Journal Of Investment Management, Vol. 18, No. 2, (2020), pp. 1–4
© JOIM 2020
www.joim.com

PRACTITIONER'S DIGEST

The "Practitioner's Digest" emphasizes the practical significance of manuscripts featured in the "Insights" and "Articles" sections of the journal. Readers who are interested in extracting the practical value of an article, or who are simply looking for a summary, may look to this section.

ON THE STABILITY OF MACHINE LEARNING MODELS: MEASURING MODEL AND OUTCOME VARIANCE PAGE 9

Vasant Dhar and Haoyuan Yu

The appeal of a systematic approach in finance, relative to being discretionary, is repeatability. With a handful of exceptions, human performance tends not to be repeatable. In contrast, a well-constructed systematic strategy holds the promise of repeatability if it is well designed.

Machine Learning offers a new way of discovering insights from data, where the machine plays an active role in discovering what is repeatable. But investing that involves holding overnight positions is subject to significant doses of randomness, driven by political, natural, and economic events. Because of this, there is an inherent level of uncertainty that must be expected in the knowledge that is discovered by the learning algorithm. Portfolio managers often wonder why reality doesn't match the backtest, despite the utmost care in following the scientific method to avoid overfitting. With a machine learning approach, it is particularly important to measure the stability of models learned from data as an input to automated model selection. Understanding the variance in behavior of learned models is essential for making this decision.

This article quantifies the uncertainty in decision making and performance the operator of a machine learning based algorithm should expect as a function of the type of problem. This type is defined along two axes. The first is the predictability of a problem, which varies between randomness and determinism. The second axis is the "base rate" of the phenomenon of interest. How often does it occur? The less frequently it tends to occur, the harder it is to predict correctly all else being equal. The two axes taken together form the basis for inquiry. We constructed datasets corresponding to discrete points on this grid for which we record two types of variance resulting from the learning process. The first is what we call model variance, which measures the instability of its decisions. The second is outcome variance in performance.

The experiments were motivated by our experience using machine learning algorithms on autopilot over the last ten years. The reasons for why the programs behaved the way they did, for example, their

Not for Distribution

ratio of long and short trades only became apparent after years of observation, which we ultimately linked to the two dimensions we varied in our experiments.

The results should be of interest to researchers and practitioners. For researchers, they quantify the variance expectations from machine learning algorithms for any kind of problem. For operators of systematic programs, they quantify the risks associated with running machine learning based models.

CAN MACHINES "LEARN" FINANCE? PAGE 23

Ronen Israel, Bryan Kelly and Tobias Moskowitz

We show that liquidity is important in describing stock returns by constructing a liquidity-based model that outperforms leading benchmark models in describing the most well-known anomalies. We use both earnings and liquidity factors to construct this competing model, both of which are highly intuitive to practitioners. Stock trading volume, as a measure of liquidity has been reported in conjunction with stock prices in newspapers at least since the 1800's. Today, institutional investors and index providers adjust weights by free float; they already act according to their belief that market cap weighting should be adjusted by investors' ability to freely trade their positions. Unlike some other weighting schemes, earnings and liquidity factors are readily amenable to asset pricing theory; investors desire to hedge their own earnings and liquidity state risk along with Sharpe's (1964) market risk. Fama and French (1995) themselves have discussed their use of HML as a proxy for an earnings state variable. We use a direct earnings variable.

The notion of replacing market cap with a measure of liquidity such as trading activity has the potential to materially impact asset allocation. Much of the market is characterized by large and small relative to market capitalization. A more practical measure for an investor who wants to buy a or sell a position is the asset's liquidity, and we use the well-known measure of stock's trading volume. Existing retail products can be adjusted to better capture this measure and new products can be constructed with a better alignment to whether the stock is investible.

We also collect and demonstrate state-of-the-art statistical tests for practitioners to use to gauge model performance. These include comparisons of factor models in terms of GRS-statistics, cross-sectional r-squared statistics, and the maximum squared Sharpe ratio for the intercept's statistic, as well as the average over the absolute values of the intercepts. For fair comparison, we reconstruct and test three separate factor models on the same universe on their ability to explain a set of portfolio anomalies over a long time period, and over a rolling time period.

DYNAMIC GOALS-BASED WEALTH MANAGEMENT USING REINFORCEMENT LEARNING PAGE 37

Sanjiv R. Das and Subir Varma

Several large global financial firms are engaged in managing wealth in retirement portfolios amounting to around $45-50 trillion [2019 ICI FactBook]. Almost all of this money is *statically managed*, i.e., it is allocated to assets to optimize risk versus return one year at a time, taking into account the risk tolerance

Not for Distribution

of the investor. The state of the art here is "life-cycle" funds, i.e., funds that dial down risk as the retiree grows older and this a fixed and clearly non-optimal approach. It does not account for the fact that different people have different goals and risk preferences. So why have wealth management companies not moved to a *dynamically* optimal approach to manage retirement portfolios? Mostly, because they have not had tools for solving this problem at scale, and also because they have not had to. Such firms have always collected healthy revenue from retirement portfolio fees. Now, with the introduction of robo-advisors, the fees these companies collect are being threatened and the entire industry is experiencing large-scale fee compression. These firms will be looking to differentiate themselves by using dynamically optimal solutions to the problem. Reinforcement learning is an important approach that is being developed to solve large-scale dynamic retirement problems and this paper offers an introduction with an example implementation to showcase how this is done.

USING MACHINE LEARNING TO PREDICT REALIZED VARIANCE PAGE 57

Peter Carr, Liuren Wu and Zhibai Zhang

Forecasts of variance are needed for risk assessment, which in turn affects portfolio composition. A well known forecast of S&P500 variance is the square of VIX, which is the value of a particular long-only semi-static portfolio of near-dated SPX option prices. When squared VIX is used to forecast S&P500 variance over a 30 day horizon, the SPX option prices are combined in a linear fashion.

One of the main attractions of machine learning algorithms is their potential to both improve linear forecasts and to synthesize information in non-linear way. In this paper, we investigate whether or not using machine learning algorithms improves our ability to forecast the variance of S&P500 over a 30 day horizon. In particular, we employ ridge regression, feed forward neural networks, and random forests directly on SPX option prices.

We also employ these techniques on the residuals from a squared VIX forecast in an effort to combine machine learning and human understanding represented by (squared) VIX. Our results favor this hybrid approach suggesting that artificial intelligence can supplement human intelligence, but is not yet ready to replace human intelligence for this forecasting problem.

LOCAL, GLOBAL, AND INTERNATIONAL CAPM: FOR WHICH COUNTRIES DOES MODEL CHOICE MATTER? PAGE 73

Demissew Ejara, Alain Krapl, Thomas J. O'Brien and Santiago Ruiz de Vargas

Modern financial markets are internationally integrated to the point that an international CAPM is conceptually superior to the traditional "local" CAPM, which in principle is appropriate only for a segmented financial market. Moreover, an international CAPM that includes currency risk is conceptually superior to one that ignores it. However, to a practitioner who wants to use a risk-return model to estimate a discount rate for a valuation analysis, the local CAPM is easier to apply than an international CAPM, and an international CAPM that ignores currency risk (termed a global CAPM) is easier to apply than one that includes it.

Not for Distribution

Because the effort needed to apply each model varies, it is relevant for practitioners to know how much difference the model choice makes in discount rate estimates. This empirical study shows that the model choice tends to make a small difference for some countries and a large difference for others. Therefore, practitioners in some countries can apply an easier model and estimate discount rates that tend to reasonably approximate those of an international CAPM that includes currency risk. Practitioners in countries where model choice makes a substantial difference should beware that applying an easier model may result in substantial errors in discount rate estimates for valuation.

Not for Distribution