Applied Robust Performance Analysis for Actuarial Applications: A Guide

Blanchet, J., Lam, H., Tang, Q., and Yuan, Z.

Introduction

- The goal of this project is to provide systematic tools for the quantification of model error (or model misspecification) in actuarial risk analysis.

- The systematic approach proceeds as follows:
  - The methodology starts from a baseline model which was calibrated using an arbitrary procedure.
  - Using the baseline model, we compute the worst case risk among all possible models which are within a "plausible" distance of the baseline model.
  - The methodology can be implemented using Monte Carlo experiments.

A Toy Example: Illustrating the Impact of Model Error

- A potential loss, X, is *assumed* to be distributed Exponential.
- The actuary *estimates* the mean to be 1 (say $1M=1 million).
- Capital requirement = b is computed to withstand losses with .995 probability.
- *Solve for b so that P(X>b)=exp(-b)=.005, yielding b≈5*
- BUT: X *was only ASSUMED to be* exponentially distributed.
- Section 5.1, equation (16) uses the methodology to show that *b≈20* in this example reflects the quantification of model error.
- So, what is the methodology and how to use it? Can we add certainty to the model to reduce b? If so, how?

The Methodology: What's the Main Idea?

- How does the methodology work?
- **Ans:** We want to estimate

$$E_{true}(h(X))$$

The expectation E_{true} is unknown. We consider a baseline model P_0 (like the exponential on the previous page), which is convenient (maybe because of tractability).

We then evaluate the following optimization problem:

$$\max\{E(h(X)): over\ P\ such\ that\ D(P||P_0) \leq \delta\}.$$

In words, we maximize (or minimize depending on the context) the worst case expectation and the optimization is performed over ALL models which differ from P_0 by an amount less than δ.

We call this optimization problem the "Basic Distributionally Robust" (BDR) formulation.

Section 3 of the paper explains how find the worst case P maximizing the expectation in the BDR formulation and also the value of maximization problem.

The Methodology: Comparing Models…

- What does it mean to optimize over all models?
- **Ans:** This means that we must assess the difference between two models, say P and P_0, using a criterion which is non-parametric (because we don't want to make specific assumptions on alternative models). This is why choosing $D(P||P_0)$ in terms of the Kullback-Leibler divergence is useful. The definition is given in page 4 of the paper.
- Intuitively, as explained in Section 5.2 Kullback-Leibler reweights the probabilities to favor outcomes that have higher adverse impact in the expectation of interest.
- Other notions of discrepancy can also be used. Kullback-Leibler is advantageous because it has been extensively studied in engineering and economics.
- The methodology we advocate has its roots in the robustness approach introduced by the Nobel prize winners Hansen & Sargent.

The Methodology: The Size of the Feasible Region of Models...

- How does one select δ?

- **Ans:** This is the hardest part of the procedure. The advantage is that *at least we are not aiming at pinning down the true model, but only finding δ such that $D(P_{true}||P_0) \leq \delta$.*

- In Section 6.1, equation (19) we establish a connection between this method and Empirical Likelihood. If there are n observations used for non-parametric inference, then $2n\delta$ should be the 95% quantile of a chi-squared distribution with 1 degree of freedom (because there is only one expectation to estimate, the one appearing in the objective function of the optimization problem).

- More generally, we also discuss how to calibrate δ in Section 6.2.

The Methodology: How to Improve the Bounds?

- But what if my bound just "feels" too pessimistic?
- **Ans:** The ONLY way to get too pessimistic bounds is if there is a significant possibility for model misspecification or model error. Observe that if δ=0 we recover model P_0, and implicitly we are assuming that there is NO model error.
- The ONLY way then to reduce the bound is by adding available information on the true model.
- For example, if we know the additional information that $E_{true}g_1(\mathrm{X}) = a_1,...,$ $E_{true}g_m(\mathrm{X}) = a_m$ then solve instead the expanded optimization problem:
$$\max\{E(h(X)): D(P||P_0) \leq \delta, Et_{ruegi}(X) = ai, 1 \leq i \leq m\}.$$
- Now 2nδ is chosen as the 95% quantile of a chi-squared distribution with (m+1) degrees of freedom (again due to the connection to empirical likelihood).

What About the Impact of Dependence?

- Can we use the methodology to assess the impact of dependence?

- **Ans:** Yes, a situation that arises in practice is when one is interested in E_{true}h(X,Y), for a pair of risks (X and Y), say X and Y represent the time-until-death of a couple. The mortality of X is well understood marginally, and the same is true for Y, but the joint mortality might be less certain. This example is considered in Section 4.1 in the paper.

- The optimization problem has fixed marginal distributions, thus only the dependence structure is to be optimized.

How Does One Solve the Optimization Problem?

- How to solve the optimization problem?

- **Ans:** As explained in Section 3 of the paper, the Basic Distributionally Robust formulation (even with additional information to improve bounds) leads to a convex optimization problem. This is important because there is widely available computational packages (many of them free) which can be used to solve convex optimization problems. In the case of discrete distributions with finite support the optimization problem is "standard" and we provide a quick summary of the relevant aspects (for what we study here) of the Karush-Kuhn-Tucker conditions in the paper.

- The general case, in which the models possess density and therefore are not discrete, leads to a infinite dimensional optimization problem which falls under calculus of variations theory. The results are completely analogous to the discrete case and they are also discussed in the paper (see Section 5).

Can One Solve These Optimization Problems Practically?

- Can one solve these problems say in Excel?

- **Ans:** Yes, these problems can be solved by Monte Carlo sampling in Excel. We have implemented several examples to illustrate how to do this. A companion Excel file has been submitted with the paper, with worked out examples, which are explained in Sections 5.1, 7, and 8.1. The examples include the evaluation of Conditional Value at Risk and examples involving the use of t-copulas.