Risk and uncertainty

Research Impact

'Ask five different economists a question, and you will get six different answers'. This joke is often used to put economists in a bad light, while if anything, the economist's response is closest to the academic ideal. We deal with a limited amount of data, which is often poorly recorded and driven by the beliefs of an impossible number of decision makers. Even worse, many of the official figures are subject to political pressure, and are often revised. In order to obtain a perfect setting in which we can get accurate and proper data from which it is easy to infer relations, we have to resort to the controlled setting of experiments, where we can only answer small questions that are impossible to generalize.

It is therefore no wonder that economists cannot agree on what to make of the economy. No two datasets are truly comparable and each will lead to different conclusions. The conclusions drawn are in the hand of the researcher or econometrician. The results of any study are a result of both the model specification and the data used to test the model on. Careful thought therefore has to be put in both of these aspects, and each specification and dataset has its pros and cons. These are very tough decisions to make.

And it gets worse from there, as the process of quantifying meaningful economic relations is a matter of econometrics and statistics. Since we are dealing with probabilities, there is always the chance that we find an economic relationship where there is none, and similarly, that we are unable to disentangle anything meaningful while in reality one variable causes the other. A large part of the econometrics literature concerns itself with accurately measuring the uncertainty in our estimates. Accurate assessment of the precision is vital, as it controls the degree of evidence a dataset provides for a certain claim.

This thesis provides new methods to (i) better quantify this uncertainty in order to be able to make more accurate claims, (ii) reduce the uncertainty in order to be able to make claims we were not able to make with less precise estimates, and (iii) use this precision in order to assess the reliability of data over time, such that we can better react to new information by disregarding poor estimates. Incorporating these techniques, and many other techniques developed by others, can hopefully lead to the fifth economist at the very least agreeing with himself, and only giving one answer to the question.

The applications in this thesis mainly revolve around risk management, which has clear applicability outside of academia. The methods presented may be of interest to banks, asset managers, pension funds, but also for regulatory institutions. All these institutions face the exact same problem of interpreting data to decide on their operations and policies. These decisions may have a large impact on their clients' investments, retirement savings or the economy as a whole in the case of policy makers.

(i) Quantifying uncertainty

In chapter 2 we propose a method to classify different financial institutions in terms of their 'systemic risk', or the degree of how 'too-big-to-fail' they are. For regulators this is a daily concern, and many ideas have been proposed to somehow curb the banks' behavior to prevent another global financial crisis. Their accurate assessment of these firms' risks is of great importance, first, to hopefully be effective in their policies, and second, as they will impose restrictions on these firms. Academics have also tried to shed light on this issue and have developed an enormous literature of all kinds of measures of this systemic risk, the ultimate goal being to rank these firms. We quantify the uncertainty in these measures (and potentially many others), in order to make more accurate claims. We find that the uncertainty in these measures is of considerable magnitude and one has to be very careful in making bold claims about their ability to distinguish firms in terms of their riskiness.

(ii) Reducing uncertainty

Chapter 3 proposes a new method to more accurately estimate the common risks of stocks. An accurate estimate is of great importance for at least two reasons. First, it can be used to improve portfolio allocations in terms of their risk-return balance, and second, it provides more accurate estimates of the risks inherent for a given holding. A precise estimator, such as the one proposed in this thesis, can therefore be of benefit to any company who deals with stocks.

With a similar goal, but in a very different way, chapter 6 proposes a method to improve the hedging of long-term liabilities. The problem is clearly of great interest to financial institutions such as pension funds. Here, the method to reduce uncertainty is to more accurately capture the time-variation in the dynamics of interest rates. The chapter provides a simple method that can be used by any practitioner to think about ways to improve their hedging decisions, and reduce surprises in future costs.

(iii) Exploiting uncertainty

Chapters 4 and 5 propose to explicitly use the amount of uncertainty that we know is present to make better decisions. By using this information different data points can get different weights. These papers use this in a dynamic setting, where intuitively, new

information is only incorporated when the observation is sufficiently precise. The research shows that by properly taking into account this uncertainty, more accurate decisions can be made. Chapter 4 demonstrates improved volatility forecasts and the illustrations in chapter 5 include portfolio formation and beta hedging. However, there are many more potential applications. The methods can be used in any time series where there we have a general idea on the precision of the estimates, such as time series of analyst forecasts.

All of these three elements allow the economist, academic or practitioner, to either better extract economic implications from data, or better forecast important variables. Apart from the immediate appeal to risk management presented herein, most of the presented methods offer general applicability in a wide variety of set-ups, which for now are left unexplored.