EXPERT ADJUSTMENTS of MODEL FORECASTS

Theory, Practice and Strategies for Improvement

PHILIP HANS FRANSES

Expert Adjustments of Model Forecasts

Theory, Practice and Strategies for Improvement

PHILIP HANS FRANSES

CAMBRIDGE
UNIVERSITY PRESS

Expert Adjustments of Model Forecasts

To what extent should anybody who has to make model forecasts generated from detailed data analysis adjust their forecasts based on their own intuition? In this book, Philip Hans Franses, one of Europe's leading econometricians, presents the notion that many publicly available forecasts have experienced an 'expert's touch', and questions whether this type of intervention is useful and if a lighter adjustment would be more beneficial. Covering an extensive research area, this accessible book brings together current theoretical insights and new empirical results to examine expert adjustment from an econometric perspective. The author's analysis is based on a range of real forecasts and the datasets upon which the forecasters relied. The various motivations behind experts' modifications are considered, and guidelines for creating more useful and reliable adjusted forecasts are suggested. This book will appeal to academics and practitioners with an interest in forecasting methodology.

PHILIP HANS FRANSES is Professor of Applied Econometrics and Professor of Marketing Research at the Erasmus University Rotterdam. Since 2006 he has served as the Dean of the Erasmus School of Economics. His research interests concern the development and application of econometric methods for problems in marketing, finance and macroeconomics.

'We are all aliens to ourselves,
and if we have any sense of who we are,
it is only because we live inside the eyes of others.'
(Auster, 2012)

Figures

Tables

Preface

This monograph aims to collect together recent theoretical insights and various empirical results for a rapidly developing area concerning the analysis of business and economic forecasts. Strictly following econometric textbooks, it is tempting to assume that forecasts are the result of a careful modelling exercise, and that the econometric model-builder delivers the forecast to an end-user. However, since the mid 1980s the literature has contained various studies which suggest that the model-builder is not always the last person to deal with the forecast, and there can be someone in between the model-builder and the end-user, whom I will call the expert. This expert can modify or adjust the model forecast, after having received it and evaluated it, and it is this adjusted forecast which is typically delivered to the end-user. The early literature contains some scattered examples where such an adjustment occurs, sometimes to the benefit of the quality of the forecast, but sometimes not. Recent literature shows a revived interest in these expert-adjusted forecasts, for various reasons. First, and as will be argued in this book, it turns out that many, if perhaps not all, economic forecasts seem to undergo some tweaking from an expert. The recent availability of very large databases with expert-adjusted forecasts and model forecasts emphasizes this outcome. Second, there is a revived interest in analysing the quality of economic forecasts, and when experts have adjusted econometric model forecasts it may be necessary to rethink how such final forecasts should be evaluated. Third, it will be interesting to understand what it is that an expert does, in particular in cases where they themselves do not say. Below we will see that there are many possible reasons for experts to modify forecasts, and in this book I aim to put the experts' behaviour into an econometric perspective. I will argue that adjusting model

forecasts can be quite a good idea, and there are some potentially useful guidelines. It is hoped that this book will rouse the interest of academics and practitioners to pursue further research and obtain practical experience in order to learn how to create useful and reliable forecasts.

My interest in this area basically started with the analysis of a large database that I was able to acquire from a large pharmaceutical company. So, in addition to an overview of the earlier studies in the forecasting literature and more recent accounts, this book also covers most of my findings and insights in a single volume. My research has been spurred by the availability of various very valuable databases, which I (of course with the help of students and colleagues) was allowed to analyse. Based on these analyses, theories were developed, and with these we were able to make recommendations to practitioners. Most of our results have been published in articles in international journals, which I will not replicate, although I will highlight the outcomes. The chronology of my articles on this topic shows how our insights were obtained, and in which order. However, with hindsight this chronology should be different, and in fact, if the reader were to read my research work in chronological order, it would be quite easy to lose the main insights. Hence, I decided to write this book to put all my results into a proper sequential perspective.

It all began with a lecture that I gave in 2006 in Brussels. In 2004 I had published an article in *Interfaces* entitled 'Do we think we make better forecasts than in the past? A survey of academics' (Franses, 2004). One of the conclusions of the survey was that half the interviewed academics believed that econometric or statistical model forecasts could be improved by (somehow) including the domain knowledge of experts. Until that moment, I had rarely considered this notion, because as a trained econometrician I had always believed that forecasting only from econometric models was always the best option. Sometime in 2005, I was approached by people from a company called Marcus Evans, who were looking for a speaker at a conference to be organized in Brussels, on 16 and 17 March 2006, with the theme 'Making supply meet demand'. In the midst of hectic times in our Erasmus School of

Economics, I thought it would be nice to do something different, and I accepted the invitation. My presentation, accordingly entitled 'Forecasting demand, can we do better these days?', was scheduled right before lunch on 16 March. The last slide mentioned that 'models may benefit from expert adjustment', but, I concluded, 'how to add such knowledge and how to evaluate adjusted forecasts is still an open question'. I was invited to join the lunch and at the table was a fellow Dutchman, who introduced himself as Sander Demouge from the Netherlands-based pharmaceutical company Organon. He said, 'We have the same problem at Organon where we have data on forecasts from experts and from an automated forecast support system (FSS), and we want to know how to interpret those expert forecasts. Can we perhaps talk to you one day to see if together we can find a way out?' My 'Yes, of course, with pleasure' marked the start of the research project, various highlights of which are now summarized in this book.

Sander and I met again in November 2006 when he presented the issues at stake at Organon. They had a version of ForecastPro™ – forecasting software, which they used to generate statistical model-based sales forecasts each month for all their products sold in forty-plus countries. At the same time, local managers quoted their own forecasts, and these were also recorded and stored. Organon's key problem was that they had announced bonus payments for these managers depending on how much more accurate their forecasts were compared to the statistical model, but the company did not know how to measure that quality effectively. I said that I was more than happy to help, and if I did, would Organon allow me and my team to publish the outcomes if these were of enough interest to international publications? A few days later Sander confirmed that publishing was no problem as long as we did not mention the product names and the countries. On 20 December 2006 I received the spreadsheet with the data, and this file turned out to be huge. In modern-day language one would call this database an example of 'big data'. It contained all the information that was needed to analyse the differences between model forecasts and expert forecasts, for hundreds of products, sold in various

months from 2004 to 2006 in a range of countries on all continents for horizons one month to twenty-four months. This was an amazing database, and I could not wait to start the analysis. One nice feature of it was that I also received data on some of the traits of the managers (experts) who were responsible for making the forecasts.

In September 2006 I started as the Dean of the Erasmus School of Economics, and this hampered me from making a quick start on analysing Organon's database. To be honest, my complete lack of skills with programs such as Excel, in which the data were presented, also did not help progress. Luckily I was able to hire Rianne Legerstee, one of the most gifted students ever to attend our Econometric Institute. She spent almost the first six months of 2007 sorting the data so that they could be analysed. This meticulous work turned out to be extremely useful, as we could put together various papers. We wrote various reports for Organon and in the autumn of 2007 we presented our results, which Organon also shared with their managers during a training session. They were very happy with our results, which among other things showed that the expert forecasts were rarely better than the model forecasts. So they now also had a reason to stop linking forecast performance with bonuses. The training session turned out to be very helpful, as we demonstrated in a 2011 report and a 2014 paper (summarized in Chapter 5 below). In January 2008, Organon was taken over by an American company, and a little later Sander Demouge informed me that he had moved to another job.

Rianne and I kept working on these data, and we tried to publish our results in international academic journals. This turned out not to be immediately very successful, and it was not until early 2009 that we had our first piece accepted. In 2008 Rianne accepted my offer to study for a PhD, and in 2012 she graduated with a thesis that included two papers that also dealt with the Organon data. It is fair to say that without Rianne Legerstee there would have been no papers, without Sander Demouge there would have been no data, and without my lecture in Brussels, there would have been no book.

In the meantime I got more and more involved in the topic of the interaction between experts and models, and this led to contacts with

KLM Royal Dutch Airlines, the Netherlands Bureau of Economic Policy Analysis (CPB) and Bayer in Leverkusen, Germany. I was allowed to analyse expert forecasts and model forecasts (if they were available). I thank Pieter Bootsma and Stefan ten Haaf of KLM, Henk Kranendonk and Debby Lanser of CPB, and Prasad Saraph, Christopher Baron and Achim Siegert of Bayer for their generous help with their data. Amazingly, Henk and Debby of the CPB were so kind as to re-create the CPB model forecasts using the older versions of the model, which turned out to be an enormous effort.

The Econometric Institute of the Erasmus School of Economics is by far the best place in the world to work as an econometrician. It hosts the best students, and I would specifically like to mention Bert de Bruijn and Marjolein van Baardwijk for their assistance. The institute also hosts the best colleagues, and here I would like to thank Richard Paap, Dennis Fok, Dick van Dijk, Patrick Groenen, Christiaan Heij, Michael McAleer, Michel van der Velden and Alex Koning for their helpful suggestions over the years, and for their co-authorship on various projects, including some referred to in this book.

I am very proud that this book will be incorporated in the catalogue of Cambridge University Press. It is a great honour to me that they want to publish this book, and I am very thankful to my editor Chris Harrison for his ongoing trust in my academic endeavours. Also, four anonymous reviewers gave detailed and very helpful suggestions, which seriously improved this book. Parts of the book were presented as lectures at the Econometric Institute of the Erasmus School of Economics, the Netherlands Bureau for Economic Policy Analysis (CPB), Organon, the University of Groningen, and at the International Symposia on Forecasting in 2011 (Prague) and 2013 (Seoul).

To end this preface, I would like to dedicate this book to two individuals who have been very important to me and my career. The first is my PhD supervisor Teun Kloek. When I finished my PhD thesis in 1991 I had learned that econometrics was the best academic subject there is, and that, basically, everything is interesting to study. The second is my incomparable senior colleague Clive Granger who

unfortunately died way too early. In 2007, when we were enjoying the sea winds around Neeltje Jans, he asked me 'Why would someone like you be so stupid as to become a Dean?' I replied that I hoped to have enough time to do some research, and with this book I hope that I would have convinced him that I did.

Anyway, the main lesson that I can give to anyone who reads this book is that if you are asked to give a lecture, do it, and also pose a few questions to the audience: you never know what will happen.

Contents

I   Introduction

This monograph deals with the situation where an analyst evaluates expert forecasts and model forecasts, and where it is known that the expert has seen the model forecast and thus that the expert forecast potentially amounts to an adjustment of the model forecast. More precisely, the analyst assumes that

$$\text{Expert-Adjusted Forecast} = \alpha \text{ times Model Forecast} + \text{Adjustment.} \qquad (1.1)$$

This additive expression is chosen for analytical convenience, as will become clear in Chapter 2, and also to easily allow for the possibility that the model forecast and the expert forecast have opposite signs.

It is important to stress that the analyst only observes the model forecast and the final expert forecast, and of course also the realized observation, but that the analyst does not observe the value of α, nor the size of the adjustment. In many practical settings, the analyst is usually not the same individual as the expert who adjusts the model forecast, nor is the analyst the same person as the model-builder. In fact, the analyst may have to report to management or to policymakers on the usefulness and relevance of the final expert forecasts, perhaps relative to the model forecast. Such expert forecasts can concern business and economic variables: for example, sales of durable products, earnings of companies or macroeconomic variables like gross domestic product (GDP) or inflation. The forecasts may have to be generated very frequently, for example, hourly, or they may also be quoted just once every half-year.

There is one particular feature that is very important here and that is that the experts are assumed to quote their forecasts *given* that they have received model-based forecasts. It is, however, uncertain if

and how they actually incorporate the model forecasts into their decision process, as usually there is no written documentation. So it may be that an expert sets α equal to 0, and fully bases the final expert forecast on his or her own judgement, but this is unknown to the analyst. Indeed, typically, experts do not document how they decide on the values of α and the adjustment.

It is irrelevant whether the model forecasts originate from multi-equation macroeconomic models or from simple extrapolation tools, or anything in between. It may very well be that an expert does not know what the model or forecast algorithm looks like, and in most practical cases it also holds that the expert cannot exercise any influence on how the model forecasts are created. It is usually the case that the expert is not the same person who designs the model, but no specific assumptions on this feature have to be made. The models and their parameters can be updated every single hour, or they may be taken as constant for a long period of time. The statistical tools with which the models are calibrated are largely irrelevant, and it may well be that the expert in fact does not have a clear-cut idea of how the model forecasts were created. In the end, the situation is that the analyst observes an expert forecast, a model forecast and a realization, and the analyst has to evaluate the expert forecast using some criterion.

A key premise of the analysis in this book is that the analyst does not know α or the size of the adjustment and that the analyst is also unaware of *how* the expert has chosen a value of α and the adjustment. The size of these two features, that is, α, and the adjustment, can be set by the expert using his or her own intuition or model, but how that is done is usually unknown. The definition of 'intuition' is obtained from the *Oxford Dictionary* and is 'the ability to understand something instinctively, without the need for conscious reasoning'. At the same time, it may also be that the expert uses knowledge that can be documented and evaluated. For example, a known future regime shift may not be incorporated in the model and thus not in the model forecast, and the expert may use this knowledge to assign a value to the adjustment.