Risk, uncertainty and investment decisions

By CIO Mikkel Petersen, AI Alpha Lab ApS

"Uncertainty is an uncomfortable position. But certainty is an absurd one."

Voltaire

Is it risk or is it uncertainty

Frank Knight was an idiosyncratic economist who formalized a distinction between risk and uncertainty in his 1921 book, "Risk, Uncertainty, and Profit". As Knight saw it, an ever-changing world brings new opportunities for making profitable investment decisions, but it also means we have imperfect knowledge of future events. Knight distinguished between two types of uncertainty, risk and genuine uncertainty. According to Knight, risk applies to situations where we know the potential outcomes in advance and where we can accurately measure the odds associated with the individual outcomes.

An example of risk is rolling a pair of dice. Before we roll, we know in advance what the odds are for each possible outcome (provided that the dice are fair). When we engage in a game of dice, one can make optimal decisions with regards to the outcome space, since the total future outcome space is known in advance as well as the associated probabilities. However, one face the risk that adverse random outcomes will materialize over finite time horizons.
In any game of chance, in the short run, we are always faced with the risk of the right decision leading to the wrong outcome. However, convergence to the mean (the law of large numbers) ensures that managing risk in the long run is pretty straightforward. You match up your investment to the odds of it paying off. When

uncertainty is probabilistically measurable risk, it is possible to design policies that are optimal on average or in some quantile sense.

Policy design under risk is based on first principles as expressed by economic theory and enables investors to distinguish between good, bad, winning and losing investments.

Unfortunately, most real-life problems do not just face risk, they are also facing a much trickier endeavor, uncertainty.

Uncertainty occurs when we don't know the possible outcomes in advance, let alone their probabilities. Genuine uncertainty occurs in complex systems, where lots of actors interact over time and arises from the fact that in the physical reality more things can happen than will happen.

Investment professionals has so far been limited in their decision making by the choice of model that they employ. The models have been solely based on historical risk assessments (data risk), relying on maximum likelihood estimation and have not taken on the complex task of doing proper inference by also estimating the confidence in future model performance, i.e. defining model uncertainty.

Predictive uncertainty

In order to make optimal decisions under uncertainty, one needs a set of future scenarios and the associated probability of realization. Within investing the scenario could be the return expectations of a set of stocks and to calculate the future value of an investment in any given stock, one also needs what we call predictive uncertainty.

Predictive uncertainty is the sum of two types of uncertainty, aleatoric and epistemic uncertainty. Aleatoric uncertainty or data uncertainty is the inherent uncertainty in non-stationary data that translate into any model applied on the data. This is the only type of uncertainty investors have been able to quantify until today and therefore investors have, consciously or not, made the assumption that aleatoric uncertainty can approximate total or predictive uncertainty. The approximation is only true if we further assume that the model of choice is the right model at all times, i.e. the model can be expected to perform equally well at all times. This is hardly true and therefore the approximation of historical risk equaling predictive uncertainty is flawed at best.

We need to calculate epistemic or model uncertainty in order to get a better understanding of optimal investment decision making. This is admittedly a daunting task as it involves the calculation of all possible model or parameter specifications that are able to explain a given dataset in order to find the probability that the model of choice, is the right one. So far this has not been possible, but the advent of machine learning models such as probabilistic multi-layer perceptrons, based on Bayesian thinking, has enabled us to quantify epistemic uncertainty and as such quantify predictive uncertainty. The result is that we can now make investment decisions on a much more informative ground, creating value that was not possible just 5 years ago.

The way it should be done

Below is Bayes Theorem, a way to do proper hypothesis testing or inference, when faced with uncertainty. Even though many within the fields of natural and social sciences agree that this is indeed the right way to observe the physical reality in which we live, historically there was a problem or actually two problems with the below theorem.

Likelihood
How probable is the evidence given that our hypothesis is true?

Prior
How probable was our hypothesis before observing the evidence?

$$P(H|e) = \frac{P(e|H)P(H)}{P(e)}$$

Posterior
How probable is our hypothesis given the observed evidence?

Marginal
How probable is the evidence under all possible hypotheses?

The first problem is the prior knowledge (red text). Incorporating prior knowledge required scientists to state their minds, which is not easy, and furthermore to formalize this knowledge.

The second problem is the marginal probability (blue text) which requires the estimation of all possible hypotheses or models given the observed data, an exponentially difficult estimation procedure with potentially billions of parameters to estimate.

The solution to the two problems became to naively assume that we have no prior knowledge of the problems we try to solve, removing the prior term in the numerator, and also not treat problems as probabilistic, thereby reducing the marginal probability in the denominator to just a normalizing factor, which can be removed. The only term remaining on the right side of Bayes theorem is the likelihood and as such we can infer things about the world by calculating only the maximum likelihood. But to get there we made a lot of questionable assumptions!

Maximum likelihood

Maximum likelihood basically means a procedure that best-fits the parameters of a model to a given dataset. This is problematic when the data used to fit the model is different from the data the model is supposed to explain. Dynamic data, such as most financial data, may not contain a good representation of the future and since maximum likelihood models only use data, the ability to extrapolate into the future is limited. The maximum likelihood models do not incorporate any prior knowledge that we might have and at the same time the model is forced to provide a best fit. This combination makes the models prone to finding correlations and not causality.

The key to solve this problem lies in the concept of Bayesian machine learning, which combine data, domain knowledge and probabilistic modelling. By incorporating prior knowledge of the problem at hand, coding it as a statistical distribution and then combining it with the information in data, significantly reduces the risk of fitting the model to statistical only truths. On top of that, Bayesian machine learning models are allowing the model to doubt its conclusions, forcing it to disregard spurious correlations and quantify the validity of its predictions. The combination of all this results in a model that actually has an understanding of the environment in which it operates.

In order to harvest the benefits of machine learning, we as humans need to be smart in our use of this powerful new tool. Most popular models from machine learning today, like neural nets, are merely fitting a line to existing data i.e. an advanced linear regression. Regardless of how complex or deep the network is, the network does the exact same thing. You may configure it differently to get representations in different layers and make it more effective, but you are still only looking at what is

part of your existing training sample. A simple perception model only gives you one answer that fits your model without questioning it, much like classical maximum likelihood models.
What you can do with Bayesian predictive inference machines using probabilistic programming is to give the algorithm room for doubt and the possibility for the model to disproof your assumption. In this way we minimize the risk of fitting our model to the spurious correlation in data and maximize the probability that we fit our model to the true causal dependencies corresponding to the physical reality that we live in.

The new doctrine
Mathematicians, statisticians, physicists and investment professionals have been indoctrinated with models that were never meant to solve complex real life problems. Models that relied on strict assumptions in order to fit a reality that never existed, but so far this was justifiable since there were no alternative. Today there is!

Machine learning provides us with the ability to make optimal decisions in a world of uncertainty. Today we can model the world without relying on unrealistic assumptions resulting in statistical-only results. Yes, investors are forced to step into unknown waters, educate themselves in fields not formerly required by the finance community and accept that computers can now solve certain tasks much better than humans. However, those who take the leap of faith, educate themselves and incorporate machine learning, will be greatly rewarded. They will be at the forefront of active investing in the coming decades, taking more informed investment decisions and be armed with the knowledge of what they don’t know!

“It ain’t what you know that gets you into trouble

It’s what you know for sure that just ain’t so”

Mark Twain

This material is provided for information purposes only and does not constitute, and shall not be considered as, an offer, solicitation or invitation to engage in investment operations or as investment advice. All reasonable precautions have been taken to ensure the correctness and accuracy of the information. However, the correctness and accuracy are not guaranteed and we accept no liability for any errors or omissions. The material may not be reproduced or distributed, in whole or in part, without our prior written consent.

It is emphasized that investment returns shown are simulated and do not represent actual performance of assets during a period. If the simulated strategy had been implemented during the period, the actual returns may have differed significantly from the simulated returns presented. Past performance, whether actual or simulated, is not a reliable indicator of future results and the return on investments may vary as a result of currency fluctuations.

AI Alpha Lab ApS

CVR 40 41 55 99