Accepted Manuscript

Robust empirical optimization is almost the same as mean–variance optimization

Jun-ya Gotoh, Michael Jong Kim, Andrew E.B. Lim

PII: S0167-6377(17)30514-X
DOI: https://doi.org/10.1016/j.orl.2018.05.005
Reference: OPERES 6367

To appear in: *Operations Research Letters*

Received date : 7 September 2017
Revised date : 23 May 2018
Accepted date : 23 May 2018

Robust empirical optimization is almost the same as mean-variance optimization

Jun-ya Gotoh[a], Michael Jong Kim[b], Andrew E.B. Lim[c]

[a]*Department of Industrial and Systems Engineering, Chuo University, Japan. Email: jgoto@indsys.chuo-u.ac.jp*
[b]*Sauder School of Business, University of British Columbia, Canada. Email: mike.kim@sauder.ubc.ca*
[c]*Department of Analytics and Operations & Department of Finance, NUS Business School, National University of Singapore, Singapore. Email: andrewlim@nus.edu.sg*

Abstract

We formulate a distributionally robust optimization problem where the deviation of the alternative distribution is controlled by a ϕ-divergence penalty in the objective, and show that a large class of these problems are essentially equivalent to a mean-variance problem. We also show that while a "small amount of robustness" alway reduces the in-sample expected reward, the reduction in the variance, which is a measure of sensitivity to model misspecification, is an order of magnitude larger.

Key words: Robust empirical optimization, mean-variance optimization, data-driven optimization, relative entropy, ϕ-divergence, regularization, bias-variance trade-off.

1. Introduction

Empirical optimization (or sample average approximation) is a data-driven approach to decision making where decisions are chosen by optimizing the (in-sample) expected reward under the empirical distribution. While it has some desirable asymptotic properties, out-of-sample performance can be poor when the sample size is not sufficiently large or the statistical assumptions that form the basis of empirical optimization are not valid, and robust (worst-case) empirical optimization has emerged as one approach to addressing these limitations [2, 5, 6, 12, 15, 21]. We add to this literature by (i) providing insight into *how* worst-case empirical optimization comes up with decisions with lower sensitivity to model misspecification for a broad class of models, (ii) characterizing the impact of robustness on the solution of the robust problem, and (iii) quantifying the impact of "robustifying" an empirical optimization problem on the mean and variance of the reward.

Specifically, we show that a large class of robust empirical optimization problems are (almost) equivalent to an empirical mean-variance problem with the equivalence becoming exact in the regime of vanishing model uncertainty. Intuitively, a decision is susceptible to model misspecification if a small mismatch between the in- and out-of-sample reward distributions results in a large difference between the respective expected rewards, which can occur if the in-sample distribution has a large spread and there is a mismatch between the tails of the two distributions. The mean-variance relationship shows that robust optimization reduces sensitivity to misspecification by favoring decisions that reduce the spread of the