk modeling does not achieve this objective. The latter construct relies upon an underlying probability generating distribution of errors. It should be noted that in applying relative entropy to model risk measurement, this assumption is not needed but rather a test of whether this assumption is valid.

Meeting the previously stated objective in the context of modeling is achieved by bounding a measure of loss that can, within reason, reflect a level of model error. It's observed that while amongst practitioners, one alternative means of measuring model risk is to consider challenger models, an assessment of estimation error or sensitivity in perturbing parameters is a more prevalent means of accomplishing this objective, which captures only a very narrow dimension of model risk. In contrast, our methodology transcends the latter aspect to quantify potential model errors, such as incorrect specification of the probability law governing the model without assuming that it is correct.

As these types of model errors under consideration all relate to the likelihood of such error, which is connected to the perturbation of probability laws governing the entire modeling construct, the principle of relative entropy is applied. In Bayesian statistical inference, relative entropy between a posterior and a prior distribution is a measure of information gain when incorporating incremental data. In the context of quantifying model error, relative entropy has the interpretation of a measure of the additional information required for a perturbed model to be considered superior to a champion or null model. Said differently, relative entropy may be interpreted as measuring the credibility of a challenger model. Another valuable feature of this construct is that within a relative entropy constraint, the so-called *worst-case alternative* (e.g., in this case, a divergence in the distributions of a loss estimate between the models due to ignoring some feature of the alternative model) can be expressed as an *exponential change of measure*.

Illustrating this application to alternative types of corporate PD models, a comparison of loss distributions is performed of 1- and 3-dear default horizon *through-the-cycle* (TTC) models suitable for credit underwriting, vs. *point-in-time* (PIT) models used in early warning systems. These models are built from a long history of borrower-level data sources from Moody's, COMPUSTAT, and CRSP. The quantification of model risk is performed with respect to the following modeling assumptions:

- Omitted variable bias - leaving out the Merton *distance-to-default* (DTD) risk factor.
- Misspecification according to neglected interaction effects.
- Misspecification according to an incorrect link function - the *Cumulative Log-Log* (CLL) as opposed to the Logit.

Distributions of the relative *proportional deviation in AIC* (RPD-AIC) from the base specifications through a simulation exercise is developed. It is observed that omitted variable bias in relation to DTD results in the highest model risk, an incorrectly specified link function has the lowest measured model risk, and neglected interaction effects are intermediate in the quantity of model risk (see the table below).

Type of Model	Model Specification	Model Assumption	Min.	25th Prcntl.	Median	Mean	75th Prcntl.	Max.	Std. Dev.
Through-the-Cycle	Model 1	Omitted Variable Bias	0.0093	0.1137	0.2009	0.2290	0.3208	0.8328	0.1461
		Neglected Interaction Effects	0.0221	0.1116	0.1626	0.1759	0.2267	0.5262	0.0861
		Incorrectly Specified Link Function	0.0134	0.0721	0.0960	0.1005	0.1233	0.2714	0.0380
	Model 2	Omitted Variable Bias	0.0079	0.1010	0.1746	0.1962	0.2687	0.7362	0.1251
		Neglected Interaction Effects	0.0081	0.0830	0.1203	0.13389	0.1719	0.5239	0.0699
		Incorrectly Specified Link Function	0.0158	0.0606	0.0821	0.0866	0.1077	0.24061	0.03541
Point-in-Time	Model 1	Omitted Variable Bias	0.0044	0.0816	0.1306	0.1759	0.2149	0.5528	0.0995
		Neglected Interaction Effects	0.0123	0.0572	0.0876	0.0978	0.1266	0.4128	0.0543
		Incorrectly Specified Link Function	0.0062	0.0352	0.0486	0.0635	0.0685	0.1783	0.0256
	Model 2	Omitted Variable Bias	0.0113	0.0873	0.1414	0.1587	0.2118	0.5911	0.0945
		Neglected Interaction Effects	0.0033	0.0500	0.0765	0.0869	0.1131	0.3436	0.0505
		Incorrectly Specified Link Function	0.0077	0.0304	0.0414	0.0461	0.0580	0.1621	0.0222

3 Conclusions

Consistent with this objective in model risk measurement, we calculated bounds on measures of loss that can result over a range of model errors within a certain distance of a nominal model for a range of alternative models. This was accomplished by quantifying such changes according to the principle of relative entropy. The application of this principle was illustrated through a case study where the measure of loss in models for corporate PD considering alternative use cases is the AIC.

There are various implications for model development and validation practice, as well as supervisory policy which can be gleaned from this analysis. First and foremost, this exercise of simulating a loss measure shows we should exercise caution in over-reliance on measures of model fit derived from a single historical dataset. Even if out-of-sample performance is favorable, there could be an unpleasant surprise when adding to the reference datasets when re-estimating the models. Second, from a fitness-for-purpose perspective, it is a better practice to consider the use case for any credit model in establishing the model design.

The proposal for measuring model risk better supports this objective in contrast with other approaches. Considering the observations and contributions to the literature, applying the principle of relative entropy provides valuable guidance to model development, model validation, and supervisory practitioners. Additionally, the discourse has contributed to resolving the debates around which class of credit model is best fit for purpose in large corporate PD modeling applications. This better performance is manifested broadly as both an improved fit to the data and lower measured model risk due to model misspecification.