Essays in the Study and Modelling of Exchange Rate Volatility

Genaro Sucarrat

5 September 2006

Contents

Acknowledgements

In the course of researching and writing this thesis I have acquired a debt to many people. The person I owe the most is my supervisor Luc Bauwens, who is also co-author on two of the essays that make up this thesis. My collaboration with Luc goes back to my Master dissertation and I have thoroughly enjoyed learning from and working with him throughout this period. Repeatedly you will find statements like "I owe a very very special thank to Luc, one of the best persons I have ever met in my life" (Veredas 2002, p. 2), "I would have never started a PhD if I had not enjoyed working with him so much" (Beltran 2005, p. i) and I would like to thank Luc for his "dedicated help throughout my thesis" (Giot 1999, p. i), in the acknowledgment parts of the theses of his former PhD students. I can testify that these are not pure politeness phrases, they are genuine expressions of gratitude. They convey the appreciation of Luc as a scholar and as a person that I and his other PhD students share. Luc, thank you very much.

Another person that I am substantially indebted to is my second co-author, Dagfinn Rime. My collaboration with Dagfinn goes back to my stay in Oslo during the Autumn 2003 - Winter 2004, and he is the originator of two of the studies in the sense that they grew out of ideas that originally were his. Baxter and Rennie (1996, p. vii) have stated that in academia "effort is too often expended on finding precise answers to the wrong question..with no reference to the concern of practitioners". Dagfinn is one of the exceptions. He has a sense for what the interesting questions are, and I am grateful for the opportunity to work with and learn from him.

A colleague once joked that the only persons that will ever read my thesis is my thesis committee. Admittedly this is an outcome which is not unlikely, so I am particularly grateful to them for devoting their time, effort and patience in reading and evaluating it. In addition to Luc the committee consists of Farooq Akram, Vincent Bodart, Sébastien Laurent and Fatemeh Shadman.

I and my coauthors have also received many useful comments, suggestions and questions from numerous other people, including Yacine Aït-Sahalia, Helena Beltran, Charles Bos, Ana C. Cebrian, Geert Dhaene, Neil Ericsson, Fausto Galli, Javier Gil-Bazo, Jesús Gonzalo, Eric Ghysels, David F. Hendry, Alfredo Ibañes, Eilev Jansen, Jan-Tore Klovland, Alois Kneïp, Hans-Martin Krolzig, Michel Mouchart, Ragnar Nymoen, Walid Ben Omrane, Andrew Patton, Eric Renault, Jeroen Rombouts, Rickard Sandberg, Enrique Sentana, David Veredas, Fredrik Wulfsberg, three anonymous referees, participants at the bi-annual doctoral workshop in economics at Université catolique de Louvain in May 2004 (Louvain la Neuve), participants at the MICFINMA summer school in June 2004 (Konstanz), participants at the bi-annual doctoral workshop in economics at Université catolique de Louvain in February 2005 (Louvain la Neuve), participants at the poster session following the joint CORE-ECARES-KUL seminar in April 2005 (Brussels), participants at the bi-annual doctoral workshop in economics at Université catolique de Louvain in May 2005 (Louvain la Neuve), participants at the Journal of Applied Econometrics conference in June 2005 (Venice), participants at the 3rd OxMetrics Conference in August 2005 (London), seminar participants at Statistics Norway in September 2005 (Oslo), seminar participants at the Department of Economics at the University of Oslo in September 2005 (Oslo), seminar participants at the Norwegian School of Economics and Business Administration in November 2005 (Bergen), conference participants at the 16th $(EC)^2$ in December 2005 (Istanbul), seminar participants at CEMFI in February 2006 (Madrid), participants at the International Conference on High Frequency Finance in May 2006 (Konstanz), seminar participants at Departamento de Economia de Empresa at Universidad Carlos III de Madrid in May 2006, participants at the International Symposium on Forecasting in June 2006 (Santander), and participants at the European Meeting of the Econometric Society in August 2006 (Vienna).

The sponsors of my PhD education have been several. Initially I was funded by The Norwegian State Educational Loan Fund, and Lise and Arnfinn Heje's Fund (Norway), then by The Finance Market Fund (”Finansmarkedsfondet”, Norway) before finally going to Universidad Carlos III de Madrid under contract HPRN-CT-2002-00232, MICFINMA, of the European Community's Human Potential Programme. Financial support from CORE, Universidad Carlos III de Madrid and Norges Bank's fund for economic research (2006) for conference attendance, and from CORE for data purchase, is also gratefully acknowledged.

The writing and researching of this thesis has taken place at several institutions. Most of the time was spent at CORE, so I am particularly grateful to the Econometrics Group, its friends and the CORE administration for all the nice times I had there. At Universidad Carlos III de Madrid, where the thesis by and large was completed, I am grateful to Álvaro Escribano, Jesús Gonzalo and Mikel Tapia for making the transition to the Spanish way of life ("mañana-mañana") smooth. I also spent substantial time at the Economics Department at the University of Oslo and at the Research Department at Norges Bank. My sincerest gratitude goes to them for their hospitality. At the latter institution I am also indebted to Erik Meyer and Janett Skjelvik at the Statistics Department for preparing the volume data studied in chapter 5, and for taking the time to answer questions regarding them.

Finally I would like to thank my family and friends who, literally, are scattered all over the world. At times this makes it difficult to catch up, but I hope to improve in the future.

Chapter 1

Introduction

When John Maynard Keynes and Harry Dexter White—the chief negotiators of the UK and the USA—negotiated over the financial world order in Bretton Woods in 1944, they both shared a common vision. The "plan accords", in the words of John Maynard Keynes, "every member government the explicit right to control all capital movements. What used to be heresy is now endorsed as orthodox" (Keynes 1980b, p. 17, cited in Helleiner 1994, p. 25). Although the UK and the US were only two out of the more than forty allied nations that attended the conference, the participants nevertheless shared the view that the unstable world financial order had played an important part in bringing about two world wars.[1] As the then US Treasury Department Secretary Henry Morgenthau told the conference, the goal of the Bretton Woods Agreement is to "drive the usurious moneylenders from the temple of international finance" (Gardner 1980, p. 76, cited in Helleiner 1994, p. 4). Everyone's mind revolved around the question of how a stable international financial order could be put in place so that future wars could be avoided, and to this end a dominant view was that international financial flows had to be restricted in order to promote trade among states, and so that Keynesian demand management and stabilisation policies could be pursued successfully. As White stated in early drafts of the Bretton Woods Agreement, capital controls "would give each government much greater measure of control in carrying out its monetary and tax policies" (Horsefield 1969, p. 67, cited in Helleiner 1994, p. 33). Similarly, Keynes argued that the "management of the domestic economy depends upon being free to have the appropriate rate of interest without reference to the rates prevailing elsewhere in the world. Capital control is a corollary to this" (Keynes 1980a, p. 149, cited in Helleiner 1994, p. 34).

Although the strict capital controls envisaged by the Bretton Woods Agreement were never implemented, the period that followed has nevertheless come to be known as the Bretton Woods System of exchange rates because its nature was very much in spirit with the Agreement. The years from 1945 to 1970s constitute the longest period in modern times with comprehensive capital control enforced by the major Western economies. During the 1960s the system came increasingly under stress and ultimately collapsed in the beginning of the 70s. The US and the UK subsequently led the way in a wave of financial market deregulation, and states had to adopt new ways in handling exchange rate variability, including coordinated exchange rate management between governments, joint currency market interventions by central banks, monetary unions, currency boards, etc. Faced with such a situation financial econometric models can be of use for both policy makers and businesses alike. For businesses they may be useful as tools in risk management, whereas policymakers may use them to acquire knowledge about what and how economic factors impact upon financial variability for informed policymaking.

Throughout this thesis a distinction is made between the *variability* of exchange rate changes on the one hand and the *volatility* of exchange rates on the other. According to the Merriam-Webster Online dictionary the etymological origin of the word "volatility" is the latin *volatilis*, which is a derivative of *volare*, "to fly". Although one of the meanings of volatility still is "flying" or "having the power to fly", the term typically carries a rather different and specific meaning in financial econometrics, namely the conditional or unconditional variance (or standard deviation) of a financial price increment. Exchange rate variability on the other hand is here defined as squared exchange rate return. Contrary to what the title of this thesis might suggest the main focus of this thesis is actually on variability. But since exchange rate volatility effectively is a model or prediction of exchange rate variability in most of the models employed in this thesis (chapter 2 provides further discussion and makes the distinction more precise), and since it is customary to speak of "volatility models" in this way in the literature although it strictly speaking is incorrect—the correct would be "models of the squared error term" or something in this vein, the terms will at times be used interchangeably as if they stood for the same phenomenon. Nevertheless, as will become apparent—in particular in chapter 4 where a modelling strategy is evaluated, whether one's main focus is on one or the other can lead to important methodological choices.

There are three main themes in this thesis. The first theme is the relation between exchange rate variability and market activity, and more precisely there are two questions within this theme that will receive special attention. First, what is the relation between period to period—or "short-term"—changes in variability and period to period changes in market activity? And second, what is the relation between the general —or "long term"—level of variability and the general level of market activity? When Karpoff (1987) surveyed the relationship between financial variability and trading volume—a measure of market activity—during the mid-eighties, only one out of the nineteen studies he cited was on exchange rates. This was partly because of a lack of data. The increased availability of data brought by the nineties has changed this and today there are numerous studies that shed light, directly or indirectly, on the relation between exchange rate variability and market actitivity. The empirical studies in this thesis, I believe, contribute to this literature.

The second theme concerns the modelling of exchange rate volatility. It is fair to say that most of the financial econometric volatility literature has been driven by a business perspective rather than a policymaking perspective. This is not to say that questions that are interesting for businesses are not interesting for policymakers. Nor am I saying that the business perspective does not help us to understand the economic reasons for the variation in variability. On the contrary. What I am saying though is that questions more often would have been addressed differently if the literature had been more influenced by a policymaking or explanatory perspective rather than a business perspective. For instance, the seminal work by Clark (1973), which models price increments as random but volatility as dependent on "information arrival", had as its main objective to explain the leptokursis of financial returns rather than variation in volatility. A more recent example is the continuous time volatility literature, spurred by Andersen and Bollerslev's (1998) highly influential article. It is easy to see the usefulness of this literature for those who price derivatives by means of continuous time models, but it is not apparent that the insights of this literature are equally useful for policymakers. Agents' loss functions are typically in terms of observed prices, so I thus find it natural that the variability of those observed price increments occupy the centre of attention, or at least constitute the ultimate yardstick to evaluate estimates from continuous time models. Moreover, in addition to a range of theoretical and practical issues, there are serious philosophical objections to the view that low frequency models of volatility and/or variability should be evaluated against estimates derived from continuous time models in explanatory modelling in general. (These philosophical objections will be briefly alluded to in subsection 4.4.1 in chapter 4.)

The third theme of the thesis concerns the nature of human reality. I take as my starting point, my human ontology so to say, that the course of history is indeterministic, that history does not repeat itself, that the past has a bearing upon the future, that people differ from each other at each point in time, that economic events have temporal extension, and that causal connections between economic events supervene on processes made up of chains of economic events, each with temporal extension.[2] Empirical dynamic econometrics is thus an activity that aims at developing history-repeats-itself representations (models), that is, approximate generalisations, that hold over time. An own chapter in this thesis is devoted to how these ideas can be reconciled within the framework of David F. Hendry's (1995, chapter 9) reduction theory, and that chapter appears towards the end of the thesis as chapter 6. Although seemingly unrelated to the rest its content actually underpins the whole thesis, and proposes some solutions to certain issues in the financial volatility literature. The direct relation between this chapter and other parts of the thesis is most obvious in chapter 4 which evaluates the socalled general-to-specific (GETS) methodology applied to the modelling of exchange rate volatility, since the GETS methodology is based upon and tries to mimic reduction theory. Moreover, section 4.4.1 of that chapter raises serious philosophical objections—motivated by my human ontology—to the view alluded to above that low frequency predictions of volatility should be evaluated against estimates derived from continuous time models. Also, many other choices, points and arguments of greater or smaller importance in the thesis are based upon the argument and assumptions of chapter 6. Logically it should therefore appear at the beginning of the thesis, but for pedagogical reasons I have placed it at the end.

A further consequence of my viewpoint on the nature of human reality concerns the particularity of theories and models. The empirical studies in this thesis are all on Norwegian data, which means that the object of study differs substantially from the object of study that most frequently appear in international journals. For this reason it seems pertinent to learn more about Norway.

1.1 The Norwegian economy

This section gives a broad overview of the Norwegian economy over the period studied in this thesis, namely 8 January 1993 to 25 February 2005, paying particular attention to characteristics of relevance for the study and modelling of exchange rate variability.[3] More details are provided in bits and pieces later in the thesis, in particular in chapters 2 and 3.

Norway is a small and open economy with only four and a half million inhabitants, and has one of the highest ratios of export plus import to GDP in the world. Accordingly, with its own money and no formal peg or exchange rate arrangement against other currencies, the variability of Norwegian exchange rates is of major importance. Over the period 8 January 1993 - 25 February 2005 one may distinguish between three different exchange rate management regimes. The first regime may be labelled an "exchange rate stabilisation" regime and lasted to around the middle or end of 1998 (a discussion on when this regime ended appears below). Until the end of 1992 the Norwegian krone was pegged to the European Currency Unit (ECU) by Government resolution, which obliged Norges Bank—the Central Bank of Norway—to buy and sell Norwegian kroner at values within +/- 2.25% of the reference value 7.9940 NOK = 1 ECU. However, Norges Bank also used interest rate changes actively as a means of keeping the value within the bands. The obligation to buy and sell Norwegian kroner within these margins could only be suspended through Government resolution for a limited period of time (maximum 30 days), and such a resolution was issued 10 December 1992 following a request from Norges Bank due to the currency market turmoil in the Autumn 1992. On Friday 8 January 1993, 29 days after the 10 December resolution, a new Government resolution was issued which freed Norges Bank from the obligation to buy and sell the krone at any value. Effectively this implied a change from a fixed to a floating exchange rate management regime, whose new objective was to stabilise—in the sense of avoiding large swings—the Norwegian krone.[4] In order to stabilise the Norwegian krone Norges Bank made use of both interest rate changes and currency interventions, but it is not publicly known to what extent the latter was used since Norges Bank currency interventions are confidential.

Whether the second exchange rate regime, which may be labelled "partial inflation targeting", started in May 1998 or in the beginning of 1999 with the arrival of the current Governor of the Bank is not evident. In May 1998, after a formal letter exchange with Norges Bank, the Ministry of Finance affirmed that the best way to achieve exchange rate stability was to pursue an inflation policy that did not differ substantially from the European Monetary Union (EMU) countries. This differed slightly from previously, since previously it was comparable inflation level to Norway's main trading partners—which comprises more countries than EMU—that had been specified as one of the main means of achieving exchange rate stability. So according to one view a shift in exchange rate management took place in May 1998. But even if the Ministry of Finance affirmation meant a *de facto* change in May 1998 rather than in the beginning of 1999, it would nevertheless be difficult to detect the change in the data because of the currency market turmoils that followed during the summer and autumn of 1998 due to the Russian Moratorium. Moreover, according to several analysts the arrival of the new Governor in January 1999 implied greater weight than his predecessors on comparable inflation policy (with the EMU countries) as a means of achieving exchange rate stability.[5] For these reasons we find it convenient to define the start of the partial inflation regime as to coincide with the arrival of the new Governor in the beginning of 1999 rather than from the formal letter exchange in May 1998.

The third exchange rate regime started 29 March 2001 when Norges Bank was instructed by the Ministry of Finance to fully pursue an inflation target of 2.5%. The period after March 2001 may thus be termed a "full" inflation targeting regime. Although analysts agree that a formal change took place on this date, they disagree to what extent there were learning effects present before and/or after the change.

The most important exchange rate for the actors that regularly trade the Norwegian krone (NOK) in the spot market is the krone against the

Euro (NOK/EUR), and before 1999 the most important exchange rate was the krone against the Deutsche Mark (NOK/DEM).[6] For simplicity reasons we transform the latter into Euro-equivalents using the official conversion rate 1.95583 DEM = 1 EUR, so that we can talk of a single exchange rate rather than two. Figure 1.1 contains a graph of the level of the Bid NOK/EUR at the end of the last trading day of the week over the study period.[7] An increase in the exchange rate means a depreciation in the value of the NOK, and a decrease the opposite. Level expectations appear to be present in the sense that there are no values above or below 9.2 and 7.2, respectively, and in the sense that rarely do we find values above 8.8 or below 7.6. Figure 1.2 contains a graph of the log-return of NOK/EUR in percent over the study period, and exhibits two apparent characteristics.[8] First, there seems to be a notable break in the general level of variability around the end of 1996 or beginning of 1997, and second other structural breaks do not seem to occur henceforth (for instance in relation with the transition from one regime to another)—or at least they are not apparent by just looking at the graphs. It would be of great interest if one could explain the general increase from 1997 and onwards by a general increase or decrease in NOK/EUR trading, but as we will see things are not so straightforward.

1.2 Overview of thesis

The rest of the thesis is divided into six chapters and one appendix. Chapter 2 describes and motivates the definitions and models of variability that will be entertained in this thesis. In particular, the chapter describes and motivates the exponential model of variability (EMOV) as particularly useful for explanatory financial variability modelling and hypothesis testing. Chapter 3 makes full use of the EMOV and the definitions of variability from chapter 2 in a study of the relation between exchange rate variability and market activity. Chapter 4 undertakes an out-of-sample forecast evaluation of general-to-specific (GETS) modelling of exchange rate volatility, a modelling strategy that has proved powerful in the explanatory econometric modelling of many other economic series. Chapter 5 approaches the relation between exchange rate variability and market activity from a slightly different angle compared to that of chapter 3. A Norwegian dataset on currency transaction volume enables us to study the role played by heterogeneity in the relation between variability and market activity. Chapter 6 proposes a solution to some shortcomings in the probabilistic reduction theory that underpins GETS modelling. Chapter 7, the final chapter of the thesis, contains the conclusions and proposes an agenda for future research. References and endnotes follow the conclusions, whereas tables and figures follow each chapter. The data appendix which contains details of the data transformations and sources of the original untransformed data is placed at the very end of the thesis.

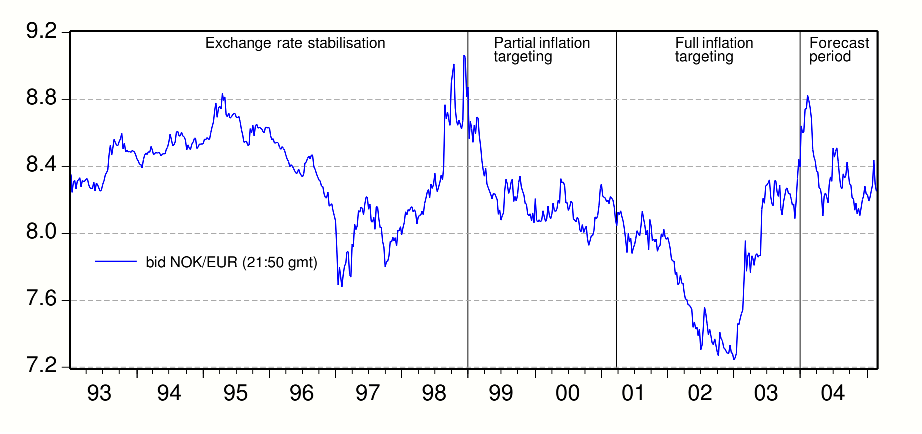

Figure 1.1: Weekly bid NOK/EUR from 8 January 1993 to 25 February 2005 (NOK/DEM in Euro equivalents before 1999)

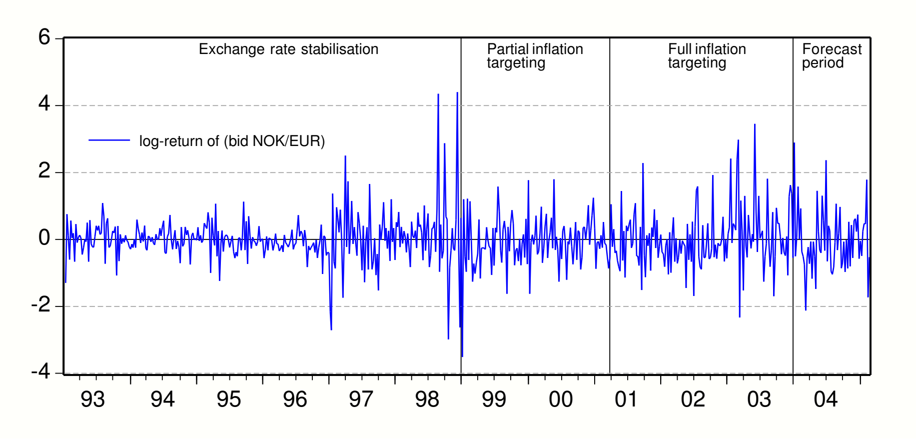

Figure 1.2: Weekly log-return of bid NOK/EUR in percent from 8 January 1993 to 25 February 2005

Chapter 2

Definitions and models of exchange rate variability

This chapter is a substantially revised amalgam of parts in Bauwens, Rime and Sucarrat (2006), and Bauwens and Sucarrat (2006).

The main purposes of this chapter is to introduce and motivate the definitions and models of exchange rate variability that will be employed in the thesis, and to relate them to the literature. The chapter contains two sections. In the first section a distinction between period and within-period variability is made, a distinction which is of particular usefulness when studying variability across different exchange rate regimes as in the case of Norway in this thesis. The second section presents and describes the exponential model of variability (EMOV), a model that is especially useful and flexible in explanatory exchange rate variability modelling, and relates the model to the more common autoregressive conditional heteroscedasticity (ARCH) and stochastic volatility (SV) families of models.

2.1 Definitions of exchange rate variability

Conceptually we may distinguish between period variability on the one hand and within or intra-period variability on the other. If $S_t = \{S_{0(t)}, S_{1(t)}, \ldots, S_{n(t)}, \ldots, S_{N-1(t)}, S_{N(t)}\}$ denotes a sequence of exchange rates between two currencies at times $0, 1, \ldots, N$ in period t, then the squared (period) log-return $[\log(S_{N(t)}/S_{0(t)})]^2$ is an example of a period definition of variability. Range variability, defined as $[\log(\max S_t) - \log(\min S_t)]^2$ where $\log(\max S_t) - \log(\min S_t)$ is the range log-return, and realised volatility, defined as $\sum_{n(t)=1}^{N(t)}[\log(S_{n(t)}/S_{n-1(t)})]^2$, are examples of within-period definitions of variability. The main difference between period and within-period definitions of variability is straightforward. In addition to time 0 to time N variation the latter is also capable of capturing variation between 0 and N. For example, if S_n fluctuates considerably between 0 and N but ends up close to S_0 at N, then the two types may produce substantially different results. Under certain assumptions the three definitions essentially provide estimates of the same thing, see amongst other Parkinson (1980), Garman and Klass (1980), Andersen and Bollerslev (1998), Andersen et al. (2001), Andersen et al. (2005) and Aït-Sahalia (2006). However, the reader should be aware that nowhere do we rely upon restrictive continuous time models (see subsection 4.4.1 for an argument for why and under which circumstances continuous time models serve as restrictions on discrete models).

To fully appreciate the distinction between period and within-period variability, recall that Norway experienced three different exchange rate regimes (see subsection 1.1) over the sample studied in this thesis. It is not at all clear at the outset that period and within-period definitions of exchange rate variability behave and react similarly across the regimes, even though they should population-wise according to some statistical models. Two definitions of variability will play center stage in this thesis, period and range variability. The main characteristics of Norwegian weekly period variability and weekly range variability and their log-transformations are contained in tables 2.1 and 2.2, and in figures 2.1 and 2.2. Weekly period variability is denoted V_t^w and weekly range variability is denoted V_t^{hl} (the superscript "hl" is supposed to evoke the association "high-low"). Their log-counterparts are in small letters, that is, v_t^w and v_t^{hl}, and for their exact constructions the reader is referred to the data appendix at the end of the thesis. There are at least five characteristics worth noting. First, over the whole period the sample standard deviation of V_t^{hl} is more that the double than that of V_t^w. Second, the log-transformation, which makes pairs of large observations (in absolute value) less influential, matters for both correlations and standard deviations. For instance, over the whole sample the sample correlation between V_t^w and V_t^{hl} is 0.85, whereas the sample correlation between v_t^w and v_t^{hl} is only 0.65. With respect to standard deviations, the standard deviations of range variability are lower than those of period variability when the log is applied, whereas the opposite is the case when the log is not applied. Note also that the downward spikes in the figure of v_t^w is due to the log-transformation being applied on squared returns when returns are close to zero. Third, the log-transformed definitions are less correlated than one might have expected, with a minimum of 0.60 attained over the period January 1999 to March 2001. Fourth, general increases or shifts upward in variability around 1 January 1999 (the beginning of the partial inflation targeting regime) and 29 March 2001 (the beginning of the full inflation targeting regime) are absent—or at least seemingly so. In probabilistic terms, if variability is covariance stationary with mean μ then a change to μ' does not seem to take place around the regime changes. One might have expected that the policy regime changes (see section 1.1) would have resulted in general shifts upwards in variability around these dates. However, if this is the case then this is not evident by just looking at the graphs. Alternatively, the apparent absence of shifts in variability might be due to the fact that the markets had expected these changes and already adapted to them. Fifth and finally, graphically there seems to be a marked and lasting increase in variability around late 1996 or in the beginning of 1997. This is partly in line with Giot (2003) whose study supports the view that the Asian crisis in the second half of 1997 brought about a sustained increase in the variability of financial markets in general. In the case of Norwegian exchange rate variability, however, the shift upwards seems to have taken place earlier, namely towards the end of 1996 or in the beginning of 1997. This may be attributed to the appreciatory pressure on the Norwegian krone in late 1996 and early 1997. But another hypothesis, put forward by van Dijk et al. (2005), is that the general shift upwards is due to a European Council meeting in December 1996 in which a decision concerning the European Monetary Union (EMU) was made. Van Dijk et al. (2005) show that several daily European non-EMU exchange rates (denominated in US dollars) exhibit a shift upwards in the general level of variability around this date, including the NOK/USD exchange rate. This is partly consistent with the argument of Bjønnes et al. (2005). Their results are compatible with an increase in NOK-trading by foreigners, but they attribute it to a "speculative attack" rather than a general increase in NOK-speculation, which would be more in line with the explanation given by van Dijk et al. (2005).

2.2 Models of exchange rate variability

If s_t denotes the log of an exchange rate and r_t is either the period or range log-return, then we will refer to r_t^2 as variability. The exponential model of variability (EMOV) is given by

$$r_t^2 = \exp(\mathbf{b}'\mathbf{x}_t + u_t), \tag{2.1}$$

where $\mathbf{b}$ is a parameter vector, $\mathbf{x}_t$ is a vector of conditioning variables that may (or may not) contain values prior to t, and $\{u_t\}$ is a sequence of mutually uncorrelated and homoscedastic variables each with conditional mean equal to zero.[9] The linear specification in the exponent is motivated by several reasons. The most straightforward is that it results in simpler estimation compared with the more common ARCH and SV models, in particular when many explanatory variables are involved. Under the assumption that $\{r_t^2 = 0\}$ is an event with probability zero, then consistent and asymptotically normal estimates of $\mathbf{b}$ can be obtained almost surely with OLS under standard assumptions, since

$$\log r_t^2 = \mathbf{b}'\mathbf{x}_t + u_t \qquad \text{with probability 1.} \tag{2.2}$$

Another motivation for the exponential specification is that large values of r_t^2 become less influential. A third motivation, pointed to by (amongst others) Engle (1982), Geweke (1986) and Pantula (1986), and which subsequently led Nelson (1991) to formulate the exponential general ARCH (EGARCH) model, is that it ensures positivity. This is particularly useful in empirical analysis because it ensures that fitted values of variability are not negative. Finally, another attractive feature of the exponential specification is that it produces residuals closer to the normal in (2.2) and thus presumably leads to faster convergence of the OLS estimator. In other words, the log-transformation is likely to result in sounder inference regarding $\mathbf{b}$ in (2.2) when an asymptotic approximation is used. Applying the conditional expectation operator in (2.1) gives

$$E(r_t^2|\mathcal{I}_t) = \exp(\mathbf{b}'\mathbf{x}_t) \cdot E[\exp(u_t)|\mathcal{I}_t], \tag{2.3}$$

where $\mathcal{I}_t$ denotes the information set in question (note that $\mathcal{I}_t$ not only contains $\mathbf{b}'\mathbf{x}_t$, it may also contain additional information). An estimate of conditional observed volatility is readily obtained if either $\{u_t\}$ is IID or if $\{\exp(u_t)\}$ is a mean innovation, that is, $E[\exp(u_t)|\mathcal{I}_t] = E[\exp(u_t)]$ for $t = 1, \ldots, T$, since the formula $\frac{1}{T}\sum_{t=1}^{T}\exp(\hat{u}_t)$ then provides a consistent estimate of the proportionality factor $E[\exp(u_t)|\mathcal{I}_t]$.

To see the relation between the EMOV and the ARCH and SV families of models, recall that the latter two decompose returns into a conditional mean μ_t and a remainder e_t

$$r_t = \mu_t + e_t, \tag{2.4}$$

where e_t is commonly decomposed into $e_t = \sigma_t z_t$ if $\{e_t\}$ is heteroscedastic.[10] In principle r_t can be both period and range return, but admittedly most scholars would prefer to think of it as period return. Also, the literature has evolved having period return in mind. The better μ_t is specified the smaller e_t is in absolute value, and the better σ_t is specified the smaller z_t is in absolute value. If σ_t^2 follows a non-stochastic autoregressive process and if $Var(r_t^2|\mathcal{I}_t) = \sigma_t^2$, then (2.4) belongs to the ARCH family. A common example is the GARCH(1,1) of Bollerslev (1986)

$$\sigma_t^2 = \omega + \alpha e_{t-1}^2 + \beta\sigma_{t-1}^2, \tag{2.5}$$

with $z_t \sim \text{IN}(0, 1)$. Explanatory terms, say, $\mathbf{c}'\mathbf{y}_t$, would typically enter additively in (2.5). If σ_t^2 on the other hand follows a stochastic autoregressive process, then (2.4) belongs to the SV family of models, and in the special case where σ_t and z_t are independent the conditional variance equals $E(\sigma_t^2|\mathcal{I}_t)$.

The EMOV can bee seen both as an approximation to the ARCH and SV families of models of volatility, and as a direct model of variability. To see this consider the specification

$$r_t = \sigma_t z_t. \tag{2.6}$$

Squaring yields (2.1) above if $\sigma_t^2 = \exp(\mathbf{b}'\mathbf{x}_t)$ and $z_t^2 = \exp(u_t)$, and applying the log gives (2.2) with $u_t = \log z_t^2$. Now, recall that expected variability within the ARCH family[11] is

$$E(r_t^2|\mathcal{I}_t) = \mu_t^2 + \sigma_t^2. \tag{2.7}$$

In words, the total expected exchange rate variation consists of two components, the squared conditional mean μ_t^2 and the conditional variance σ_t^2. As Jorion (1995, footnote 4 p. 510) has noted σ_t^2 typically dwarfs μ_t^2 with a factor of several hundreds to one,[12] so the "de-meaned" approximation

$$\mu_t^2 + \sigma_t^2 \approx \sigma_t^2 \tag{2.8}$$

is often reasonably good in practice. As a consequence, the expression $\exp(\mathbf{b}'\mathbf{x}_t) \cdot E[\exp(u_t)|\mathcal{I}_t]$ can be interpreted both as a model of variability r_t^2 *and* as a model of volatility.

Table 2.1: Descriptive statistics of weekly period and range variability

		1993/1–2005/2 ($T = 633$)	1993/1–1998/12 ($T = 311$)	1999/1–2001/3 ($T = 118$)	2001/4–2005/2 ($T = 204$)
V_t^w	*Avg.*	0.60	0.49	0.54	0.81
	Med.	0.14	0.07	0.15	0.25
	St.dev.	1.66	1.84	1.27	1.56
	Max.	19.36	19.36	12.33	11.91
	Min	0.00	0.00	0.00	0.00
V_t^{hl}	*Avg.*	1.76	1.46	1.69	2.26
	Med.	0.83	0.38	1.15	1.52
	St.dev.	3.80	4.83	2.35	2.38
	Max.	69.06	69.06	22.76	15.58
	Min	0.02	0.02	0.23	0.20
v_t^w	*Avg.*	-2.40	-2.96	-2.02	-1.76
	Med.	-2.00	-2.60	-1.92	-1.38
	St.dev.	2.44	2.54	2.05	2.31
	Max.	2.96	2.96	2.51	2.48
	Min	-10.76	-10.76	-8.97	-10.23
v_t^{hl}	*Avg.*	-0.26	-0.87	0.16	0.43
	Med.	-0.19	-0.97	0.14	0.41
	St.dev.	1.31	1.41	0.78	0.86
	Max.	4.23	4.23	3.13	2.75
	Min	-3.92	-3.92	-1.45	-1.60

Note: Zero-values are due to rounding.

Table 2.2: Sample correlations between weekly period and range variability

Sample		V_t^w	V_t^{hl}		v_t^w	v_t^{hl}
1993/1-	V_t^w	1.00		v_t^w	1.00	
2005/2	V_t^{hl}	0.85	1.00	v_t^{hl}	0.65	1.00
1993/1-	V_t^w	1.00		v_t^w	1.00	
1998/12	V_t^{hl}	0.85	1.00	v_t^{hl}	0.65	1.00
1999/1-	V_t^w	1.00		v_t^w	1.00	
2001/3	V_t^{hl}	0.92	1.00	v_t^{hl}	0.60	1.00
2001/4-	V_t^w	1.00		v_t^w	1.00	
2005/2	V_t^{hl}	0.90	1.00	v_t^{hl}	0.64	1.00

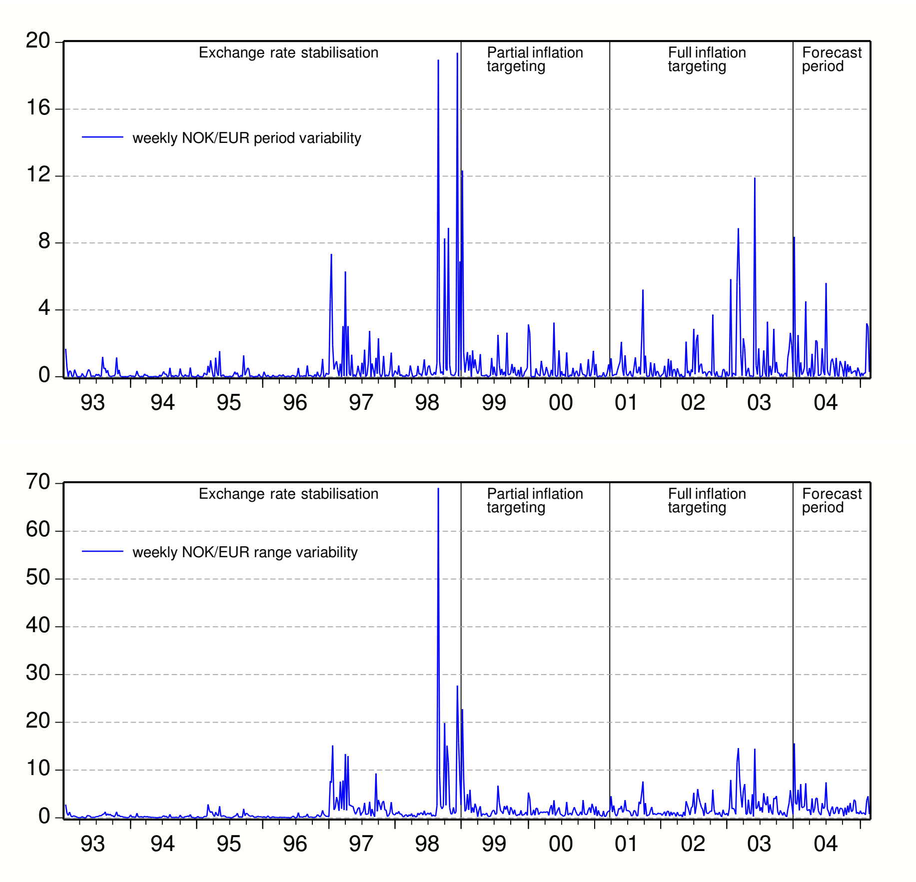

Figure 2.1: Weekly NOK/EUR period and range variability from 8 January 1993 to 25 February 2005 (633 weekly non-missing observations)

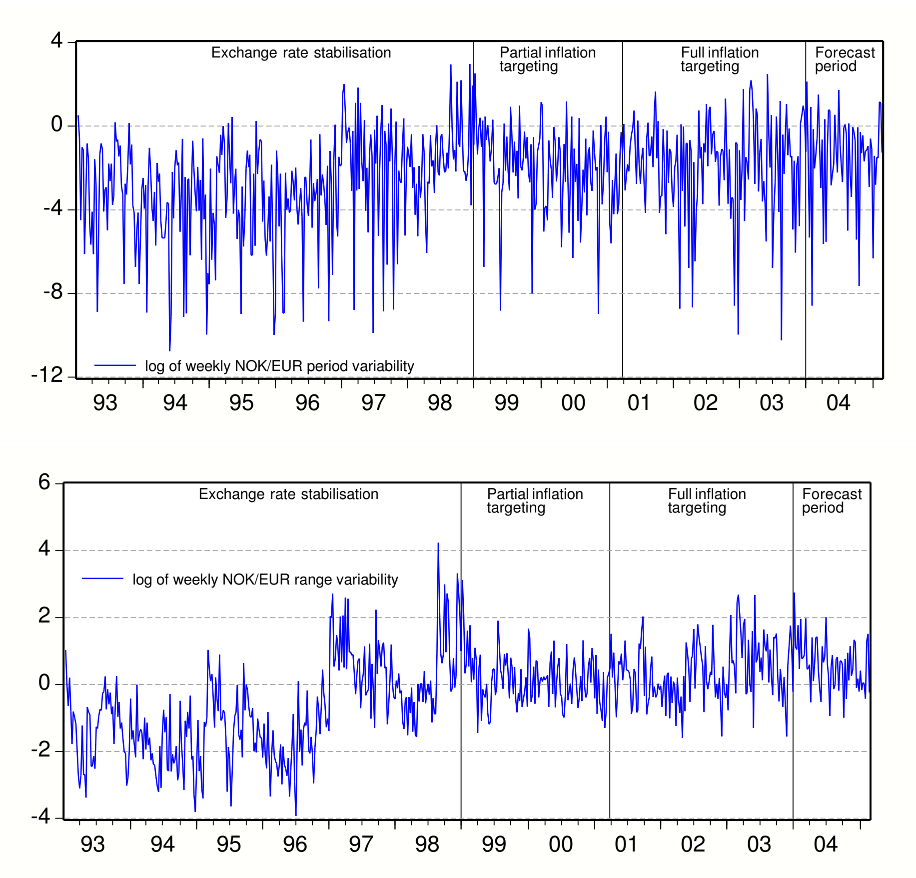

Figure 2.2: Log of weekly NOK/EUR period and range variability from 8 January 1993 to 25 February 2005 (633 weekly non-missing observations)

Chapter 3

Exchange rate variability and market activity

This chapter is a substantially revised version of Bauwens, Rime and Sucarrat (2006).

3.1 Introduction

If exchange rates walk randomly and if the number of steps depends positively on the extent of market activity, then increased market activity should increase exchange rate variability. This chain of reasoning is the economic essence of the socalled "mixture of distribution hypothesis" (MDH) associated with Clark (1973), who suggested that market activity—measured by volume—acts as a proxy for the number of information events. Since changes in market activity also reflect calendar effects (holidays, say) and institutional changes we will not restrict ourselves to Clark's explanation, however. Moreover, as argued by Tauchen and Pitts (1983), a general increase in market activity might have the opposite effect on volatility than that suggested by Clark (1973). In terms of the random walk metaphor, a general increase in market activity might reflect increased liquidity which could lead to smaller steps. So the direction of the overall effect is not certain.

When Karpoff (1987) surveyed the relationship between financial price variability and trading volume during the mid-eighties, only one of the nineteen studies he cited was on exchange rates. The increased availability of data brought by the nineties has changed this, and the ten studies that we summarise in table 3.1 are only a subset of the currently available studies that directly or indirectly investigate the relationship between exchange rate variability and market activity. Nevertheless, our study of Norwegian weekly exchange rate variability from 1993 to 2003 adds to the literature in several ways. First, our study spans more than a decade covering three different exchange rate regimes. Second, not only do we find that the impact of week-to-week changes in market activity on exchange rate variability is positive and statistically significant, but in addition parameter stability analysis suggests the impact is relatively stable across the different exchange rate regimes. Finally, our results do not support the hypothesis that an increase in the general level of market activity—for example due to an increase in the number of traders—reduces exchange rate variability. On the contrary, some of our results suggest within-period variability increases.

The rest of this chapter is organised as follows. In section 3.2, we review the link between exchange rate variability and market activity, discuss measurement issues, and present our data and the other economic variables that we include in our statistical analysis. Section 3.3 contains our empirical results, whereas section 3.4 concludes.

3.2 Theory and data

This section contains three subsections. Subsection 3.2.1 reviews the link between variability and market activity, whereas subsection 3.2.2 presents our market activity data (quote frequency) and explains how we use them in measuring variation in market activity. Finally, subsection 3.2.3 motivates and describes the other variables which we include in the empirical part.

3.2.1 Exchange rate variability and market activity

Denoting by s_t the log of an exchange rate, a simple formulation of the framework in which the relation between financial variability and market activity often is analysed can be stated as

$$\Delta s_t = \sum_{n=1}^{N(t)} \Delta s_n, \qquad n = 1, \ldots, N(t), \qquad s_0 = s_{N(t-1)}, \tag{3.1}$$

$$\{\Delta s_n\} \text{ IID}, \qquad \Delta s_n \sim N(0,1), \tag{3.2}$$

$$\frac{\partial E[N(t)|\nu_t]}{\partial \nu_t} > 0, \tag{3.3}$$

see subsection 2.1 for a fuller explanation of the subscript notation. The first line (3.1) states that the price increment of period t is equal to the sum of the intra-period increments where $N(t)$ is the number of increments in period t, (3.2) is a random walk hypothesis (any "random walk" hypothesis would do), and (3.3) states that the mean of the number of intra-period increments $N(t)$ conditioned on the extent of market activity ν_t in period t is strictly increasing in ν_t. Examples of variables that are believed to increase market activity are an increase in the number of traders and the arrival of relevant information. However, circumstances that are likely to increase market activity might also have a converse effect by affecting the size—in absolute value—of the increments. For example, Tauchen and Pitts (1983) argue (in a nutshell) that an increase in the number of traders, which is believed to result in increased liquidity and thus increased market activity in terms of volume, reduces the size of the intra-period increments. Here this is akin to replacing (3.2) with (say)

$$\Delta s_n = \sigma_n(\nu_n) z_n, \qquad \sigma_n' < 0, \qquad \{z_n\} \text{ IID}, \qquad z_n \sim N(0,1). \tag{3.4}$$

In other words, increased market activity produces two counteracting effects. One effect would tend to reduce period variability through the negative impact on the size of the intra-period increments, whereas the other effect would tend to increase period variability by increasing the number of increments. Taking (3.1) together with (3.3) and (3.4) as our starting point there is thus two possibilities:

$$\frac{\partial \text{Var}(\Delta s_t | \nu_t)}{\partial \nu_t} > 0 \tag{3.5}$$

$$\frac{\partial \text{Var}(\Delta s_t | \nu_t)}{\partial \nu_t} < 0. \tag{3.6}$$

In words, the first hypothesis states that increased market activity increases period variability, whereas the second holds the opposite. That (3.5) is the case is generally suggested by table 3.1, whereas (3.6) is suggested by Tauchen and Pitts (1983). However, it should be noted that the empirical results of Jorion (1996) and Bjønnes et al. (2003) do not support the hypothesis that a general increase in market activity reduces variability.

3.2.2 Measuring variation in market activity

Several types of data have been used in order to construct measures of exchange rate market activity and of variables that are believed to have an impact on exchange rate market activity. Some of these data include transaction volume, the number of transacted contracts, quote frequency and samples from the news-screens of Reuters or Telerate. In this study we make use of quote frequency as our measure of market activity. More precisely, before 1 January 1999 our quote series consists of the number of BID NOK/DEM quotes per week, and after 1 January 1999 it consists of the number of BID NOK/EUR quotes per week. The source of the quote frequency series is Olsen Financial Technologies and we denote the log of the number of quotes in week t by q_t. However, it should be noted that we have adjusted the series for two changes in the underlying data collection methodology—see the data appendix for details. Graphs of q_t and Δq_t are contained in figure 3.1.

Empirical studies of the relationship between financial variability and variation in market activity goes back at least to Clark (1973). Clark and subsequent studies, however, included market activity measures or its log-transformation directly in their regressions, see for example Epps and Epps (1976), and Tauchen and Pitts (1983). The drawback of this is that one may not distinguish between the impacts two different types of variation in market activity may have on financial variability. The two different types of variation in market activity may be referred to as "period-to-period" or "short-term" or "unexpected" variation in market activity on the one hand, and "general" or "long-term" or "expected" variation in market activity on the other. The "expected" vs. "unexpected" terminology is due to Bessembinder and Seguin (1992). Week-to-week variation in market activity is an example of the first type and is compatible with Clark's (1973) suggestion that variation in information arrival is a major determinant of the variation in market activity. A general increase in liquidity due to (say) an increase in the number of traders, which Tauchen and Pitts (1983) suggested would have a negative effect on variability and which would typically manifest itself as a general shift upwards in volume, is an example of the second type. If q_t denotes the log of market activity measure in week t, then a straightforward decomposition is to define short-term variation as $\Delta q_t = q_t - q_{t-1}$ and long-term variation as q_{t-1}, since by definition $q_t = \Delta q_t + q_{t-1}$. The short-term component Δq_t has a straightforward and intuitive economic interpretation, namely the relative increase or decrease in market activity compared with the previous period. Similarly, if q_{t-1} is sufficiently serially correlated with previous lags, that is, with q_{t-2}, with q_{t-3} and so on, then its economic interpretation is the "general" or "long-term" level of market activity. In terms of a conditional heteroscedasticity model of a financial return series the objective of q_{t-1} is to explain autoregressive heteroscedasticity, whereas the objective of Δq_t is to explain non-autoregressive or short-term heteroscedasticity.

The drawback of using q_{t-1} as a measure of long-term variation in market activity is that it might be a noisy measure. One solution is therefore to replace q_{t-1} with a smoothed expression, see for example Bessembinder and Seguin (1992), Hartmann (1999), and Bjønnes et al. (2005). Specifically, these studies use a two-step ARMA decomposition proposed by Bessembinder and Seguin (1992) which in our context would consist of first fitting a model to Δq_t and then using the fitted values of q_t (not the fitted values of Δq_t) as a measure of long-term variation. There are at least two drawbacks with this procedure. First, it might lead to a socalled "generated regressor" problem, see Pagan (1984). Second, it is not given that the optimal predictor of the level of market activity using past values of market activity is the optimal predictor of the general level of financial volatility. For these reasons we also consider a simple alternative measure of long-term variation, namely simple averages of past values. For instance, the average of log of quote frequency q_t using two past values is equal to $(q_{t-1} + q_{t-2})/2$ and is denoted $\bar{q}^2_{t-1}$, the average using three values is equal to $(q_{t-1} + q_{t-2} + q_t)/3$ and is denoted $\bar{q}^3_t$, and so on.

3.2.3 Other impact variables

Other economic variables may also influence variability and should be controlled for. The first economic variable is a measure of general currency market turbulence and is measured through EUR/USD-variability. If $m_t = \log\ (\text{EUR/USD})_t$, then Δm_t denotes the weekly return of EUR/USD, M^w_t stands for weekly variability, and m^w_t is its log-counterpart.[13] The petroleum sector plays a major role in the Norwegian economy, so it makes sense to also include a measure of oilprice variability. If the log of the oilprice is denoted o_t, then the weekly return is Δo_t, and weekly variability is O^w_t with o^w_t as its log-counterpart. We proceed similarly for Norwegian and US stock market variables. If x_t denotes the log of the main index of the Oslo stock exchange, then the associated variables are Δx_t, X^w_t and x^w_t. In the US case u_t is the log of the New York stock exchange (NYSE) index and the associated variables are Δu_t, U^w_t and u^w_t. The motivation behind these variables is that financial markets are interlinked directly through international investors who continuously compare the returns and risk opportunities associated with each market, and psychologically since bearish and bullish sentiments may transmit from one market to another.[14] To account for the possibility of skewness in r_t, that is, that exchange rate depreciations tend to be larger in absolute value than appreciations, and asymmetries in r_t, that is, that major falls in the value of the exchange rate tends to bring about a subsequent period of higher variability (for a common stock this implies higher leverage), we use the lagged return r_{t-1} for the latter, and an impulse dummy ia_t equal to 1 when returns are positive and 0 otherwise for the former. We also include variables intended to account for the impact of holidays and seasonal variation. These are denoted h_{lt} with $l = 1, 2, \ldots, 8$, see the appendix for further details. A step dummy sd_t equal to 0 before 1997 and 1 after intended to account for what seems to be a structural break around the beginning of 1997, is also included. Admittedly we do not test for such a break (tests were carried out in earlier versions of the study, and in the current version we encountered numerical problems), but not including sd_t results in residual serial correlation and thus provides a partial justification at least. A Russian moratorium dummy id_t equal to 1 in one of the weeks following the Russian moratorium (the week containing Friday 28 August 1998 to be more precise) and 0 elsewhere is included in the range variability regressions. Its motivation is that it is needed for residual normality in all the log of range variability regressions, and in some it also removes residual serial correlation.

The foreign interest-rate variables that we include are constructed using an index made up of the short term market interest-rates of the EMU countries. Specifically, if IR_t^{emu} denotes this interest-rate index then we include a variable that is denoted ir_t^{emu} and which is defined as $(\Delta IR_t^{emu})^2$. The Norwegian interest-rate variables that are included are constructed using the main policy interest rate variable of the Norwegian central bank, and the reason we do not use market interest rates is that their impact seems to be substantially unstable—even within each exchange rate regime. The construction of the interest-rate variables reflect the regime changes that took place over the sample period, since an interesting question is whether policy interest rate changes contributed differently to exchange rate volatility in the partial and full inflation targeting periods, respectively.[15] Let F_t denote the main policy interest rate in percentages and let ΔF_t denote the change from the end of one week to the end of the next. Furthermore, let I_a denote an indicator function equal to 1 in the period 1 January 1999 - Friday 30 March 2001 and 0 otherwise, and let I_b denote an indicator function equal to 1 after 30 March 2001 and 0 before. Then $\Delta F_t^a = \Delta F_t \times I_a$ and $\Delta F_t^b = \Delta F_t \times I_b$, respectively, and f_t^a and f_t^b stand for $|\Delta F_t^a|$ and $|\Delta F_t^b|$, respectively. Effectively this means that we do not include interest rate variables in the first regime, and the reason for this is to avoid simultaneity issues since Norges Bank actively used the interest rate in their exchange rate management during the first regime.

3.3 Empirical results

This section proceeds in four steps and all estimates are of models nested within the general specification

$$\begin{aligned} v_t = b_0 + \sum_{k=1}^{5} b_k v_{t-k} + b_6 \bar{q}_{t-1}^{6} + b_7 \Delta q_t + b_8 m_t^w + b_9 o_t^w + \\ b_{10} x_t^w + b_{11} u_t^w + b_{12} f_t^a + b_{13} f_t^b + b_{14} ir_t^{emu} + b_{15} sd_t + \\ b_{16} ia_t + b_{17} r_{t-1} + \sum_{l=1}^{8} b_{17+l} h_{lt} + b_{26} id_t + e_t, \quad (3.7) \end{aligned}$$

where v_t stands for the log of variability in question (period or range) and e_t is the error term. It should be noted that in all estimations only the sample ranging from 8 January 1993 to 26 December 2003 (573 observations) is used. In the first two subsections we shed light on the impact of market activity paying particular attention to the concerns and results of the literature on the relation between market activity and financial price variability. The third subsection tests the impact of market activity controlling for the impact of other variables, whereas the final subsection explores whether the impact of market activity depends on exchange rate regime.

3.3.1 The uncontrolled impact of market activity

In this subsection we study the impact of market activity on exchange rate volatility without controlling for other variables. The motivation for this is that several important contributions (including Clark (1973), Epps and Epps (1976), and Tauchen and Pitts (1983)) have put forward the view that market activity is a main determinant of variability, and that other variables essentially work through (in the sense that they determine) or proxy market activity. One way of shedding light on this issue is to run regressions without controlling for other variables. Table 3.2 contains estimates of the specifications

$$v_t^w = b_0 + b_6\bar{q}_{t-1}^6 + b_7\Delta q_t + e_t, \tag{3.8}$$

$$v_t^{hl} = b_0 + b_6\bar{q}_{t-1}^6 + b_7\Delta q_t + e_t. \tag{3.9}$$

The variable $\bar{q}_{t-1}^6$ is chosen as measure of long term market activity, since it yields a higher R^2 in (3.8) and (3.9) than the other measures of long term market activity discussed in subsection 3.2.1. Both $\bar{q}^6_{t-1}$ and Δq_t are significant in both regressions, but one should be careful in interpreting these results since both specifications are misspecified in the sense that they exhibit substantially serially correlated residuals. The most important thing to take away from table 3.2 is to note that the market activity variables alone are unable to adequately account for the time-varying variability.

3.3.2 Variability persistence vs. market activity

The purpose of this subsection is to shed further light on the impact on market activity before controlling for other variables, and again the motivation is concerns and hypotheses put forward in important contributions to the literature. The specific hypothesis that we aim to shed light on is the suggestion that persistence in exchange rate variability can be explained by persistence in the level of market activity. For example, in a much-cited study Lamoureux and Lastrapes (1990) find that the volatility persistence coefficients in GARCH(1,1) models of common stock returns fall and often become insignificant when volume is added as an explanatory variable in the conditional variance equation.

Table 3.3 contains estimates of the autoregressions

$$v^w_t = b_0 + b_2(v^w_{t-2} + v^w_{t-3}) + b_{15}sd_t + e_t, \tag{3.10}$$

$$v^{hl}_t = b_0 + b_1(6v^{hl}_{t-1} + 3v^{hl}_{t-2} + 2v^{hl}_{t-3}) + b_{15}sd_t + b_{26}id_t + e_t, \tag{3.11}$$

and are analogous to GARCH(1,1) models with no conditional mean of financial returns. The specific lag-structures are obtained through simplification of a general specification containing five lags. Contrary to the previous table the estimation results suggest lags of variability together with the structural step dummy sd_t are capable of accounting for time-varying variability in both regressions in the sense that residuals are uncorrelated and homoscedastic. The Russian Moratorium (August 1998) impulse dummy id_t is needed in (3.11) for residuals to be serially uncorrelated, but including this variable in (3.9) does not lead to uncorrelated residuals.

In the current context insignificance of b_2 in (3.10) and b_1 in (3.11) when including $\bar{q}^6_{t-1}$ and Δq_t would be analogous to Lamoureux and Lastrapes' finding that the persistence coefficients in a GARCH(1,1) become insignificant when including volume as explanatory variable. Table 3.4 contains estimates of such autoregressions augmented by the market activity variables:

$$v^w_t = b_0 + b_2(v^w_{t-2} + v^w_{t-3}) + b_6\bar{q}^6_{t-1} + b_7\Delta q_t + b_{15}sd_t + e_t, \tag{3.12}$$

$$v^{hl}_t = b_0 + b_1(6v^{hl}_{t-1} + 3v^{hl}_{t-2} + 2v^{hl}_{t-3}) + b_6\bar{q}^6_{t-1} + b_7\Delta q_t + b_{15}sd_t + b_{26}id_t + e_t. \tag{3.13}$$

Comparing the estimates of b_1 and b_2 in table 3.3 with those of table 3.4 does not suggests that persistence falls, since b_1 and b_2 remain significant at all conventional significance levels. Indeed, contrary to Lamoureux and Lastrapes' findings the coefficient value increases in both cases, which suggests higher persistence. With respect to the market activity variables, before controlling for the impact of other variables the long-term market activity measure $\bar{q}^6_{t-1}$ is insignificant in the weekly volatility specification (3.12) and significant at 3% in the range volatility specification (3.13), whereas the short-term market activity measure

Δq_t is significant at all conventional levels and carries the expected sign in both specifications.

3.3.3 The impact of market activity on variability

In contrast with the previous two subsections the objective of this subsection is to study to what extent the market activity measures $\bar{q}^6_{t-1}$ and Δq_t explain exchange rate variability when controlling for other explanatory variables. To this end we estimate parsimonious specifications obtained through simplifications of (3.7) above while keeping $\bar{q}^6_{t-1}$ and Δq_t fixed, that is, not removing them if insignificant at any level. The resulting parsimonious specifications are

$$v^w_t = b_0 + b_2(v^w_{t-2} + v^w_{t-3}) + b_6\bar{q}^6_{t-1} + b_7\Delta q_t + b_{10}(x^w_t + u^w_t) + b_{13}f^b_t + b_{15}sd_t + e_t, \quad (3.14)$$

$$v^{hl}_t = b_0 + b_1(6v^{hl}_{t-1} + 3v^{hl}_{t-2} + 2v^{hl}_{t-3}) + b_6\bar{q}^6_{t-1} + b_7\Delta q_t + b_8 m^w_t + b_{13}f^b_t + b_{14}ir^{emu}_t + b_{15}sd_t + b_{16}ia_t + b_{18}h_{1t} + b_{25}h_{8t} + b_{26}id_t + e_t, \quad (3.15)$$

and the estimation results are contained in table 3.5. The results can be summarised in four points:

1. Short-term market activity. The estimated impacts of changes in short-term market activity Δq_t carry the hypothesised positive sign and are significant at all conventional levels in both regressions. Moreover, compared with the estimates in table 3.4 their values are relatively similar. In other words, adding regressors does not change estimates substantially.

2. Long-term market activity. The estimated impact of changes in long-term market activity as measured by $\bar{q}^6_{t-1}$ is negative and insignificant in the parsimonious weekly period variability specification (3.14), whereas it is positive and significant at 2% in the parsimonious weekly range variability specification (3.15).

3. Interest rate changes. One would expect that policy interest rate changes in the full inflation targeting period—as measured by f^b_t—increase variability, whereas the hypothesised effect in the partial inflation period—as measured by f^a_t—is lower or at least uncertain. The results in table 3.5 support this since f^a_t does not appear in any of the specifications (it has been removed due to insignificance), and since the results suggest a positive and significant contemporaneous impact (in absolute value) in the full inflation targeting period. Changes in the short term market interest rate of the EMU countries ir^{emu}_t have a significant impact in both variability specifications.

4. Other. The effect of general currency market variability m^w_t is significant only in the range variability specification, whereas the effect of oilprice variability o^w_t is not significant in either specification. This might come as a surprise since Norway is a major oil-exporting economy—currently third after Saudi-Arabia and Russia, and since the petroleum sector plays a big part in the Norwegian economy. A possible reason for this is that the impact of oilprice variability is non-linear in ways not captured by our measure, see Akram (2004).[16] The effects of Norwegian and US stock market variability are significant in the weekly variability specification (3.7), and the restriction that they are equal is not rejected. In the range variability specification on the other hand only the measure of Norwegian stock market variability is retained after simplification. The skewness variable ia_t is significant in the range variability regression but not in the weekly variability regression. This suggests range variability is higher when the NOK/EUR depreciates over the week. The New Year's Eve variable h_{1t} and the Christmas variable h_{8t} in the range variability specification are the only significant holiday variables in the parsimonious equations. The former is significant and negative as expected, which suggests Δq_t does not fully account for the impact of "hang-over" day on range variability. The latter on the other hand is significant and positive. This could be due to Δq_t having a lower effect on variability than suggested by its parameter estimate in weeks with Christmas, so that h_{8t} essentially adjusts for this parameter instability.

3.3.4 Stability analysis of the impact of market activity

Norway experienced three different exchange rate regimes over the estimation period and a question of interest is to what extent the impacts of $\bar{q}^6_{t-1}$ and Δq_t depend on regime. Figures 3.2 and 3.3, and tables 3.6 and 3.7 aim at shedding light on this question.

The figures contain recursive OLS estimates of the coefficients of $\bar{q}^6_{t-1}$ and Δq_t in (3.14) and (3.15), respectively. The recursive estimates of the coefficient of $\bar{q}^6_{t-1}$ in (3.14) stay close to the zero-line and cross it several times. However, there are no clear indication that the changes in direction correspond to or occur in the vicinity of a regime change. In (3.15) the recursive estimates of the coefficient of $\bar{q}^6_{t-1}$ are positive throughout and do not cross the zero-line, and appear to be trending upwards until just before the end of the first regime. Then they appear to change direction and exhibit a slight general trend downwards. This suggests that the impact of $\bar{q}^6_{t-1}$ was higher in the first regime, but one should be careful in pursuing this interpretation too far without further investigation, since the end of the first regime also corresponds with the introduction of the Euro. In other words, the change could be to due to institutional changes unrelated to long-term market activity. The recursive estimates of the coefficient of Δq_t are relatively stable in both specifications in the sense that they do not cross the zero line, with the estimates in the range specification being more stable than the estimates in the weekly in the sense that the difference between the maximum and minimum values is larger in the weekly case. In figure 3.2 there is some support in the recursive estimates that the impact of Δq_t in (3.14) drops in the beginning of 1999. Again, however, one should be careful in attributing this to the change in regime without further investigation, since Δq_t acquires its maximum and its 3rd lowest values in the first weeks of 1999. There is also some support in figure 3.3 that the impact of Δq_t in (3.15) increases until the end of the first regime, and trends downward. However, again and for the same reasons—the extreme value of Δq_t in the first weeks of 1999—one should be careful in attributing this to the change in regime without further investigation.

The tables 3.6 and 3.7 contain subsample estimates of (3.14) and (3.15), respectively, where each sample corresponds to an exchange rate regime. We will only comment on the coefficient estimates of $\bar{q}^6_{t-1}$ and Δq_t. For (3.14) the coefficient of $\bar{q}^6_{t-1}$ varies substantially. In the first regime its value is 0.030 and insignificant with a p-value of 93%, then it jumps to 1.009 and acquires significance at the 10% level before it becomes negative and insignificant in the third regime. This is an indication of substantial instability of the impact of $\bar{q}^6_{t-1}$ and calls for further investigation before any firm conclusions can be made. For (3.15) in table 3.7 the coefficient of $\bar{q}^6_{t-1}$ is more stable in the sense that all estimates are positive. However, there are signs of instability. They do exhibit a slight trend downwards with the regime III estimate being less than half of the estimate in the first regime, and only in the first regime is the coefficient significant at the 10% level. The coefficient estimates of Δq_t are in comparison more stable. For both (3.14) and (3.15) estimates remain positive and significant at 9% in all the subsamples. The most pronounced signs of instability occur in regime II for (3.14), since here the coefficient estimate is notably lower than in the two other regimes. For (3.15) the most pronounced sign of instability is the estimate in regime I compared with the estimate in the two other subsamples. The estimate in the first regime is more than 50% higher compared with the two other regimes.

3.4 Conclusions

Our study of weekly Norwegian exchange rate variability sheds new light on the impact of market activity in several ways. We find that the impact of short-term change in market activity, as measured by relative week-to-week changes in quoting frequency, is positive and statistically significant for two different definitions of variability, and that the impact is relatively stable across three different exchange rate regimes for both definitions of variability. One might have expected that the effect would increase with a shift in regime from exchange rate stabilisation to partial inflation targeting, and then to full inflation targeting, since the Norwegian central bank actively sought to stabilise the exchange rate before the full inflation targeting regime. In our data however there are no clear breaks, shifts upwards nor trends following the points of regime change. Indeed, the instability that there is suggests the opposite, namely that the impact was higher in the first regime when the Norwegian central bank actively sought to stabilise the exchange rate. Possible reasons for this are that traders were more reactive or less heterogeneous in their reactions to events of relevance, and that the underlying data (Reuters's quote frequency) were a better measure of market activity during the first regime, since the dealing systems of Reuters at that time had a more prominent position in the currency markets. Our results also support to some extent the hypothesis that changes in long-term market activity, as measured by the average level of quoting frequency in the previous six weeks, increases weekly range variability. However, our results do not support the hypothesis that it increases weekly period variability. We also find some evidence that impact of long-term market activity on range variability depends on exchange rate regime. In particular, that the impact is higher in the first regime and lower (and possibly insignificant) in the two subsequent regimes. Finally, our results do not suggest that the persistence in variability can be explained by persistence in the level of volume.

Our results suggest several areas for further investigation. First, there are several limitations with the quote frequency data in constructing measures of market activity. Using alternative data in shedding light on the same questions would therefore be beneficial (one such study is contained in chapter 5 of this thesis). Second, our variability data exhibit what seems to be a structural break, that is, a shift upwards, around the end of 1996 and/or beginning of 1997. The exact nature and timing of this event is not well understood. According to van Dijk et al. (2005) several non-Euro exchange rates against the USD experienced a break in unconditional volatilities (their study was conducted by means of a dynamic conditional correlation framework), and the NOK/USD exchange rate is the one that exhibits the largest shift upwards (50%).

They attribute the break to a European Council meeting in December 1996 in which a decision regarding the EMU was taken, and that this was pronounced in the Norwegian case because of a change in the intervention policy of the Norwegian Central Bank. According to Bjønnes et al. (2005) on the other hand the events at the end of 1996/beginning of 1997 were due to a speculative attack by foreign speculators. More research is therefore needed in order to understand the exact nature, timing and reasons for the shift upwards in variability around the end of 1996/beginning of 1997. Finally, a third area for further research is the impact of Norwegian market interest rates. Although the policy interest rate affect interest rate bearing securities, the actors in foreign exchange markets are mainly concerned with the money market interest rates. We failed to include Norwegian money market interest rate variables because their impact are very unstable and not robust to slight changes in specification even within each regime. Further understanding of the relation between policy interest rate changes, money market interest rates and exchange rate variability is therefore necessary.

Table 3.1: A summary of empirical studies that investigate directly or indirectly the impact of market activity on exchange rate variability.

Publication	Data	Period	Impact?
Grammatikos and Saunders (1986)	Daily currency futures contracts (DEM, CHF, GBP, CAD and JPY) denominated in USD	1978-1983	Yes
Goodhart (1991)	Intradaily quotes (USD against GBP, DEM, CHF, JPY, FRF, NLG, ITL, ECU) and Reuters' news-headline page	14/9-15/9 1987	No
Goodhart (2000)	Intradaily quotes (USD against GBP, DEM, JPY, FRF, AUD) and Reuters' news-headline pages	9/4-19/6 1989	No
Bollerslev and Domowitz (1993)	Intradaily USD/DEM quotes and quoting frequency	9/4-30/6 1989	No
Demos and Goodhart (1996)	Intradaily DEM/USD and JPY/USD quotes and quoting frequency	5 weeks in 1989	Yes
Jorion (1996)	Daily DEM/USD futures and options	Jan. 1985-Feb. 1992	Yes
Melvin and Xixi (2000)	Intradaily DEM/USD and JPY/USD quotes, quoting frequency and Reuters' headline-news screen	1/12 1993-26/4 1995	Yes
Galati (2003)	Daily quotes (USD against JPY and seven emerging market currencies) and trading volume	1/1 1998-30/6 1999	Yes
Bauwens et al. (2005)	Intradaily EUR/USD quotes, quoting frequency and Reuters' news-alert screens	15/5 2001-14/11 2001	Yes
Bjønnes et al. (2003)	Daily SEK/EUR quotes and transaction volume	1995-2002	Yes

Table 3.2: Regressions of log of variability on market activity variables

Regressor	(3.8)		(3.9)	
	Est.	*Pval.*	*Est.*	*Pval.*
Const.	-7.740	0.00	-8.486	0.00
$\bar{q}^6_{t-1}$	0.697	0.00	1.082	0.00
Δq_t	1.147	0.00	0.703	0.00
R^2	0.04		0.19	
AR_{1-10}	46.4	0.00	281.08	0.00
$ARCH_{1-10}$	8.72	0.56	117.03	0.00
Het.	5.86	0.21	17.12	0.00
Hetero.	6.68	0.25	18.57	0.00
JB	95.56	0.00	3.01	0.22
Obs.	567		567	

Note: The estimation period is 8 January 1993 - 26 December 2003. Computations are in EViews 5.1 and estimates are OLS with robust standard errors of the Newey and West (1987) type. *Pval* stands for p-value and corresponds to a two-sided test with zero as null, AR_{1-10} is the χ^2-form of the Lagrange-multiplier test for serially correlated residuals up to lag 10, $ARCH_{1-10}$ is the χ^2-form of the Lagrange-multiplier test for serially correlated squared residuals up to lag 10, *Het.* and *Hetero.* are White's (1980) heteroscedasticity tests without and with cross products, respectively, *JB* is the Jarque and Bera (1980) test for non-normality in the residuals, and *Obs.* is the number of non-missing observations.

Table 3.3: Autoregressions of log of variability

Regressor	(3.10)		(3.11)	
	Est.	*Pval.*	*Est.*	*Pval.*
Const.	-2.966	0.00	-0.666	0.00
$v^w_{t-2} + v^w_{t-3}$	0.088	0.00		
$6v^{hl}_{t-1} + 3v^{hl}_{t-2} + 2v^{hl}_{t-3}$			0.052	0.00
sd_t	1.429	0.00	0.798	0.00
id_t			4.038	0.00
R^2	0.13		0.59	
AR_{1-10}	3.74	0.96	12.60	0.25
$ARCH_{1-10}$	10.01	0.44	7.40	0.69
Het.	5.72	0.13	5.79	0.22
Hetero.	6.74	0.15	5.92	0.31
JB	117.04	0.00	21.19	0.00
Obs.	569		569	

Notes: Standard errors are of the White (1980) type. Otherwise see table 3.2.

Table 3.4: Auto-regressions of log of variability augmented with market activity variables

Regressor	(3.12)		(3.13)	
	Est.	*Pval.*	*Est.*	*Pval.*
Const.	-2.825	0.07	-1.909	0.00
$v^{w}_{t-2} + v^{w}_{t-3}$	0.091	0.00		
$6v^{hl}_{t-1} + 3v^{hl}_{t-2} + 2v^{hl}_{t-3}$			0.054	0.00
$\bar{q}^{6}_{t-1}$	-0.020	0.92	0.176	0.03
Δq_t	1.077	0.00	0.738	0.00
sd_t	1.446	0.00	0.685	0.00
id_t			3.926	0.00
R^2	0.15		0.63	
AR_{1-10}	6.02	0.81	12.24	0.27
$ARCH_{1-10}$	7.98	0.63	11.26	0.34
Het.	10.22	0.18	7.69	0.46
Hetero.	14.48	0.34	19.98	0.13
JB	110.56	0.00	13.52	0.00
Obs.	567		567	

Notes: Standard errors are of the White (1980) type. Otherwise see table 3.2.

Table 3.5: Parsimonious specifications of log of variability obtained through simplification of (3.7)

Regressor	(3.14)		(3.15)	
	Est.	*Pval.*	*Est.*	*Pval.*
Const.	-2.413	0.13	-2.134	0.00
$v^w_{t-2} + v^w_{t-3}$	0.079	0.00		
$6v^{hl}_{t-1} + 3v^{hl}_{t-2} + 2v^{hl}_{t-3}$			0.051	0.00
$\bar{q}^6_{t-1}$	-0.086	0.69	0.188	0.02
Δq_t	1.053	0.00	0.789	0.00
m^w_t			0.032	0.03
x^w_t			0.030	0.02
$x^w_t + u^w_t$	0.119	0.00		
f^b_t	3.837	0.00	1.193	0.00
ir^{emu}_t	4.650	0.01	2.878	0.00
sd_t	1.279	0.00	0.727	0.00
ia_t			0.144	0.03
h_{1t}			-0.855	0.00
h_{8t}			0.258	0.01
id_t			3.729	0.00
R^2	0.20		0.66	
AR_{1-10}	3.26	0.97	8.40	0.59
$ARCH_{1-10}$	8.70	0.56	8.60	0.57
Het.	13.42	0.42	20.02	0.46
Hetero.	20.78	0.95	50.93	0.94
JB	118.75	0.00	8.63	0.01
Obs.	567		567	

Note: Standard errors are of the White (1980) type. Otherwise see table 3.2.

Table 3.6: Subsample estimates of (3.14) in each exchange rate regime

Regressor	I		II		III	
	Est.	*Pval.*	*Est.*	*Pval.*	*Est.*	*Pval.*
Const.	-3.424	0.14	-9.618	0.05	0.203	0.94
$v^w_{t-2} + v^w_{t-3}$	0.058	0.08	0.058	0.33	0.127	0.04
$\bar{q}^6_{t-1}$	0.030	0.93	1.009	0.10	-0.238	0.50
Δq_t	1.199	0.01	0.810	0.09	1.272	0.06
$x^w_t + u^w_t$	0.111	0.01	0.136	0.01	0.102	0.13
f^b_t					3.950	0.00
ir^{emu}_t	5.184	0.03	2.298	0.45	6.536	0.07
sd_t	1.503	0.00				
R^2	0.19		0.10		0.16	
AR_{1-10}	5.18	0.88	7.54	0.67	5.44	0.86
$ARCH_{1-10}$	14.32	0.16	4.50	0.92	7.80	0.65
Het.	6.31	0.85	4.66	0.91	11.59	0.48
Hetero.	16.34	0.93	7.84	0.99	37.97	0.08
JB	54.67	0.00	29.48	0.01	31.15	0.00
Obs.	306		118		143	

Note: Standard errors are of the White (1980) type, otherwise see table 3.2. Regime I corresponds to the exchange rate stabilisation regime from the week containing Friday 8 January 1993 to the week containing Friday 25 December 1998, regime II corresponds to the partial inflation targeting regime from the week containing Friday 1 January 1999 to the week containing Friday 30 March 2001, and regime III corresponds to the full inflation targeting regime from the week containing Friday 6 April 2001 to the week containing Friday 26 December 2003. The specifications vary across regimes in order to avoid co-linearity between regressors.

Table 3.7: Subsample estimates of (3.15) in each exchange rate regime

Regressor	I		II		III	
	Est.	*Pval.*	*Est.*	*Pval.*	*Est.*	*Pval.*
Const.	-3.242	0.00	-2.164	0.21	-1.044	0.34
persistence	0.055	0.00	0.045	0.00	0.035	0.00
$\bar{q}^6_{t-1}$	0.353	0.01	0.278	0.22	0.137	0.33
Δq_t	0.983	0.00	0.618	0.00	0.613	0.02
m^w_t	0.043	0.02	0.049	0.12	0.01	0.76
x^w_t	0.027	0.11	0.027	0.30	0.04	0.29
f^b_t					1.389	0.00
ir^{emu}_t	3.000	0.01	1.973	0.09	5.079	0.00
sd_t	0.772	0.00				
ia_t	0.106	0.27	0.225	0.10	0.174	0.25
h_{1t}	-1.361	0.00	-0.783	0.05	-0.295	0.13
h_{8t}	0.416	0.05	0.266	0.05	0.126	0.24
id_t	3.583	0.00				
R^2	0.69		0.29		0.24	
AR_{1-10}	3.14	0.98	19.98	0.03	8.02	0.63
$ARCH_{1-10}$	5.08	0.89	6.04	0.81	15.26	0.12
Het.	24.93	0.13	19.37	0.20	26.65	0.06
Hetero.	55.05	0.32	42.31	0.25	79.29	0.00
JB	5.10	0.08	2.15	0.34	0.68	0.71
Obs.	306		118		143	

Note: The term *persistence* is defined as $6v^{hl}_{t-1} + 3v^{hl}_{t-2} + 2v^{hl}_{t-3}$. For regime II standard errors of the Newey and West (1987) type are used in the tests regarding the coefficients, otherwise see tables 3.2. Regime I corresponds to the exchange rate stabilisation regime from the week containing Friday 8 January 1993 to the week containing Friday 25 December 1998, regime II corresponds to the partial inflation targeting regime from the week containing Friday 1 January 1999 to the week containing Friday 30 March 2001 and regime III corresponds to the full inflation targeting regime from the week containing Friday 6 April 2001 to the week containing Friday 26 December 2003. The specifications vary across regimes in order to avoid co-linearity between regressors.

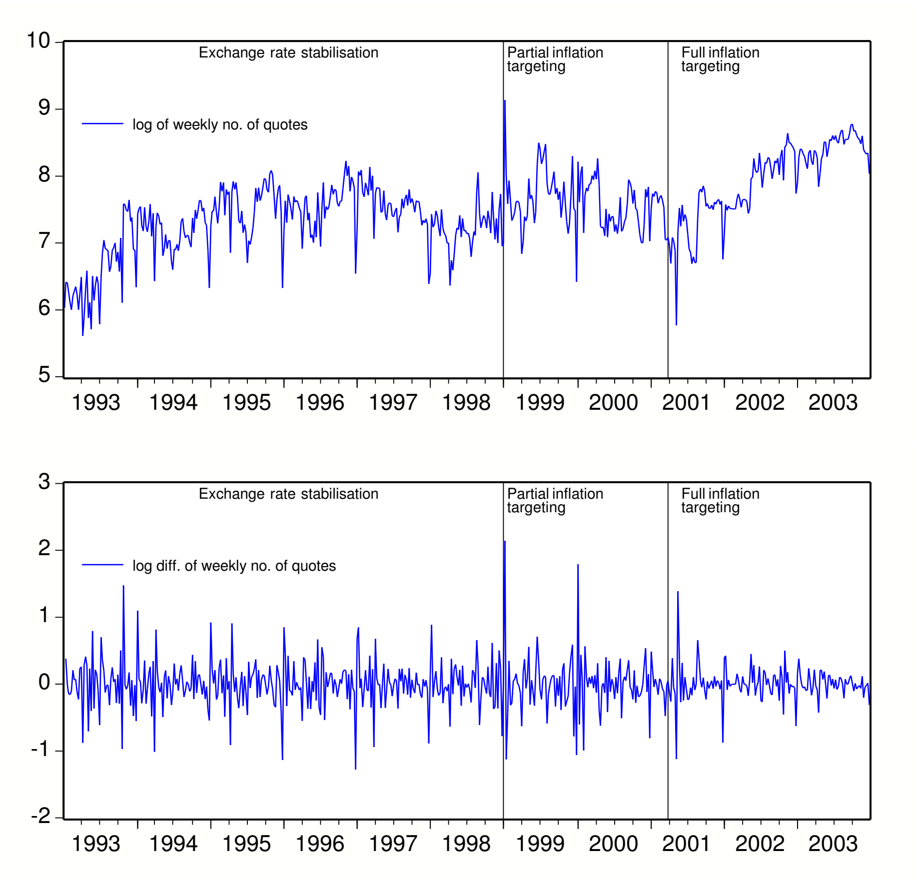

Figure 3.1: The log of weekly number of BID NOK/EUR quotes (BID NOK/DEM before 1999) in the upper graph and the log-difference of weekly quoting in the middle graph.

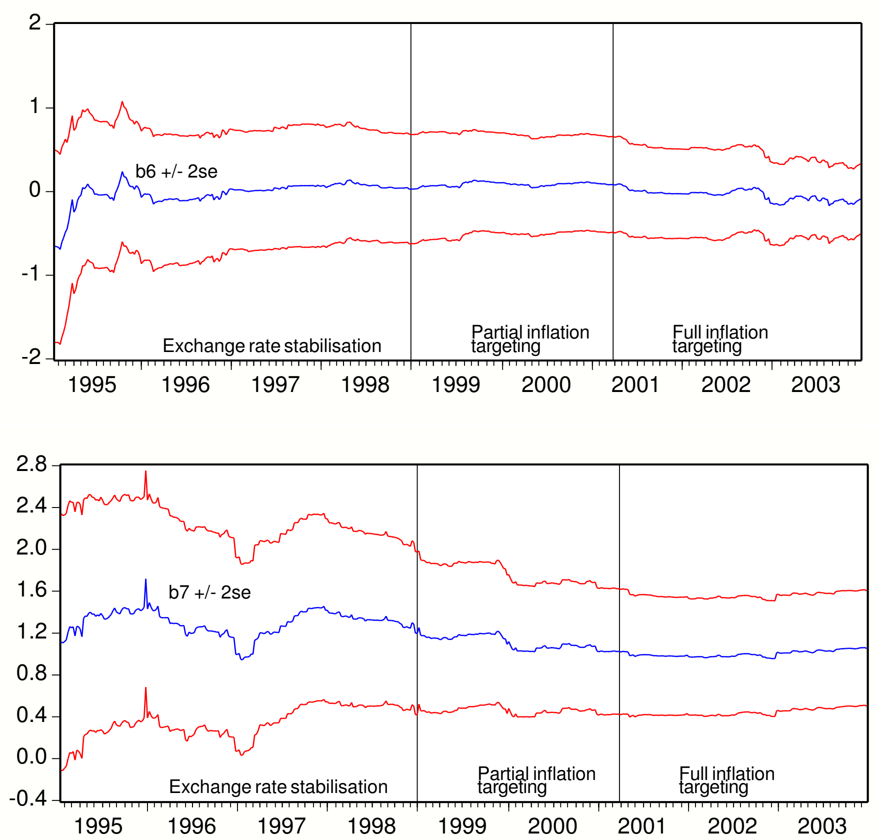

Figure 3.2: Recursive OLS estimates of the impact of $\bar{q}^{6}_{t-1}$ and Δq_t in (3.14). Computations are in PcGive 10.4 with initialisation at observation number 100, which corresponds to the week containing Friday 2 December 1994.

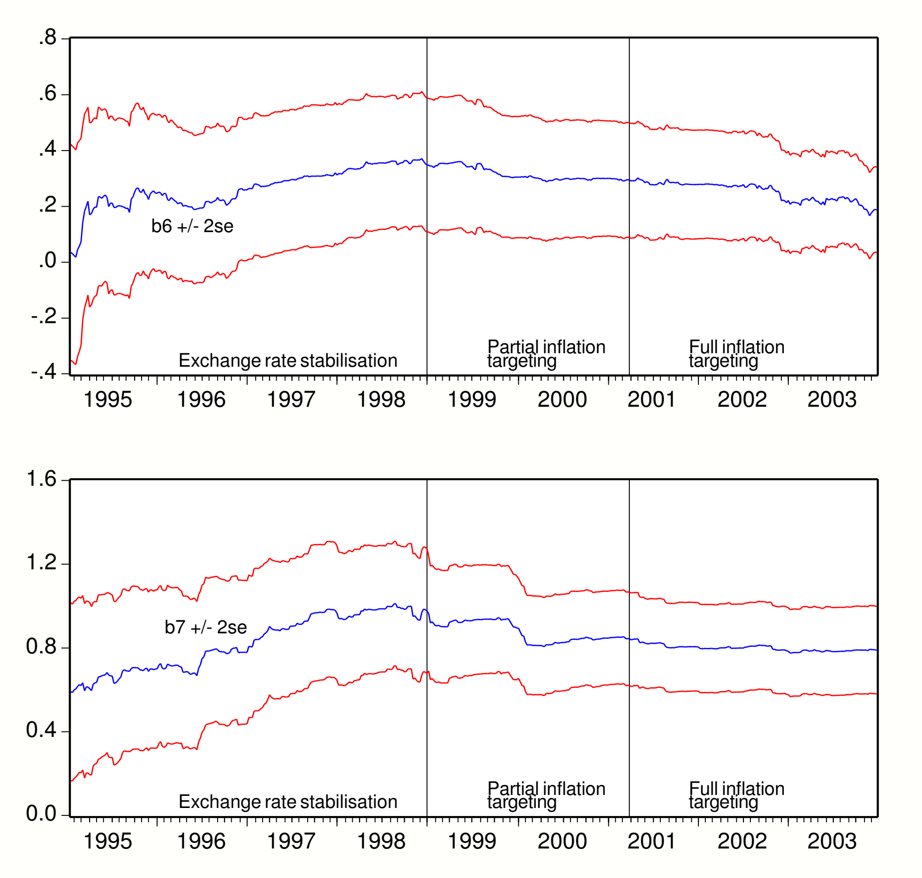

Figure 3.3: Recursive OLS estimates of the impact of $\bar{q}^{6}_{t-1}$ and Δq_t in (3.15). Computations are in PcGive 10.4 with initialisation at observation number 100, which corresponds to the week containing Friday 2 December 1994.

Chapter 4

General to specific modelling of exchange rate volatility: A forecast evaluation

This chapter is a substantially revised version of Bauwens and Sucarrat (2006).

4.1 Introduction

Exchange rate variability and models thereof—in particular volatility models—are of great importance for both businesses and policymakers alike. Businesses for example use volatility models as tools in their risk management and as input in derivative pricing, whereas policymakers use them to acquire knowledge about what and how economic factors impact upon exchange rate variability for informed policymaking. Most volatility models are highly non-linear and thus require complex optimisation algorithms in empirical application. For models with few parameters and few explanatory variables this may not pose unsurmountable problems. But as the number of parameters and explanatory variables increases the resources needed for reliable estimation and model validation multiply. Indeed, this may even become an obstacle to the application of certain econometric modelling strategies and methodologies, as for example argued by McAleer (2005) regarding automated general-to-specific (GETS) modelling of financial volatility.[17] GETS modelling is particularly suited for explanatory econometric modelling since it provides a systematic framework for statistical economic hypothesis-testing, model development and model (re-)evaluation, and the methodology is relatively popular among large scale econometric model developers and proprietors. However, since the initial model formulation typically entails many explanatory variables this poses challenges already at the outset for computationally complex models.

In this study we overcome the computational challenges traditionally associated with the application of the GETS methodology in the modelling of financial volatility by modelling volatility within a single equation exponential model of variability (EMOV), where variability is defined as squared returns. The parameters of interests can thus be estimated with ordinary least squares (OLS) under rather weak assumptions. This enables us to apply GETS to a general specification with, in our case, a constant and twenty four regressors, including lags of log of variability, an asymmetry term, a skewness term, seasonality variables, and economic covariates. Compared with models of the autoregressive conditional heteroscedasticity (ARCH) and stochastic volatility (SV) classes we estimate and simplify our specification effortlessly, and obtain a parsimonious encompassing specification with uncorrelated homoscedastic residuals and relatively stable parameters. Moreover, our out-of-sample forecast evaluation suggests that GETS specifications are especially valuable in conditional forecasting, since the specification that employs actual values on the uncertain information performs particularly well.

Another contribution of this study concerns the evaluation of explanatory economic models of financial volatility. An argument that has gained widespread acceptance lately is that discrete time models of financial volatility should be evaluated against more efficient estimates derived from continuous time models, see for example Andersen and Bollerslev (1998), Andersen et al. (1999), Andersen et al. (2005). In this essay we qualify this view by arguing that the approach is particularly inappropriate in the evaluation of discrete time explanatory economic models of financial volatility.

The rest of this chapter is divided into four sections. In the next section we outline the main ingredients of the GETS methodology. Then we present the models in section 4.3, whereas section 4.4 contains the results of the out-of-sample forecast exercise. Finally we conclude in section 4.5.

4.2 GETS modelling

A fundamental cornerstone of the GETS methodology is that empirical models are derived, simplified representations of the complex human interactions that generate the data. For a suggestion of how a probabilistic formulation of such a process may look like and from which both continuous and discrete models can be obtained as reductions, see definition 4 in chapter 6 of this thesis. Accordingly, instead of postulating a uniquely "true" model or class of models, the aim is to develop "congruent" encompassing models within the statistical framework of choice. The exact definition of congruency is given below, but in brief a congruent model is a theory informed specification that is data-compatible with weakly exogenous conditioning variables with respect to the parameters of interest, and which constitutes a "history-repeats-itself" representation (stable parameters, innovation errors).[18]

In econometric practice GETS modelling proceeds in cycles of three steps: 1) Formulate a general unrestricted model (GUM) which is congruent, 2) simplify the model sequentially in an attempt to derive a parsimonious congruent model while at each step checking that the model remains congruent, and 3) test the resulting congruent model against the GUM. The test of the final model against the GUM serves as a parsimonious encompassing test, that is, a test of whether important information is lost or not in the simplification process. If the final model is not congruent or if it does not parsimoniously encompass the GUM, then the cycle starts all over again by re-specifying the GUM. As such the GETS methodology treats modelling as a process, where the aim is to derive a parsimonious congruent encompassing model while at the same time acknowledging that "the currently best available model" (Hendry and Richard 1990, p. 323) can always be improved.

GETS modelling derives its basis from statistical reduction theory in general and Hendry's reduction theory (1995, chapter 9) in particular,[19] which is a probabilistic framework for the analysis and classification of the simplification errors associated with empirical models. The "theory offers", in Hendry's own words, "an explanation for the origin of all empirical models" (1997, p. 174) in terms of twelve "reduction operations conducted implicitly on the DGP..." (1995, p. 344), and GETS modelling seeks to mimic reduction analysis by evaluating at each reduction whether important information is lost or not. Evaluation of any empirical model can take place against six types of information-sets, namely 1) past data, 2) present data, 3) future data, 4) theory information, 5) measurement information and 6) rival models, and with each of these types we may delineate an associated set of properties that a model should exhibit in order to be considered as a satisfactory, simplified representation of the DGP:[20]

1. Innovation errors. For a model to be a satisfactory representation of the process that generated the data, what remains unexplained should vary unsystematically, that is, the errors should be innovations. In practice this entails checking whether the residuals are uncorrelated and homoscedastic.

2. Weak exogeneity. This criterion entails that conditioning variables are weakly exogenous for the parameters of interest.

3. Constant, invariant parameters of interest. Models without stable parameters are unlikely to be successful forecasting models, so this is a natural criterion if successful forecasting is desirable.

4. Theory consistent, identifiable structures. To ensure that a model has a basis in economic reality it should be founded in economic argument.

5. Data admissibility. In the current context, an example of a volatility model that violates this criterion is one that produces negative volatility forecasts.

6. *Encompassing of rival models.* A model encompasses another if it accounts for its results. Within the three-step cycle of GETS modelling sketched above, a parsimonious encompassing test is undertaken when the final model is tested against the GUM. If no or sufficiently little information is lost then the final model accounts for the results of the GUM.

Models characterised by the first five criteria are said to be congruent, whereas models that also satisfy the sixth are said to be encompassing congruent.

It is important to distinguish between two aspects of the GETS methodology, namely the properties a model (ideally) should exhibit on the one hand, that is, congruent encompassing, and the process of deriving it on the other, that is, general-to-specific search. Contrary to what the name of the GETS methodology may suggest it is actually the former that is of greatest importance. In the words of Hendry, "the credibility of the model is not dependent on its mode of discovery but on how well it survives later evaluation of all of its properties and implications..." (1987, p. 37). However, there is no secret that general-to-specific search for the "currently best available" specification is the preferred approach by the proponents of the GETS methodology. In addition to the fact that it mimics reduction analysis at least four additional important reasons can be listed:[21] The search for the currently best available specification is ordered since any specification obtained in the search is nested within the GUM; in statistical frameworks where adding regressors reduces the residual variance—as for example in the linear model with OLS estimation—the power in hypothesis testing increases; the GETS methodology provides a systematic approach to economic hypothesis testing; and finally compared with unsystematic searches GETS search is resource efficient, see Hendry and Krolzig (2004).

4.3 Empirical models

This section presents the empirical forecast models and proceeds in three steps. Compared with the previous chapter the same sample 8 January 1993 to 26 December 2003 is used for estimation (573 observations). The forecast sample is 2 January 2004 to 25 February 2005 (61 observations). The first subsection contains specifications that condition on both "certain" and "uncertain" information. With certain information we mean information that is either known, for example past values, or which is predictable with a high degree of certainty, for example holidays. With uncertain information we mean information that is not predictable with a high degree of certainty, and typical examples would be contemporaneous values of economic variables. The motivation behind the distinction between certain and uncertain information is that it enables us to gauge the potential forecast precision in the ideal case where the values of the uncertain information are correct. This is of particular interest since the GETS methodology often is championed for its ability to develop models appropriate for scenario analysis (counterfactual analysis, policy analysis, conditional forecasting, etc.), where conditioning on uncertain information plays an important part. The distinction is also of practical interest, since it enables us to investigate whether GETS models with uncertain information improve upon the forecast accuracy of models without uncertain information, since the uncertain information would have to be forecasted in a realistic forecast setting. The second subsection contains specifications with certain information only, whereas the third and final subsection contains the benchmark or ”simple” specifications that serve as a point of comparison. These models are relatively parsimonious and require little development and maintenance effort, thus the label ”simple”, and they have a documented forecasting record. Their motivation is that an important issue is whether GETS derived specifications improve upon the forecast accuracy provided by simple models.

4.3.1 Models with both certain and uncertain information

This subsection presents our models with both certain and uncertain information. Specifically they are

$$\text{GUM EMOV1:} \quad v_t^w = b_0 + b_1 v_{t-1}^w + b_2 v_{t-2}^w + b_3 v_{t-3}^w + b_4 v_{t-4}^w + b_5 \bar{q}_{t-1}^6 + b_6 \Delta q_t + b_7 m_t^w + b_8 o_t^w + b_9 x_t^w + b_{10} u_t^w + b_{11} f_t^a + b_{12} f_t^b + b_{13} ir_t^{emu} + b_{14} sd_t + b_{15} ia_t + b_{16} r_{t-1} + \sum_{l=1}^{8} b_{16+l} h_{lt} + e_t \quad (4.1)$$

$$\text{GETS EMOV1:} \quad v_t^w = b_0 + b_2(v_{t-2}^w + v_{t-3}^w) + b_6 \Delta q_t + b_9(x_t^w + u_t^w) + b_{12} f_t^b + b_{13} ir_t^{emu} + b_{14} sd_t + e_t, \quad (4.2)$$

$$\text{GETS EMOV2:} \quad v_t^w = b_0 + b_2(v_{t-2}^w + v_{t-3}^w) + b_9(\bar{x}^w + \bar{u}^w) + b_{13} \bar{ir}^{emu} + b_{14} sd_t + e_t, \quad (4.3)$$

$$\text{GETS EMOV3:} \quad v_t^w = b_0 + b_2(v_{t-2}^w + v_{t-3}^w) + b_9(\dot{x}^w + \dot{u}^w) + b_{13}\dot{ir}^{emu} + b_{14}sd_t + e_t, \quad (4.4)$$

where $\{e_t\}$ is a sequence of innovation errors. The first specification GUM EMOV1 is a general and unrestricted model with both known and unknown information, whereas the second (GETS EMOV1) is the GETS derived counterpart. Of these two only the second will be used in our out-of-sample study, and it should be noted that it differs slightly from its counterpart (3.14) in the previous chapter. In particular, $\bar{q}_{t-1}^6$ is not retained in the simplification process due to insignificance. The second specification GETS EMOV1 is obtained by setting the first as the general unrestricted specification, and then testing restrictions regarding the parameters with Wald-tests before the final specification is tested against the GUM. It should be noted that we only perform a single specification search where, at each step, we remove the regressor with highest p-value. Hoover and Perez (1999) have pointed out that performing only a single simplification search might result in "path dependence", in the sense that a relevant variable is removed early on in the search whereas irrelevant variables that proxy its role are retained. However, the software PcGets version 1.0 (see Hendry and Krolzig 2001), which automates GETS multiple-path simplification search, produces a specification almost identical to (4.2), the only difference being that v_{t-2}^w is not retained.[22] So path-dependence does not appear to be a major problem in our case. This is consistent with White's (1990) theorem, which implies that the path dependence problem reduces as the size of the sample increases.[23] In the generation of GETS EMOV1 forecasts two steps ahead and onwards we use forecasted values of v_t^w and observed values of the other covariates. In other words, the forecasts of

GETS EMOV1 are generated as if the uncertain conditioning information is known. As such the accuracy of GETS EMOV1 constitutes an indication of its potential for scenario analysis (policy analysis, conditional forecasting, counterfactual analysis, etc.), since its accuracy will reflect its potential of yielding accurate forecasts under the assumption that the uncertain information is correct. The third and fourth specifications serve as a contrast to this hypothetical situation and try to mimic more realistic circumstances by using the parameter estimates of GETS EMOV1, and by using simple forecasting rules for the uncertain information. In GETS EMOV2 the variables Δq_t and f_t^b are set equal to zero, and x_t^w, u_t^w and ir_t^{emu} are set equal to their sample averages $\bar{x}^w$, $\bar{u}^w$ and $\bar{ir}_t^{emu}$ over the period 1 January 1999 - 26 December 2003.[24] In other words, variables that would have to be forecasted in a practical setting are either set to zero or to their recent sample averages. GETS EMOV3 proceeds similarly with a single difference. Instead of averages the medians of x_t^w, u_t^w and ir_t^{emu}, denoted $\dot{x}^w$, $\dot{u}^w$ and $\dot{ir}^{emu}$, are used.[25]

Estimation results and recursive parameter stability analysis of the first two specifications are contained in table 4.1, and in figures 4.1 and 4.2. Both GUM EMOV1 and GETS EMOV1 exhibit innovation errors in the sense that the nulls of no serial correlation, no autoregressive conditional heteroscedasticity and no heteroscedasticity are not rejected at the 10% significance level, and the recursive parameter stability analysis suggests parameters are relatively stable. For both GUM EMOV1 and GETS EMOV1 the Chow forecast and breakpoint tests do not signify at the 1% level, but the 1-step forecast tests on the other hand show some signs of instability.[26] The number of spikes that exceeds the 1% critical value in the break-point tests is 11 and 13, respectively. This suggests the presence of some structural instability since on average we would expect only 5 spikes to exceed the 1% critical value (1% of 473 is just below 5).[27] We do not provide residual diagnostics and coefficient stability analysis of GETS EMOV2 and GETS EMOV3, since these use the parameter estimates of GETS EMOV1.

4.3.2 Models with certain information

This subsection contains our specifications with known or relatively certain information. Specifically they are

GUM EMOV4:
$$v_t^w = b_0 + b_1 v_{t-1}^w + b_2 v_{t-2}^w + b_3 v_{t-3}^w + b_4 v_{t-4}^w + b_{14} sd_t + b_{16} r_{t-1} + \sum_{l=1}^{8} b_{16+l} h_{lt} + e_t, \quad (4.5)$$

GETS EMOV4:
$$v_t^w = b_0 + b_2(v_{t-2}^w + v_{t-3}^w) + b_{14} sd_t + b_{18} h_{2t} + e_t, \quad (4.6)$$

GARCH(1,1)+:
$$r_t = b_0 + b_1 r_{t-1} + e_t, \qquad e_t = \sigma_t z_t,$$
$$\sigma_t^2 = \omega + \alpha e_{t-1}^2 + \beta \sigma_{t-1}^2 + \gamma_1 h_{2t}, \quad (4.7)$$

EGARCH(1,1)+:
$$r_t = b_0 + b_1 r_{t-1} + e_t, \qquad e_t = \sigma_t z_t,$$
$$\log \sigma_t^2 = \omega + \alpha \left|\frac{e_{t-1}}{\sigma_{t-1}}\right| + \beta \log \sigma_{t-1}^2 + \gamma_0 \frac{e_{t-1}}{\sigma_{t-1}} + \gamma_1 h_{2t}, \quad (4.8)$$

where σ_t is the conditional standard deviation of r_t, and $\{z_t\}$ is a sequence of random variables each with mean equal to zero conditional on the information set in question, and each with variance equal to one conditional on the same information set. The first specification GUM EMOV4 is a general formulation nested within GUM EMOV1 but containing only "certain" conditioning information, that is, past and relatively certain contemporaneous information (holiday variables). The second specification GETS EMOV4 is obtained through GETS-analysis of GUM EMOV4. In the third and fourth specifications a constant b_0, lagged return r_{t-1} and h_{2t} are added to "plain" GARCH(1,1) and EGARCH(1,1) specifications. In addition to the fact that the conditional variance σ_t^2 is modelled exponentially, the EGARCH differs from the GARCH by the inclusion of an asymmetry term $\frac{e_{t-1}}{\sigma_{t-1}}$ in the conditional variance specification. A value of γ_0 unequal to zero implies asymmetry and $\gamma_0 < 0$ in particular implies leverage, that is, that returns are negatively correlated with last period's volatility. The higher $|\beta|$ the higher persistence, and a necessary condition for covariance stationarity is $|\beta| < 1$, see Nelson (1991).

The estimation results of the four specifications are contained in tables 4.2 and 4.3, and recursive parameter stability analysis of GUM EMOV4 and GETS EMOV4 in figures 4.3 and 4.4. Both specifications exhibit innovation errors in the sense that the nulls of no serial correlation, no autoregressive conditional heteroscedasticity and no heteroscedasticity are not rejected at conventional significance levels, and the recursive parameter stability analysis for GUM EMOV4 and GETS EMOV4 are similar to those of GUM EMOV1 and GETS EMOV1 above. Both GARCH(1,1)+ and EGARCH(1,1)+ exhibit uncorrelated standardised residuals and squared standardised residuals according to the diagnostic tests, and the impact of lagged return r_{t-1} in the mean equation is negative as commonly found for exchange rates, but not significant. The estimates of $\alpha + \beta$ (0.129 + 0.877 = 1.006) and β (0.983) are very close to 1. This is usually interpreted as an indication of integrated variance, which implies infinite persistence of shocks on the conditional variance, but in this case it is probably due to the structural break around the beginning of 1997, see figure 1.2 and Starica and Mikosch (2004). Finally, the value of γ_0 is insignificantly different from zero which suggests that the symmetry imposed by the GARCH model is not restrictive, a common finding for exchange rate return.

4.3.3 Simple models

Our benchmark or simple models are all ARCH-specifications, and specifically they are

$$\text{Historical:} \qquad r_t = e_t = \sigma_t z_t, \qquad \sigma_t^2 = \omega \tag{4.9}$$

$$\text{RiskMetrics:} \qquad r_t = e_t = \sigma_t z_t, \qquad \sigma_t^2 = 0.06 e_{t-1}^2 + 0.94 \sigma_{t-1}^2 \tag{4.10}$$

$$\text{EWMA:} \qquad r_t = e_t = \sigma_t z_t, \qquad \sigma_t^2 = \alpha e_{t-1}^2 + \beta \sigma_{t-1}^2 \tag{4.11}$$

$$\text{GARCH(1,1):} \qquad r_t = e_t = \sigma_t z_t, \qquad \sigma_t^2 = \omega + \alpha e_{t-1}^2 + \beta \sigma_{t-1}^2 \tag{4.12}$$

$$\text{EGARCH(1,1):} \quad r_t = e_t = \sigma_t z_t,$$
$$\log \sigma_t^2 = \omega + \alpha \left| \frac{e_{t-1}}{\sigma_{t-1}} \right| + \beta \log \sigma_{t-1}^2 + \gamma \frac{e_{t-1}}{\sigma_{t-1}}, \quad (4.13)$$

where $\{z_t\}$ is characterised as above. The first specification labelled Historical is a GARCH(0,0) estimated on the sample 1/1/1999 - 26/12/2003 (261 observations). In other words, it is the ARCH-counterpart of the sample variance since it models volatility as non-varying, and the sample was chosen because volatility appears relatively stable graphically over this period. Failure to beat the historical variance is detrimental to models of the ARCH-class, since this essentially undermines their *raison d'être*. The second specification is an exponentially weighted moving average (EWMA) with parameter values suggested by RiskMetrics (Hull 2000, p. 372).[28] RiskMetrics proposed these values after having compared a range of combinations on various financial time series. The third specification is an EWMA with estimated parameters whereas the fourth specification is a plain GARCH(1,1) which nests the EWMAs within it, since they can be obtained through parameter restrictions. The fifth and final specification is a plain EGARCH(1,1).

Estimates and residual diagnostics of the simple models are contained in tables 4.4 and 4.5. The estimate of the Historical specification yield standardised residuals that are uncorrelated according to the AR_{1-10} test. Although this is not the case for the AR_{1-1} test which is not reported, the failure of the AR_{1-10} test to reject the null nevertheless suggests that the historical variance might be difficult to beat out-of-sample. In the RiskMetrics specification the diagnostic tests suggest the values of α and β are suboptimal, since both the standardised residuals and the squared standardised residuals are serially correlated. Indeed, the diagnostic tests of the EWMA supports this picture since there the nulls of uncorrelated and homoscedastic standardised residuals are not rejected. The α, β estimates and diagnostics of the plain GARCH(1,1) specification are almost identical, and the estimate of ω is almost zero. In other words, the two specifications will produce almost identical forecasts. In the EGARCH(1,1) model residuals are also uncorrelated whereas the estimate of the volatility persistence parameter β is high and almost 1 (it is equal to 0.981). The asymmetry parameter γ is not significant at conventional significance levels, thus suggesting the symmetry of the GARCH(1,1) is not so restrictive. Finally, compared with the estimates of ω, α and β in (4.7) and (4.8) they are virtually identical here. In other words, adding a mean specification and h_{2t} does not seem to affect the estimates of the variance equation noteworthy.

4.4 Out-of-sample forecast evaluation

The out-of-sample evaluation is undertaken on the period 2 January 2004 to 25 February 2005 (61 observations), and the section proceeds in four steps. The first subsection qualifies the widespread view that financial volatility models should be evaluated against estimates based on continuous time theory, and serves as a justification of our evaluation criteria. The second subsection contains socalled Mincer-Zarnowitz (1969) regressions of squared returns on a constant and forecasts, whereas the third contains an out-of-sample forecast accuracy comparison in terms of the mean of the squared forecast errors. The fourth and final subsection sheds additional light on the results by examining some of the 1-step forecast trajectories more closely.

4.4.1 On the evaluation of volatility forecasts

A view that has gained widespread acceptance lately is that discrete time models of financial volatility should be evaluated against estimates based on continuous time theory consistent with the discrete time model, see for example Andersen and Bollerslev (1998), Andersen et al. (1999) and Andersen et al. (2001). Typically these estimators involve high frequency data and a common example of such an estimator is the sum of squared intra-period returns (realised volatility). The motivation for using high frequency estimators derived from continuous time theory is that they are more efficient or less "noisy" than low frequency estimators based upon discrete time theory. Here we qualify the view that discrete models of volatility should be evaluated against estimates drawing on continuous time theory, by arguing that this is particularly inappropriate in the evaluation of discrete time explanatory models of financial volatility.

Consider the discrete time model

$$r_t = f(\mathbf{x}_t, \mathbf{b}) + e_t, \tag{4.14}$$

where $\mathbf{x_t}$ is a vector of variables, $\mathbf{b}$ is a parameter vector and e_t is the error term. If interpreting this as a model of the complex process of human interactions rather than the process itself—this interpretation is a cornerstone of the GETS methodology, then f is the explained part of the variation in r_t and e_t the unexplained. In other words, the $\{e_t\}$ are derived and their characteristics depend on the specfication of f. Needless to say, in such a situation diagnostics on $\mathbf{b}$ and e_t are of prime importance, and encompassing considerations are typically undertaken in terms of the $\{e_t\}$: Given congruence, the more f explains, the smaller the $\{e_t\}$ in absolute value.

If (4.14) is congruent and the $\{e_t\}$ are homoscedastic, then volatility is constant and there is no need for volatility modelling. In the case where the $\{e_t\}$ are heteroscedastic on the other hand, then volatility needs to vary for congruency to (possibly) attain. For example, the heteroscedastic model

$$r_t = f(\mathbf{x}_t, \mathbf{b}) + e_t, \qquad e_t = \sigma_t z_t, \tag{4.15}$$

that is, $\{e_t\}$ is heteroscedastic, is congruent if the $\{\sigma_t\}$ are specified in such a way that $\{z_t\}$ is an innovation and given that the other congruency criteria hold. In other words, the $\{z_t\}$ are derived and their characteristics depend on the specfication of f and σ_t. Again, diagnostics of the parameters in σ_t and of $\{z_t\}$ are of prime importance, and encompassing considerations are typically undertaken in terms of the $\{z_t\}$: Given congruence, the more f and the $\{\sigma_t\}$ explain, the smaller the $\{z_t\}$ are in absolute value.

By contrast, according to Andersen and Bollerslev (1998) and others there is little if any role for $\{z_t\}$ to play, since model comparison is to be conducted in terms of estimates $\{\sigma_t\}$ and $\{\sigma_t^*\}$, where the latter are obtained using continuous time theory. There are well known complications with this view in practical applications, including numerical approximation and sampling issues, see Aït-Sahalia (2006) for a recent overview. Another straightforward objection, however, is that it serves as a restriction. Since the continuous time model (or class of models) that serves as the point of comparison is only a model of the DGP and not the DGP itself, and since the discrete time model is compatible with many different classes of continuous time models, evaluating the discrete time model against only one class of estimates as if they were from *the* true class of models is a restriction. Indeed, given two congruent but non-nested continuous time models of the DGP the natural strategy to compare them would be in terms of (functions of) the absolute size of their estimation and/or forecast errors. Nevertheless, the most important objection against the view that explanatory discrete time models of financial volatility should be evaluated against estimates obtained from continuous time theory is philosophical. The objection is based on the commonplace view (among philosophers) that mathematics is unable to accurately depict time and space. The issue is well known in both the philosophy of time and in the philosophy of mathematics literatures, and an important source of the problem is the principle of extensionality, that is, the axiom that two elements in a set are equal if and only if they are the same. The axiom is a cornerstone of modern formal mathematics, and effectively implies that mathematics is discrete and that time and space continuity only can be approximated by making use of the axiom of infinity (or similar axioms) in one or another way.[1] A well-known example of such a mathematical structure is the set of real numbers, the mathematical structure that is most frequently used in representing continuous time, and the problem in explanatory modelling arises because mental processes like consciousness, reasoning, etc., need time—that is, temporal extension—to acquire the properties we associate with them. As a consequence, economic events have temporal extension and stand at the end of chains of economic events, each with temporal extension: It takes time for one event to bring about another. So as the time increment goes to zero, so does the potential

[1] Entries on the philosophy of time with further reading can be found in virtually any dictionary or companion to philosophy, see for example Honderich (1995) and Kim and Sosa (1995) or even the free internet cyclopedia Wikipedia (http://www.wikipedia.com). A historical introduction to the philosophy of set theory that explicitly deals with time-space issues is Tiles (1989). Bertsimas et al. (2000) have also argued that mathematics can only approximately describe time. However, their view is based on the opposite philosophical standpoint than ours. According to them time *is* discrete and mathematics is continuous.

of explanatory modelling of human events, and hence evaluating discrete explanatory models by means of continuous time theory appears especially inappropriate. For these reasons we employ forecast accuracy measures in which the forecast error is of central importance.

4.4.2 1-step Mincer-Zarnowitz regressions

A simple way of evaluating forecast models is by regressing the variable to be forecasted on a constant and on the forecasts, socalled Minzer-Zarnowitz (1969) regressions, see Andersen and Bollerslev (1998) and Patton (2005) for a discussion on their use in volatility forecast evaluation. In our case this proceeds by estimating the specification

$$r_t^2 = a + b\hat{V}_t + e_t, \tag{4.16}$$

where $\hat{V}_t$ is the 1-step forecast and e_t is the error term. Ideally, a should equal zero and b should equal one—since these constitute conditions for "unbiasedness", and the fit should be high.

Table 4.6 contains the regression output. Patton (2005, footnote on p. 6) has noted that the residuals in Mincer-Zarnowitz regressions typically are serially correlated and that this should be taken into account by using (say) Newey and West (1987) standard errors. In our case the residuals are not serially correlated according to standard tests, but admittedly it might be undetectable due to our relatively small sample. One specification stands out according to the majority of the criteria, namely GETS EMOV1. Its estimate of a is not significantly different from zero, the estimate of b is positive and significantly different from zero, the joint restriction $a = 0, b = 1$ is not rejected at conventional significance levels, and its R^2 is 0.26. This is substantially higher than any of the R^2s cited in Andersen and Bollerslev (1998, pp. 890-891) (the typical R^2 they cite is around 0.03 and the highest is 0.11), and must be very close to—if not exceeding—their population upper bound of R^2:

> "..with conditional Gaussian errors the $R^2_{(m)}$ from a correctly specified GARCH(1,1) model is bounded from above by $\frac{1}{3}$, while with conditional fat-tailed errors the upper bound is even lower. Moreover, with realistic parameter values for $\alpha_{(m)}$ and $\beta_{(m)}$, the population value for the $R^2_{(m)}$ statistic is significantly below this upper bound"—Andersen and Bollerslev (1998, p. 892).

In other words, although the unusually high R^2 of GETS EMOV1 might be due to sample specificity, it nevertheless suggests the poor forecasting performance of r_t^2 by ARCH-models can be improved upon substantially. Moreover, apart from the RiskMetrics specification, Historical beats the other five members of the ARCH-family (EWMA, GARCH(1,1), EGARCH(1,1), GARCH(1,1)+ and EGARCH(1,1)+), and the four models GETS EMOV1-4 perform better than Historical according to R^2. Also, in none of these four specifications is neither a significantly different from zero, nor is the joint restriction $a = 0, b = 1$ rejected. Apart from Historical and RiskMetrics, the restriction $a = 0, b = 1$ is rejected at the 5% level in all the ARCH-specifications, and a is significantly different from zero.

4.4.3 Out-of-sample MSE comparison

Consider a sequence of volatilities $\{V_k\}$ over the forecast periods $k = 1, \ldots, K$ and a corresponding sequence of forecasts $\{\hat{V}_k\}$. Our out-of-sample forecast accuracy measures consist of the mean squared error

(MSE)

$$MSE = \frac{1}{K}\sum_{k=1}^{K}(V_k - \hat{V}_k)^2, \tag{4.17}$$

and Modified Diebold-Mariano (Harvey et al. 1997) tests of the mean squared forecast errors against those of Historical.[29] Error-based measures are "pure" precision measures in the sense that evaluation is based solely on the discrepancy between the forecast and the actual value. One can make a case for the view that precision-based measures are the most appropriate when evaluating the forecast properties of a certain modelling strategy, since this leaves open what the ultimate use of the model is. On the other hand, this is also a weakness since considerations pertaining to the final use of the model do not enter the evaluation.[30]

The values of the MSE forecast statistics are contained in table 4.7. In the forecasting literature models with economic covariates are typically championed as producers of accurate long-term forecasts, but not necessarily of short-term forecasts better than those of "naïve" or simple models without economic covariates. Our results seem to contradict this for the short term. On short horizons up to six weeks ahead GETS EMOV1, the specification with actual values on the economic variables of the right hand side, performs well. According to the MSE it comes 1st on all horizons up to 6 weeks. On longer horizons, however, results are less encouraging. For 12 weeks ahead the GETS EMOV1 comes 8th (out of 10) according to the MSE. One might suggest that this is due to parameter instability, and the results of GETS EMOV2-3 suggest this is indeed the case. Recall that these models are the same as GETS EMOV1 in terms of parameter estimates, but use forecasted values on the right-hand variables that in a practical setting would have to be forecasted. GETS EMOV2-3 come 2nd and 3rd, which suggests the comparatively bad MSE associated with GETS EMOV1 on the 12 week horizon is due to one or several of the uncertain right-hand variables, and therefore instability in one or more of the parameters associated with them.

In a practical forecasting situation the actual values on the right hand side of the GETS EMOV1 specification would have to be forecasted, and GETS EMOV2 and GETS EMOV3 try to mimic such a situation. Both models are relatively consistent and perform comparatively well with the other models. The GETS EMOV2 comes 4th, 4th, 4th, 5th and 2nd according to the MSE, whereas the GETS EMOV3 comes 3rd, 3rd, 3rd, 4rd and 3rd. Although the MSE measures suggest that the GETS models perform relatively well compared with the other models, it should be stressed that so do some of the simple models at times. In particular, the Historical specification comes 2nd, 3rd and 1st on the 3, 6 and 12 week horizons, respectively. The RiskMetrics, GARCH and EGARCH specifications do not do particularly well at short horizons, that is, on horizons in which one would expect them to do well. Not once does any of the five specifications beat Historical 1 to 3 weeks ahead.

In terms of ranking the MSE statistics suggest indeed that the GETS models perform well out-of-sample. But are their MSEs significantly better than that of the simplest comparison model, namely Historical? Table 4.8 contains the output on such a comparison. More precisely, the table contains the p-values of the Modified Diebold-Mariano test against the forecasts of Historical, and they do not suggest that the MSE associated with any of the models—including the GETS models—is significantly lower than the MSE of Historical at any horizon. Indeed, the lowest p-value is as high as 41%. This result is somewhat surprising in light of the previous discussion and in light of the Minzer-Zarnowitz regressions in the previous subsection, and the next subsection aims at explaining this.

4.4.4 Explaining the forecast results

An important part of an out-of-sample study consists of explaining the results, and to this end figure 4.5 provides a large part of the answer. The figure contains the out-of-sample trajectories of squared NOK/EUR log-returns in percent r_t^2, the 1-step forecasts of GETS EMOV1 and the 1-step forecasts of Historical, and the figure provides some interesting insights on the forecast accuracy results. First, the series of r_t^2 seems to be characterised by some occasional large values but little persistence in the sense that large values do not tend to follow each other. Indeed, only at two instances is a large value followed by another, and for a relatively large portion of the sample r_t^2 stays rather low. This explains to some extent the forecast accuracy of Historical. Second, in the 5th and in the 11th weeks of the forecast sample Norges Bank changed its main policy interest rate. This is reflected in the large values of r_t^2 in the 5th and 11th weeks, and explains the forecast accuracy of GETS EMOV1 (it contains a variable for policy interest rate changes) and its unusually high R-squared in the 1-step forecast regressions. Finally, the other explanatory variables included in GETS EMOV1 are probably the reason why it also follows r_t^2 relatively well at other instances when r_t^2 moves substantially. However, an exception to this seems to occur on at least two occasions. In other words, although the forecast success of GETS EMOV1 is due to the explanatory variables, there are also signs that they may have the opposite effect, namely increasing the forecast error. All in all, then, the forecast results suggest the GETS EMOV1 is useful for conditional forecasting but that it does not improve upon the forecasting by simple models when the explanatory variables are unchanging or move little.

4.5 Conclusions

This study has evaluated the out-of-sample forecast accuracy of GETS derived models of weekly NOK/EUR volatility. The GETS specification that uses actual values of uncertain information performs particularly well when it is able to explain big movements in the exchange rate, but not necessarily better than simple models like the constant volatility model when the exchange rate does not move much or when it is unable to explain the movement. All in all, then, our results suggest GETS models are comparatively useful in forecasting, and that they are particularly useful in conditional forecasting. Models of the GARCH(1,1) and EGARCH(1,1) types do not fare particularly well and the reason is that large squared returns do not seem to come in pairs or longer sequences at the weekly frequency. Rather, big movements in the weekly NOK/EUR exchange rate seems more to be a ”one off” phenomenon.

This suggests several lines for further research. First, the generality of our results must be established. Is GETS-modelling of financial volatility useful on higher frequencies than the weekly which typically exhibit more volatility persistence? On other exchange rates and for other financial assets? Second, contrary to McAleer's (2005) assertion, automated GETS-modelling of financial volatility can be readily implemented and should be investigated more fully. Finally, a drawback with our approach is that the conditional mean is restricted to zero, which means that predictability in the direction of exchange rate changes can not be exploited. One interesting line of research is thus to make use of two-step OLS estimators of ARCH models so that on the one hand all the numerical issues and problems associated with GETS modelling of volatility are avoided, and on the other that conditional means also can be modelled.

Table 4.1: GUM and GETS regressions of log of weekly NOK/EUR volatility on both certain and uncertain information

Regressor	(4.1)		(4.2)	
	Est.	*Pval.*	*Est.*	*Pval.*
$const.$	-3.304	0.03	-3.035	0.00
v^w_{t-1}	0.013	0.76		
v^w_{t-2}	0.070	0.07		
v^w_{t-3}	0.093	0.03		
v^w_{t-4}	-0.001	0.99		
$v^w_{t-2} + v^w_{t-3}$			0.079	0.00
q_t	0.063	0.76		
Δq_t	1.024	0.00	1.066	0.00
m^w_t	0.067	0.13		
o^w_t	0.021	0.65		
x^w_t	0.125	0.01	0.119	0.00
u^w_t	0.113	0.01		
f^a_t	-0.256	0.83		
f^b_t	3.775	0.00	3.751	0.00
ir^{emu}_t	4.797	0.02	4.819	0.00
sd_t	1.130	0.00	1.238	0.00
ia_t	-0.127	0.52		
r_{t-1}	-0.025	0.85		
h_{1t}	-1.207	0.16		
h_{2t}	-0.141	0.62		
h_{3t}	0.330	0.64		
h_{4t}	-0.710	0.22		
h_{5t}	0.195	0.71		
h_{6t}	0.653	0.25		
h_{7t}	0.019	0.98		
h_{8t}	-0.036	0.96		
R^2	0.21		0.20	
AR_{1-10}	5.11	0.88	3.07	0.98
$ARCH_{1-10}$	7.00	0.73	8.71	0.56
$Het.$	39.72	0.44	11.46	0.41
$Hetero.$	179.98	0.92	14.60	0.95
JB	124.51	0.00	121.59	0.00
$Obs.$	568		569	

Note: See table 4.2.

Table 4.2: GUM and GETS Regressions of log of weekly NOK/EUR volatility on certain information only

Regressor	(4.5)		(4.6)	
	Est.	*Pval.*	*Est.*	*Pval.*
const.	-2.923	0.00	-2.946	0.00
v^w_{t-1}	0.030	0.48		
v^w_{t-2}	0.077	0.04		
v^w_{t-3}	0.093	0.03		
v^w_{t-4}	-0.020	0.63		
$(v^w_{t-2} + v^w_{t-3})$			0.088	0.00
sd_t	1.449	0.00	1.428	0.00
r_{t-1}	-0.023	0.86		
h_{1t}	-0.821	0.30		
h_{2t}	-0.499	0.07	-0.478	0.08
h_{3t}	0.436	0.59		
h_{4t}	-0.752	0.17		
h_{5t}	0.174	0.71		
h_{6t}	0.346	0.55		
h_{7t}	-0.446	0.55		
h_{8t}	-0.626	0.37		
R^2	0.14		0.13	
AR_{1-10}	9.03	0.53	4.22	0.94
$ARCH_{1-10}$	8.61	0.57	9.25	0.51
Het.	25.70	0.18	6.39	0.17
Hetero.	76.81	0.70	7.43	0.39
JB	112.06	0.00	121.61	0.00
Obs.	568		569	

Note: The estimation sample is 8 January 1993 to 26 December 2003 and computations are in EViews 5.1 with OLS estimation. All specifications use robust standard errors of the White (1980) type, *Pval* stands for *p*-value and corresponds to a two-sided test with zero as null, AR_{1-10} is the χ^2 version of the Lagrange-multiplier test for serially correlated residuals up to lag 10, $ARCH_{1-10}$ is the χ^2 version of the Lagrange-multiplier test for serially correlated squared residuals up to lag 10, *Het.* and *Hetero.* are the χ^2 versions of White's (1980) heteroscedasticity tests without and with cross products, respectively, *JB* is the Jarque and Bera (1980) test for non-normality in the residuals, and *Obs.* is the number of non-missing observations.

Table 4.3: ARCH models of r_t with certain information.

Regressor	(4.7) *Est.*	*Pval.*	(4.8) *Est.*	*Pval.*
$const.$	-0.017	0.42	-0.007	0.77
r_{t-1}	-0.055	0.33	-0.067	0.21
$var.const.$	0.009	0.07	-0.209	0.00
e^2_{t-1}	0.129	0.01	0.285	0.00
σ^2_{t-1}	0.877	0.00	0.983	0.00
$\frac{e_{t-1}}{\sigma_{t-1}}$			0.012	0.86
h_{2t}	-0.043	0.08	-0.171	0.35
$LogL.$	-575.68		-571.98	
AR_{1-10}	9.64	0.47	14.62	0.15
$ARCH_{1-10}$	4.80	0.90	5.75	0.84
JB	664.93	0.00	417.11	0.00
$Obs.$	571		571	

Note: The estimation sample is 8 January 1993 to 26 December 2003 and computations are in EViews 5.1 with robust standard errors of the Bollerslev and Wooldridge (1992) type. $Pval$ stands for p-value and corresponds to a two-sided test with zero as null, $LogL$ stands for log-likelihood, AR_{1-10} is the Ljung and Box (1979) test for serial correlation in the standardised residuals up to lag 10, $ARCH_{1-10}$ is the Ljung and Box (1979) test for serial correlation in the squared standardised residuals up to lag 10, JB is the Jarque and Bera (1980) test for non-normality in the standardised residuals, and $Obs.$ is the number of non-missing observations.

Table 4.4: Historical, RiskMetrics and EWMA models of r_t.

Regressor	(4.9) *Est.*	*Pval.*	(4.10) *Est.*	*Pval.*	(4.11) *Est.*	*Pval.*
var.const.	0.681	0.00				
e^2_{t-1}			0.060		0.159	0.01
σ^2_{t-1}			0.940		0.877	0.00
LogL.	-370.34		-861.22		-583.63	
AR_{1-10}	8.94	0.54	31.49	0.00	12.88	0.23
$ARCH_{1-10}$	9.49	0.49	46.38	0.00	5.32	0.72
JB	87.85	0.00	145K	0.00	565.09	0.00
Obs.	261		572		572	

Note: The estimation sample is 1 January 1999 to 26 December 2003 for (4.9), and 8 January 1993 to 26 December 2003 for (4.10) and (4.11). Computations are in G@RCH 4.0 and EViews 5.1 with robust standard errors of the Bollerslev and Wooldridge (1992) type, and "K" is short for "kilo(s)", that is, 145K = 145 000. Otherwise see table 4.3.

Table 4.5: Plain ARCH models of r_t.

Regressor	(4.12) *Est.*	*Pval.*	(4.13) *Est.*	*Pval.*
var.const.	0.006	0.14	-0.223	0.00
e^2_{t-1}	0.146	0.01	0.292	0.00
σ^2_{t-1}	0.867	0.00	0.981	0.00
$\frac{e_{t-1}}{\sigma_{t-1}}$			0.017	0.79
LogL.	-580.96		-576.75	
AR_{1-10}	11.43	0.32	12.27	0.27
$ARCH_{1-10}$	4.97	0.89	6.45	0.78
JB	635.55	0.00	377.88	0.00
Obs.	572		572	

Note: See table 4.3.

Table 4.6: Mincer-Zarnowitz regressions of r_t^2 on a constant and 1-step out-of-sample forecasts ($K = 61$)

	a	b	R^2	$Pval.$
GETS EMOV1	-0.17	1.36	0.26	0.79
	[0.63]	[0.02]		
GETS EMOV2	-0.32	1.65	0.03	0.76
	[0.78]	[0.36]		
GETS EMOV3	-0.32	1.54	0.03	0.90
	[0.78]	[0.36]		
GETS EMOV4	-0.20	1.20	0.03	0.91
	[0.82]	[0.31]		
GARCH(1,1)+	0.92	-0.08	0.00	0.00
	[0.01]	[0.80]		
EGARCH(1,1)+	0.93	-0.09	0.00	0.02
	[0.02]	[0.83]		
Historical	-	1.24	0.00	0.39
		[0.00]		
RiskMetrics	0.93	-0.10	0.00	0.13
	[0.08]	[0.86]		
EWMA	0.95	-0.10	0.00	0.00
	[0.01]	[0.70]		
GARCH(1,1)	0.96	-0.12	0.00	0.00
	[0.01]	[0.69]		
EGARCH(1,1)	0.98	-0.14	0.00	0.01
	[0.01]	[0.71]		

Note: Numbers in square brackets denote the p-values of a two-sided coefficient hypothesis test with zero as the null hypothesis, and the last column denotes the p-value of a $\chi^2(2)$ Wald test of the joint restriction $a = 0, b = 1$. Otherwise see table 4.2.

Table 4.7: MSE forecast statistics

	1-week	2-week	3-week	6-week	12-week
GETS EMOV1	**1.609**	**1.635**	**1.070**	**1.067**	1.253
GETS EMOV2	2.074	2.100	1.209	1.196	1.023
GETS EMOV3	2.061	2.086	1.208	1.192	1.023
GETS EMOV4	2.040	2.058	1.218	1.190	1.027
GARCH(1,1)+	2.307	2.226	1.248	1.236	1.200
EGARCH(1,1)+	2.287	4.817	2.806	1.348	1.202
Historical	2.132	2.165	1.198	1.192	**1.022**
RiskMetrics	2.178	2.658	1.370	1.470	1.171
EWMA	2.435	4.758	2.541	2.054	1.314
GARCH(1,1)	2.331	4.247	2.179	1.847	1.277
EGARCH(1,1)	2.288	2.186	1.475	1.434	1.335

Note: Bold value indicates minimum in its column.

Table 4.8: P-values of the Modified Diebold-Mariano forecast test against Historical using squared forecast errors

	1-week	2-week	3-week	6-week	12-week
GETS EMOV1	0.41	0.41	0.49	0.49	0.51
GETS EMOV2	0.50	0.50	0.50	0.50	0.50
GETS EMOV3	0.50	0.50	0.50	0.50	0.50
GETS EMOV4	0.49	0.49	0.50	0.50	0.50
GARCH(1,1)+	0.51	0.50	0.50	0.50	0.51
EGARCH(1,1)+	0.51	0.99	0.82	0.50	0.51
Historical	—	—	—	—	—
RiskMetrics	0.50	0.57	0.51	0.52	0.50
EWMA	0.53	1.00	0.85	0.67	0.51
GARCH(1,1)	0.51	1.00	0.71	0.60	0.51
EGARCH(1,1)	0.51	0.50	0.52	0.51	0.51

Note: The p-values are computed using a Student's t distribution with 1 degree of freedom (DF) in the 1-week test, 2 DFs in the 2-week test, and so on.

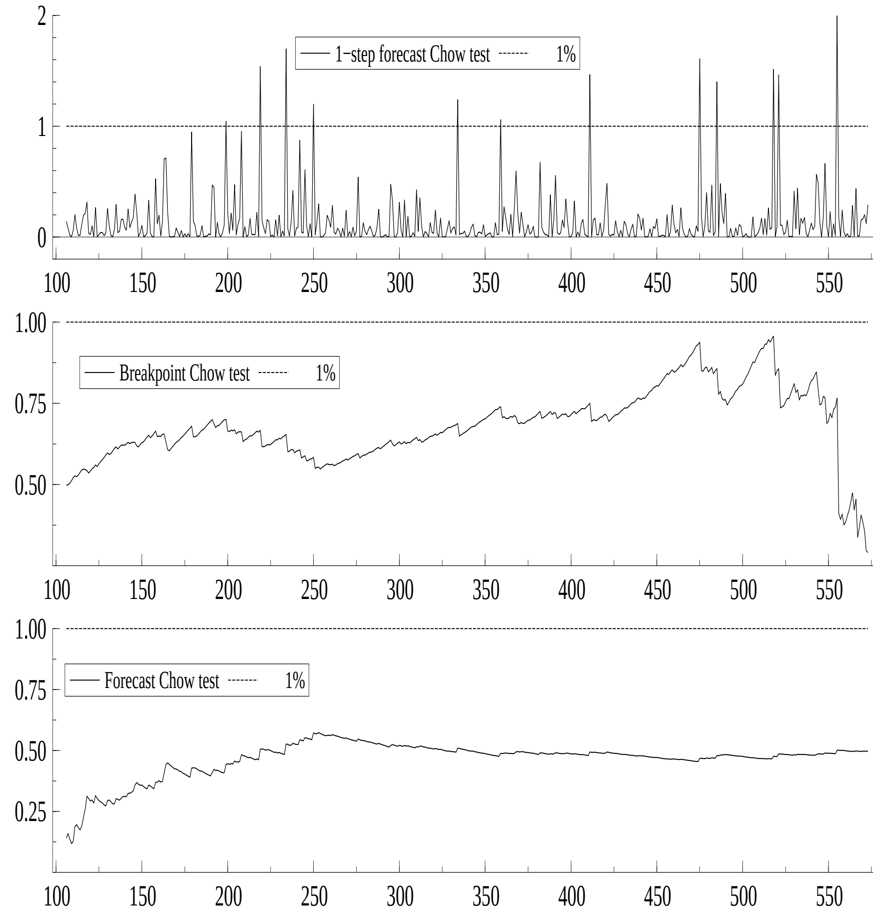

Figure 4.1: Recursive analysis of GUM EMOV1. Computations are in PcGive 10.4 with OLS and initialisation at observation number 100.

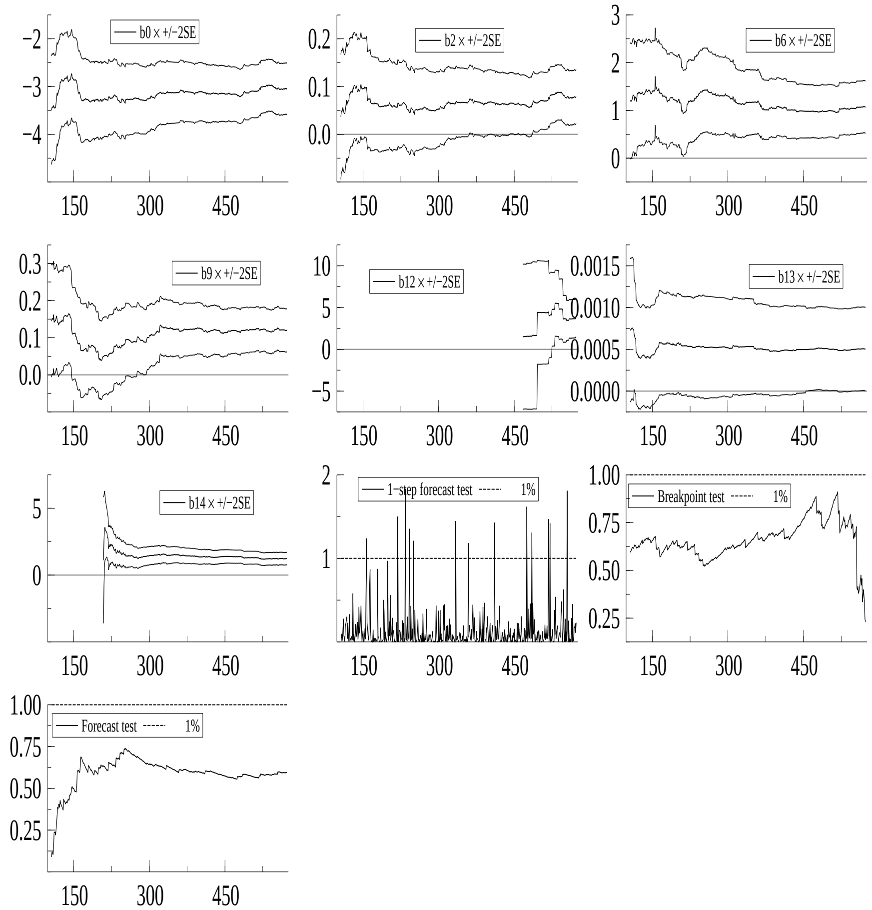

Figure 4.2: Recursive analysis of GETS EMOV1. Computations are in PcGive 10.4 with OLS and initialisation at observation number 100.

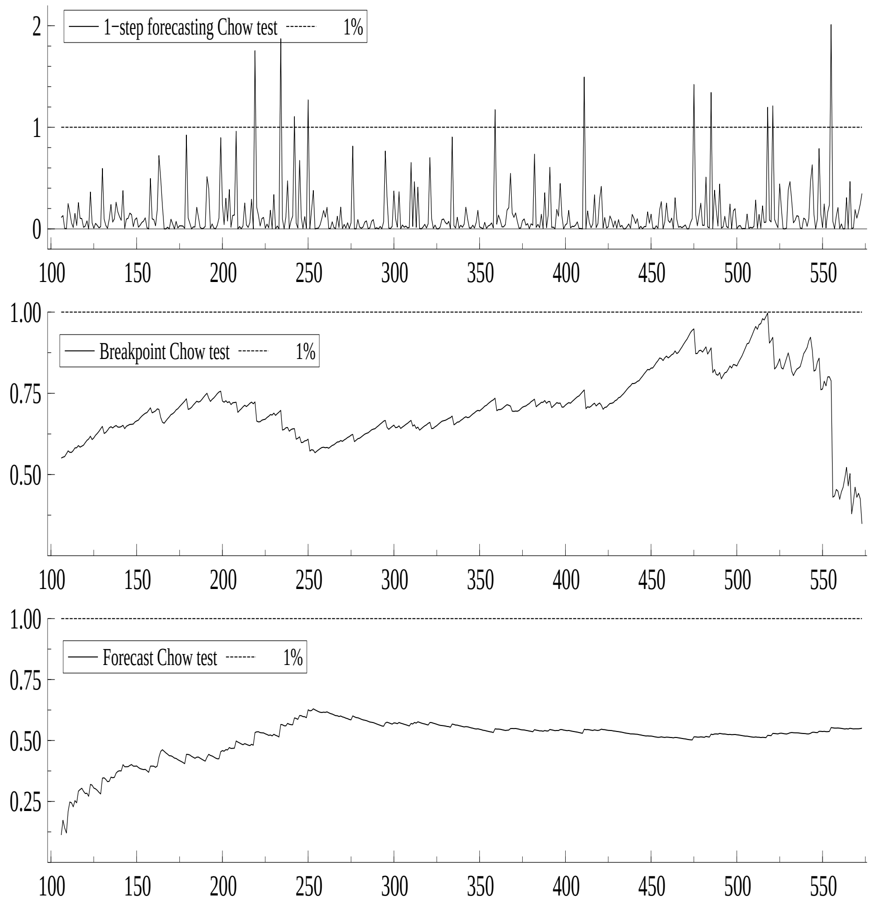

Figure 4.3: Recursive analysis of GUM EMOV4. Computations are in PcGive 10.4 with OLS and initialisation at observation number 100.

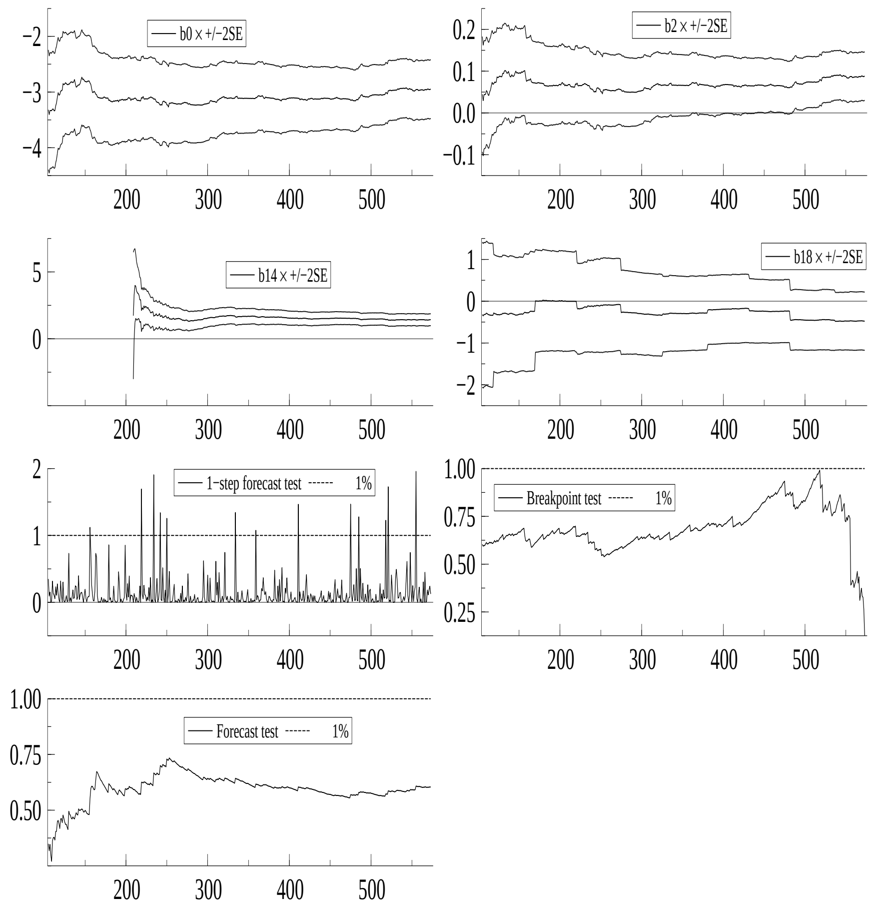

Figure 4.4: Recursive analysis of GETS EMOV4. Computations are in PcGive 10.4 with OLS and initialisation at observation number 100.

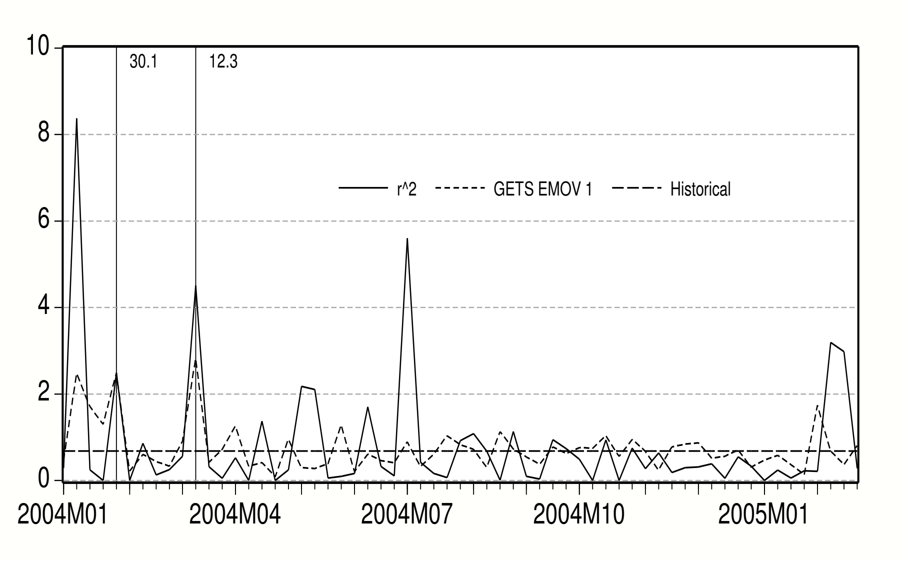

Figure 4.5: Out-of-sample trajectories of r_t^2, GETS EMOV1 and Historical. Vertical lines indicate weeks in which Norges Bank changed their main policy interest rate.

Chapter 5

Exchange rate variability, market activity and heterogeneity

This chapter draws on joint research with Dagfinn Rime.

5.1 Introduction

Currency markets are heterogeneous in many ways. Although trading in principle can take place around the clock, in practice it is concentrated around the opening hours of the major financial centres of Tokyo, London and New York. So as time go by the themes and topics considered relevant change as trading shifts from centre to centre. In each centre's opening hours it is the interbank market which dominates both in terms of volume and influence, but only a few actors have direct access to this market. This gives rise to various "trading venues" that interact with each other to different extent. For example, both the level and the evolution of the Euro per US dollar exchange rate on a specific internet trading platform may vary notably from that of the interbank market. Furthermore, participants tend to restrict their attention to only a handful of currencies. In a world with more than hundred monies but only a few very liquid ones, essentially the US dollar, the Euro and the Yen, most monies are traded against only these liquid currencies. Finally, different types of investors act differently, possess different kinds of information and are differently positioned to impact upon the price of currencies.

The study in this chapter sheds light on the role played by heterogeneity in the relation between market activity and exchange rate variability, through an investigation of the relationship between the weekly variability of the Norwegian Krone (NOK) against the Euro (EUR) and a measure of NOK/EUR spot trading volume in Norway. Compared with chapter 3 the study in this chapter uses a dataset that spans a shorter period because of the spot trading volume data. The spot trading volume data are used to study the role played by heterogeneity in two ways. First, does Norwegian volume matter? On a global scale NOK/EUR trading is one of the smaller currency pairs in terms of volume, so if banks and actors in Norway have an impact it is probably either due to their share of total trading volume being sizeable or due to their privileged proximity to Norwegian demanders and supplier of NOK. The second type of heterogeneity we investigate is whether groups of similarly sized banks have a different impact than other groups of similar size. For example, since bigger banks account for a bigger share of trading volume one might expect that their impact on variability is greater than that of the group of small banks.

The rest of the chapter consists of three parts. In the next we describe the data used, establish notation and motivate our somewhat unusual data frequency, which are due to the spot volume data. The subsequent section contains the empirical results, whereas the final section concludes.

5.2 Data and notation

In order to make efficient use of the data in order to construct a measure of spot NOK/EUR volume we need to make use of an unusual frequency. To be more precise, we make a distinction between the first part of the week on the one hand and the second part of the week on the other. In our estimations we will only make use of spot data from the second part of the week. The second part of the week comprises the last two trading days of the week, typically Thursday and Friday, and the first part of the week comprises the preceding trading days, typically Monday, Tuesday and Wednesday. Compared with the previous chapters this calls for new data transformations and a slightl change in notation. In addition to providing details of the currency transaction volume data, the purpose of this section is thus to explain the needed data-transformations (further details are provided in the data-appendices), and to introduce the associated notation. The section contains three subsections. The first subsection contains the re-definitions of variability. Then, in the second subsection we detail the currency transaction volume data and explain how we use them and the quote frequency data to construct measures of market activity at a bi-weekly frequency. Finally, the third subsection explains the transformations associated with the other variables that we include in our empirical investigations.

5.2.1 Period and range variability re-defined

As in chapter 3 we make use of both a period measure and a range measure of exchange rate variability. Let $r_{t_2}^{se}$ denote the period log-return of the NOK/EUR exchange rate from 07:00 GMT on Thursday to 21:50 GMT on Friday. The superscript "se" is intended to evoke the association "start-end". The time index t_2 stands for the period that comprise the last two trading days in week t, and due to holidays Thursday and Friday are not always the last two trading days of the week. When they are not then returns are adjusted accordingly, see data appendix for further details. Similarly, the time index t_1 is used to denote the period that comprise the trading days that precedes t_2 in week t. The trading days of week t that precedes the last two are typically Monday, Tuesday and Wednesday. When they are not they are adjusted accordingly too.

Period variability in the second part of week t is then defined as $(r_{t_2}^{se})^2$ and denoted $V_{t_2}^{se}$, whereas range variability in t_2 is defined as $[\log(\max S_{t_2}) - \log(\min S_{t_2})]^2$ and denoted $V_{t_2}^{hl}$. Their corresponding log-transformations are denoted in small letters, that is, $v_{t_2}^{se} = \log V_{t_2}^{se}$ and $v_{t_2}^{hl} = \log V_{t_2}^{hl}$. The main characteristics are contained in tables 5.1 and 5.2, and in figures 5.2 - 5.5. For comparison purposes with the figure of weekly returns in chapter 1 the evolution of $r_{t_2}^{se}$ is contained in figure 5.1. There are at least five characteristics worth noting. First, at times period and range variabilities differ notably. For instance, the average of $V_{t_2}^{se}$ over the whole period is 0.27, whereas the average of $V_{t_2}^{hl}$ over the same period is more than double. Moreover, the averages of $V_{t_2}^{se}$ and $V_{t_2}^{hl}$ are higher in the full inflation targeting period than in the partial inflation targeting regime, whereas the opposite is the case for their log-transformations. Second, the log-transformation, which makes pairs of large observations (in absolute value) less influential, also matters for the correlation between period and range volatilities. For instance, the sample correlation between $V_{t_2}^{se}$ and $V_{t_2}^{hl}$ is 0.90 over the whole sample, whereas the sample correlation between $v_{t_2}^{se}$ and $v_{t_2}^{hl}$ is only 0.58 over the same sample. The drop in correlation is similar when the two subsamples are compared. Third, the two definitions are less correlated than one might have expected, particularly between $v_{t_2}^{se}$ and $v_{t_2}^{hl}$, with a minimum of 0.47 attained in the partial inflation targeting period. Fourth, although figures 5.2 and 5.3 suggest that there are more large values of variability in the second policy period, a general increase or shift upward in variability around 29 March 2001 is absent—or at least seemingly so.

5.2.2 Measuring market activity

In order to shed light on the role played by heterogeneity we use two types of market activity data, quote (NOK/EUR) frequency in the international interbank market—using the same underlying rawdata as in the previous chapters—and a measure of spot NOK/EUR trading volume by banks within Norway's regulatory borders. The only change in notation with respect to the quote data is the subscript. For example, the number of quotes in t_2 is denoted Q_{t_2} and its log-counterpart q_{t_2}. The volume data are collected every week by Norges Bank and goes back to the beginning of the 1990s. However, due to substantial changes in the underlying data-collection methodology and definitions we opt to only use the part after 1999, which corresponds to 313 observations at the weekly frequency from 15 January 1999 to 7 January 2005. For more details regarding these data the reader is referred to the data appendix.[31] The spot volume variables in the second part of week t are denoted Z_{t_2} with a superscript, where the superscript refers to the category of the data it stands for. More precisely, we denote our measure of total spot NOK/EUR volume for $Z_{t_2}^{tot}$, the measure of spot NOK/EUR volume by big banks is denoted $Z_{t_2}^{big}$, the measure of spot NOK/EUR volume by medium sized banks is denoted $Z_{t_2}^{med}$, and the measure of spot NOK/EUR volume by small banks is denoted $Z_{t_2}^{sma}$. By definition we have that $Z_{t_2}^{tot} = Z_{t_2}^{big} + Z_{t_2}^{med} + Z_{t_2}^{sma}$, and their log-counterparts are denoted in small letters, that is, $z_{t_2}^{tot}$, $z_{t_2}^{big}$, $z_{t_2}^{med}$ and $z_{t_2}^{sma}$. The three categories of banks are "naturally" formed in the sense that the volume of "large" banks is substantially higher than that of the other banks, and in the sense that the volume of "small" banks is substantially lower than that of the others. For confidentiality reasons we cannot disclose the identities of the banks that make up which category nor the volume associated with each bank, and for further details of the data the reader is referred to the data appendix. Descriptive statistics of the Z_{t_2} variables and their log-counterparts are contained in table 5.3. On average total spot NOK/EUR transaction volume in t_2 amounts to almost 323 million NOK, and about 234 million NOK of this, more that 2/3 of the total amount, is due to the group of large banks. The group of small banks account for less than 5%. These shares are relatively stable over the sample and mean the group of large banks account for a substantive part of volume.

In order to distinguish between the different effects variation in market activity may have on variability we use the same strategy as in previous chapters. Let the symbolism $t_2 - 1$ stand for the second part of week $t - 1$, $t_2 - 2$ for the second part of week $t - 2$, and so on. If z_{t_2} denotes the log of a measure of volume in the second part of week t, then a straightforward decomposition is to define short-term variation as $\Delta z_{t_2} = z_{t_2} - z_{t_2-1}$ and long-term variation as z_{t_2-1}, since by definition

$z_{t_2} = \Delta z_{t_2} + z_{t_2-1}$. Again, as in chapter 3, the short-term component Δz_{t_2} has a straightforward and intuitive economic interpretation, namely the relative increase or decrease in volume compared with the previous period. Similarly, if z_{t_2-1} is sufficiently serially correlated with previous lags, that is, with z_{t_2-2}, with z_{t_2-3} and so on, then its economic interpretation is the "general" or "long-term" level of market activity. Just as in chapter 3 the drawback of using z_{t_2-1} as a measure of long-term variation in market activity is that it might be a noisy measure. One solution is therefore to replace z_{t_2-1} with a smoothed expression, and we employ the same strategy here as previously. The average of log of total volume $z_{t_2}^{tot}$ using two past values is equal to $(z_{t_2-1}^{tot} + z_{t_2-2}^{tot})/2$ and is denoted $\bar{z}_{t_2-1}^{tot/2}$, the average using three values is equal to $(z_{t_2-1}^{tot} + z_{t_2-2}^{tot} + z_{t_2-3}^{tot})/3$ and is denoted $\bar{z}_{t_2-1}^{tot/3}$, and so on. Similarly, the average of log quote frequency using two past values is equal to $(q_{t_2-1} + q_{t_2-2})/2$ and is denoted $\bar{q}_{t_2-1}^{2}$, the average of log quote frequency using three past values is equal to $(q_{t_2-1} + q_{t_2-2} + q_{t_2-3})/3$ and is denoted $\bar{q}_{t_2-1}^{3}$, and so on. Table 5.4 contains selected sample correlations of the various measures of long and short-term variation in market activity. The $\hat{q}_{t_2}$ and $\hat{z}_{t_2}$ variables denote measures of long-term variation obtained through the two-step ARMA method described in subsection 3.2.2. Generally the sample correlations between the measures is relatively strong (generally above 0.8), so one might ask whether using one instead of another actually matters. Our exploratory analyses suggest it does: Which measure that is included matters for the significance results.

5.2.3 Other determinants of exchange rate variability

We also include other variables in our empirical investigations. To account for the possibility of skewness and asymmetries in r_{t_2} we use the lagged return r_{t_1}—that is, the log-return of NOK/EUR in the first part of week t—for the latter, and an impulse dummy ia_{t_2} equal to 1 when returns are positive and 0 otherwise for the former. Also here the Norwegian interest-rate variables reflect the fact that Norway changed inflation policy on 29 March 2001, when the Ministry of Finance instructed Norges Bank to replace exchange rate stabilisation with inflation targeting as its main policy objective. Contrary to previous chapters, however, here we use market interest rate variables instead of the policy interest rate. If $\Delta ir_{t_2}^{no}$ denotes the change in the 3-month Norwegian market interest rate in the second part of week t, then $ir_{t_2}^{no,b}$ is defined as $(\Delta ir_{t_2}^{no})^2$ before the regime change took place and zero thereafter, whereas $ir_{t_2}^{no,c}$ is defined as $(\Delta ir_{t_2}^{no})^2$ after the regime change took place and zero before. To further distinguish between the impact of Norwegian interest rates between changes in the policy interest rate by Norges Bank and the impact when Norges Bank changes the policy interest rate, we further decompose $ir_{t_2}^{no,b}$ and $ir_{t_2}^{no,b}$. More precisely, we will add Δ as a superscript when Norges Bank changes its main policy interest rate, and we will add a 0 when it does not. Specifically, in the partial inflation targeting period, $ir_{t_2}^{no,b0}$ is equal to $ir_{t_2}^{no,b}$ when Norges Bank does not change the policy interest rate and zero otherwise, and $ir_{t_2}^{no,b\Delta}$ is equal to $ir_{t_2}^{no,b}$ when it does and zero otherwise. Similarly the corresponding variables for the full inflation period are denoted as $ir_{t_2}^{no,c0}$ and $ir_{t_2}^{no,c\Delta}$. By construction it follows that $ir_{t_2}^{no,b0} + ir_{t_2}^{no,b\Delta} = ir_{t_2}^{no,b}$, and that $ir_{t_2}^{no,c0} + ir_{t_2}^{no,c\Delta} = ir_{t_2}^{no,c}$. As a measure of changes in short-term EMU interest rates we use a 3-month market rate. Specifically, if $(\Delta ir_{t_2}^{emu})^2$ denote the square of the change in the market interest rate in the second part of week t, then we will use $ir_{t_2}^{emu}$ as a shorthand for this expression. As a measure of general currency market turbulence we use EUR/USD-variability. If Δm_{t_2} denotes the log-return of EUR/USD in the second part of week t, then $M^{se}_{t_2}$ stands for variability and $m^{se}_{t_2}$ its log-counterpart. The petroleum sector plays a major role in the Norwegian economy, so it makes sense to also include a measure of oilprice variability. If the log-return of oilprice over the second part of week t is denoted Δo_{t_2}, then its variability is $O^{se}_{t_2}$ and its log-counterpart $o^{se}_{t_2}$. We proceed similarly for Norwegian and US stock market variables. If Δx_{t_2} denotes the log-return of the main index of the Oslo stock exchange in the second part of week t, then the associated variables are denoted $X^{se}_{t_2}$ and $x^{se}_{t_2}$. In the US case volatilities are denoted as $U^{se}_{t_2}$ and $u^{se}_{t_2}$. We also include two impulse dummies to deal with two extreme (negative) observations in the regressions of $v^{se}_{t_2}$. These two observations are due to the log-transformation being applied on period variability $(r^{se}_{t_2})^2$ when return $r^{se}_{t_2}$ is so close to zero that they appear as clearly smaller values than the others, with the consequence that without the dummies no regressor is significant. These two impulse dummies are denoted $id^2_{t_2}$ and $id^3_{t_2}$, respectively, and an economic interpretation is that the quoting discreteness at these dates is unusually fine.

5.3 Empirical results

This section proceeds in four steps. In the first subsection we provide a simple comparison of the volume data with the quote data in measuring market activity. The next two subsections shed light on the role played by Norwegian market activity, whereas the fourth and final subsection addresses the question of whether Norwegian banksize matters.

All models in this section are nested within the general specification

$$v_{t_2} = b_0 + b_1\text{persistence} + b_{10}\Delta z_{t_2}^{tot} + b_{11}\bar{z}_{t_2}^{tot} + \\ b_{12}\Delta z_{t_2}^{big} + b_{13}\bar{z}_{t_2}^{big} + b_{14}\Delta z_{t_2}^{med} + b_{15}\bar{z}_{t_2}^{med} + b_{16}\Delta z_{t_2}^{sma} + \\ b_{17}\bar{z}_{t_2}^{sma} + b_{18}\Delta q_{t_2} + b_{19}\bar{q}_{t_2} + \text{the rest} + e_{t_2}. \quad (5.1)$$

The left side variable v_{t_2} stands for the log of variability in question, that is, either $v_{t_2}^{se}$ or $v_{t_2}^{hl}$, "persistence" (made explicit below) stands for the associated persistence in variability, the $\bar{z}_{t_2}$ variables stand for volume based measures of long-term market activity, $\bar{q}_{t_2}$ stands for a quote based measure of long-term market activity, "the rest" (made explicit below) stands for the other variables that are included in our regressions, and e_{t_2} denotes the error term. Log-period and log-range variability require different types of lag structures in order to account for persistence. In their specific parsimonious form they are defined as

$$\text{persistence}^{se} = v_{t_1}^{hl} \quad (5.2)$$

$$\text{persistence}^{hl} = 3v_{t_1}^{hl} + 9v_{t_2-1}^{hl} + v_{t_1-1}^{hl}, \quad (5.3)$$

respectively.[32] The presence of only $v_{t_1}^{hl}$ on the right hand side of (5.2) means that past values of log-range variability is a better predictor of log-period variability $v_{t_2}^{se}$ than past values of log-period variability. Since our main focus will be on the quote and volume variables we will at times for convenience reasons refer to the other variables as "the rest".

Specifically this term is defined as

$$\text{the rest} = b_{20}m^{se}_{t_2} + b_{21}o^{se}_{t_2} + b_{22}x^{se}_{t_2} + b_{23}u^{se}_{t_2} + b_{24}ir^{no,b0}_{t_2} + b_{25}ir^{no,b\Delta}_{t_2} + b_{26}ir^{no,c0}_{t_2} + b_{27}ir^{no,c\Delta}_{t_2} + b_{28}ir^{emu}_{t_2} + b_{29}ia_{t_2} + b_{30}r^{se}_{t_1} + b_{31}id^2_{t_2} + b_{32}id^3_{t_2}, \quad (5.4)$$

where the impulse dummies $id^2_{t_2}$ and $id^3_{t_2}$ are only included in the log-period variability regressions.

5.3.1 Quote data vs. volume data

The commonplace view is that volume is a better measure of market activity than quote frequency. However, compared with the quote data our volume data are limited in scope, since they only comprise volume within Norwegian regulatory borders. The purpose of this subsection is to compare how well the two types of data explain variability, and to this end we run separate regressions that only contain each set of market activity variables. The motivation for this is that the data are overlapping in the sense that the quote data also contain the quotes of banks within Norwegian regulatory borders. The motivation for not including other variables in the regression is the same as in subsection 3.3.1, namely to shed light on the hypothesis that variation in market activity is a major cause of variability persistence. Table 5.5 contains estimates of the period variability regressions

$$v^{se}_{t_2} = b_0 + b_{10}\Delta z^{tot}_{t_2} + b_{11}\bar{z}^{tot/2}_{t_2-1} + b_{31}id^2_{t_2} + b_{32}id^3_{t_2} + e_{t_2} \quad (5.5)$$

$$v^{se}_{t_2} = b_0 + b_{18}\Delta q_{t_2} + b_{19}\bar{q}^{15}_{t_2-1} + b_{31}id^2_{t_2} + b_{32}id^3_{t_2} + e_{t_2}, \quad (5.6)$$

whereas table 5.6 contains estimates of the range variability regressions

$$v_{t_2}^{hl} = b_0 + b_{10}\Delta z_{t_2}^{tot} + b_{11}\bar{z}_{t_2-1}^{tot/2} + e_{t_2} \tag{5.7}$$

$$v_{t_2}^{hl} = b_0 + b_{18}\Delta q_{t_2} + b_{19}\bar{q}_{t_2-1}^{15} + e_{t_2}. \tag{5.8}$$

The variables $\bar{z}_{t_2-1}^{tot/2}$ and $\bar{q}_{t_2-1}^{15}$ are chosen as measures of long term market activity on the basis of R^2, and both tables suggest the quote data fare better as measures of market activity in terms of R^2. Also, contrary to the comparable results in table 3.2 of equation (3.8) in chapter 3, the estimates of (5.5) and (5.7) do not exhibit serially correlated errors. An interpretation of this is that the market activity variables adequately account for variability persistence. However, it should be noted that a regression of $v_{t_2}^{se}$ on only a constant (not reported) does not produce serially correlated residuals neither. In other, words, there is little variability persistence in $\{v_{t_2}^{se}\}$ to explain. Moreover, the regressions (5.7) and (5.8) of $v_{t_2}^{hl}$ do produce serially correlated residuals. So there are signs also here that the market activity variables alone are unable to adequately account for time-varying variability.

5.3.2 Persistence vs. market activity

Just as in subsection 3.3.2, here we seek to shed further light on the relation between persistence and market activity. In particular, our aim is to shed light on the hypothesis that persistent financial return variability is explained by market activity. Again we first run regressions of variability on the persistence term, and then we add the market activity variables to see to what extent the persistence estimates and significance results are affected.

Table 5.7 contains estimates of the period and range variability regressions

$$v^{se}_{t_2} = b_0 + b_1 \text{persistence} + b_{31} id^2_{t_2} + b_{32} id^3_{t_2} + e_{t_2} \tag{5.9}$$

$$v^{hl}_{t_2} = b_0 + b_1 \text{persistence} + e_{t_2}. \tag{5.10}$$

In the period variability regression (5.9) the persistence term contains a single variable, namely range variability in the first part of the week $v^{hl}_{t_1}$. In the range variability regression (5.10) the persistence term is much richer, and in contrast to the previous subsection the residuals are no longer serially correlated.

Table 5.8 contains estimates of the period variability regressions

$$v^{se}_{t_2} = b_0 + b_1 \text{persistence} + b_{10} \Delta z^{tot}_{t_2} + b_{11} \bar{z}^{tot/2}_{t_2-1} + b_{31} id^2_{t_2} + b_{32} id^3_{t_2} + e_{t_2} \tag{5.11}$$

$$v^{se}_{t_2} = b_0 + b_1 \text{persistence} + b_{18} \Delta q_{t_2} + b_{19} \bar{q}^{15}_{t_2-1} + b_{31} id^2_{t_2} + b_{32} id^3_{t_2} + e_{t_2}, \tag{5.12}$$

whereas table 5.9 contains estimates of the range variability regressions

$$v^{hl}_{t_2} = b_0 + b_1 \text{persistence} + b_{10} \Delta z^{tot}_{t_2} + b_{11} \bar{z}^{tot/2}_{t_2-1} + e_{t_2} \tag{5.13}$$

$$v_{t_2}^{hl} = b_0 + b_1\text{persistence} + b_{18}\Delta q_{t_2} + b_{19}\bar{q}_{t_2-1}^{15} + e_{t_2}. \tag{5.14}$$

The results of the period variability regressions in table 5.8 are mixed. Adding the quote variables increases the persistence estimate from 0.100 to 0.111, whereas adding the volume variables decreases the persistence estimate from 0.100 to 0.087. However, a Wald coefficient restriction test of $b_1 = 0.100$ in (5.11) is not rejected at the 10% level (the p-value is above 80%). The results of the range variability regressions in table 5.9 are less mixed. The coefficient estimate of persistence falls slightly in both regressions, but the associated p-value remains below 0.01%. Moreover, again Wald coefficient restriction tests do not reject the restriction of $b_1 = 0.027$ at the 10% level (the p-values are above 50% in both cases), so the fall is not substantial in neither case. All in all, then, although three of the four persistence estimates fall when the market activity variables are added, the results do not provide statistical support of the hypothesis that market activity accounts for the persistence in variability.

5.3.3 Global vs. Norwegian market activity

In contrast with the previous two subsections here we control for the impact of other variables than those of market activity and persistence. Table 5.10 contains estimates of specifications obtained through simplification of

$$v_{t_2}^{se} = b_0 + b_1\text{persistence} + b_{10}\Delta z_{t_2}^{tot} + b_{11}\bar{z}_{t_2-1}^{tot/2} + \text{the rest} + e_{t_2} \tag{5.15}$$

$$v_{t_2}^{se} = b_0 + b_1\text{persistence} + b_{10}\Delta z_{t_2}^{tot} + b_{11}\bar{z}_{t_2-1}^{tot/2} + b_{18}\Delta q_{t_2} + b_{19}\bar{q}_{t_2-1}^{15} + \text{the rest} + e_{t_2} \quad (5.16)$$

with the constant and the market activity variables fixed, that is, they are not removed if insignificant at 10%, whereas table 5.11 contains estimates of specifications obtained through simplification of

$$v_{t_2}^{hl} = b_0 + b_1\text{persistence} + b_{10}\Delta z_{t_2}^{tot} + b_{11}\bar{z}_{t_2-1}^{tot/2} + \text{the rest} + e_{t_2} \quad (5.17)$$

$$v_{t_2}^{hl} = b_0 + b_1\text{persistence} + b_{10}\Delta z_{t_2}^{tot} + b_{11}\bar{z}_{t_2-1}^{tot/2} + b_{18}\Delta q_{t_2} + b_{19}\bar{q}_{t_2-1}^{15} + \text{the rest} + e_{t_2} \quad (5.18)$$

with the constant and the market activity variables fixed. In each pair of regressions the first only contain the volume variables as regressors, whereas the second contain both the volume and quote variables. The motivation behind the pairs of regressions is that the volume and quote variables are overlapping in the sense that the quote variables also comprise quotes from banks in Norway. Comparing the two regressions in each pair thus enable us to answer three questions of interest, namely whether spot NOK/EUR volume matters for variability, whether global interbank NOK/EUR quoting frequency matters, and to what extent they overlap.

The results of the period variability regressions in table 5.10 are unambiguous. Neither in (5.15), where only the Norwegian market activity variables enter, nor in (5.16) are any of the market activity variables significant at conventional levels of significance. In other words, our results do not support the hypothesis that there is any impact of Norwegian nor global market activity on period variability when controlled for other variables. However, the drop in the coefficient values of $\Delta z_{t_2}^{tot}$ and $\bar{z}_{t_2-1}^{tot/2}$ when the global market activity variables are added may be interpreted as an indication of overlap.

The results of the range variability regressions in table 5.11 on the other hand provide some support for the hypothesis that market activity has an impact on range variability when controlling for other variables. In specification (5.17), where only Norwegian market activity enters, long term market activity $\bar{z}_{t_2-1}^{tot/2}$ is significant at 4%. Short term market Norwegian activity $\Delta z_{t_2}^{tot}$, however, is not significant. Adding global interbank measures of market activity (specification (5.18)) reduces the parameter estimates of the Norwegian variables, and increases the p-value of Norwegian long term market activity $\bar{z}_{t_2-1}^{tot/2}$ to 11%. This could be interpreted as the presence of some overlap. Also, long-term global interbank market activity $\bar{q}_{t_2-1}^{15}$ is almost significant with a p-value of 12%. Finally, short term global interbank market activity is significant at 8%.

5.3.4 Does Norwegian bank size matter?

The purpose of this subsection is to shed light on whether groups of similarly sized banks impact differently upon variability. Differently put, whether size matters. Table 5.12 contains estimates of parsimonious specifications obtained through simplification of

$$v_{t_2}^{se} = b_0 + b_2\text{persistence} + b_{12}\Delta z_{t_2}^{big} + b_{13}\bar{z}_{t_2}^{big} + \\ b_{14}\Delta z_{t_2}^{med} + b_{15}\bar{z}_{t_2}^{med} + b_{16}\Delta z_{t_2}^{sma} + b_{17}\bar{z}_{t_2}^{sma} + \text{the rest} + e_{t_2} \quad (5.19)$$

$$v_{t_2}^{hl} = b_0 + b_2 \text{persistence} + b_{12}\Delta z_{t_2}^{big} + b_{13}\bar{z}_{t_2}^{big} + b_{14}\Delta z_{t_2}^{med} + b_{15}\bar{z}_{t_2}^{med} + b_{16}\Delta z_{t_2}^{sma} + b_{17}\bar{z}_{t_2}^{sma} + \text{the rest} + e_{t_2} \quad (5.20)$$

keeping the constant and all the volume variables although significant at all levels. The estimation output do not point in a single direction. The results of the period variability specification (5.19) suggest that none of the market activity variables are significant at the 10% level. In the range variability specification (5.20) only one of the variables exhibit significance, namely long term market activity of big banks. The estimate is positive which means that higher long term market activity of the big banks increases variability. The insignificant estimates are conflicting. For example, the estimated impact of short term market activity of large banks is positive—as expected—in the period variability specification, whereas in the range variability specification the estimate is negative. Similar differences in the sign of the estimated impacts occur with respect to the market activity variables of the medium sized banks.

Not removing insignificant variables produces higher standard errors associated with the coefficient estimates, with the possible consequence that significant impacts of market activity are not revealed. To explore this possibility table 5.13 contains estimates of parsimonious specifications obtained through simplification of the same specifications (5.19) and (5.20) with only the constant being kept. In the range variability specification there is little change in the significance results, since the only retained market activity variable is that of big banks' long term activity (notably with low p-value). However, there is a change in the significance results of the period variability specification. Not fixing the market activity variables yields a significant estimate of small banks' long term market activity $\bar{z}_{t_2}^{sma}$ in table 5.13. The impact is negative and therefore the opposite sign of the estimated impact of large banks' long term activity $\bar{z}_{t_2}^{big}$ on range variability in the same table 5.13.

5.4 Conclusions

This study has sought to shed light on the role played by heterogeneity in the relation between market activity and exchange rate variability, through an investigation of the relation between a measure of NOK/EUR spot currency transaction volume in Norway and NOK/EUR variability. Whereas an increase in short term global interbank market activity (as measured by the relative increase in quoting frequency) increases range variability, our results do not support the hypothesis that increases in short term Norwegian market activity (as measured by the relative increase in our measure of spot NOK/EUR trading volume) has a statistically significant impact on neither period nor range variability. Moreover, we do not find support for the hypothesis that some groups of banks, for example big banks, have an impact on variability through their short term market activity. With respect to the impact of long term market activity, however, our results do suggest that Norwegian NOK/EUR trading has an impact. In particular, we find some support of the hypothesis that increased long term activity by banks in Norway increases range variability through their long term spot NOK/EUR trading, and that groups of similarly sized banks have different impacts. The group of small banks' long term market activity has a negative impact on period variability, whereas the group of large banks' long term market activity has a positive impact on range variability.

An area for further research that our results suggest should be pursued in particular is the impact Norwegian long term market activity has on variability. Our results suggest a negative impact on period variability from the long term market activity of small banks, and a positive impact on range variability from the long term market activity of large banks. It is not evident why this is the case, so further investigation—possibly approaching the issue in different ways—could shed further light on the issue.

Table 5.1: Descriptive statistics of period return, and period and range volatilities

		15/1/1999–7/1/2005 ($T = 313$)	15/1/1999–30/3/2001 ($T = 116$)	6/4/2001–7/1/2005 ($T = 197$)
$r_{t_2}^{se}$	*Avg.*	-0.04	-0.08	-0.01
	Med.	-0.06	-0.08	-0.04
	Max.	3.05	1.06	3.05
	Min	-2.14	-1.58	-2.14
	S.e.	0.52	0.45	0.55
$V_{t_2}^{se}$	*Avg.*	0.27	0.21	0.30
	Med.	0.10	0.10	0.10
	Max.	9.28	2.49	9.28
	Min	0.00	0.00	0.00
	S.e.	0.66	0.34	0.78
$V_{t_2}^{hl}$	*Avg.*	0.78	0.62	0.88
	Med.	0.56	0.45	0.62
	Max.	13.21	2.63	13.21
	Min	0.10	0.10	0.10
	S.e.	0.96	0.48	1.14
$v_{t_2}^{se}$	*Avg.*	-2.71	-2.85	-2.63
	Med.	-2.30	-2.34	-2.28
	Max.	2.23	0.91	2.23
	Min	-13.43	-10.21	-13.43
	S.e.	2.17	2.11	2.20
$v_{t_2}^{hl}$	*Avg.*	-0.57	-0.73	-0.48
	Med.	-0.58	-0.80	-0.48
	Max.	2.58	0.97	2.58
	Min	-2.31	-2.31	-2.29
	S.e.	0.78	0.72	0.8

Table 5.2: Sample correlations between period and range volatilities

Sample		$V_{t_2}^{se}$	$V_{t_2}^{hl}$		$v_{t_2}^{se}$	$v_{t_2}^{hl}$
15/01/1999 -	$V_{t_2}^{se}$	1.00		$v_{t_2}^{se}$	1.00	
7/1/2005	$V_{t_2}^{hl}$	0.90	1.00	$v_{t_2}^{hl}$	0.58	1.00
(T=313)						
15/1/1999 -	$V_{t_2}^{se}$	1.00		$v_{t_2}^{se}$	1.00	
30/3/2001	$V_{t_2}^{hl}$	0.70	1.00	$v_{t_2}^{hl}$	0.47	1.00
(T=116)						
6/4/2001 -	$V_{t_2}^{se}$	1.00		$v_{t_2}^{se}$	1.00	
7/1/2005	$V_{t_2}^{hl}$	0.93	1.00	$v_{t_2}^{hl}$	0.64	1.00
(T=197)						

Table 5.3: Descriptive statistics of volume and quote data

	Average	*Median*	*Max.*	*Min.*	*S.e.*
$Z_{t_2}^{tot}$	322672	317256	611128	107163	81303
$Z_{t_2}^{big}$	234255	232059	498021	65667	75207
$Z_{t_2}^{med}$	78972	77861	183174	3515	38336
$Z_{t_2}^{sma}$	9446	8377	27042	954	5548
Q_{t_2}	10304	2771	54917	75	15314
$\Delta z_{t_2}^{tot}$	0.0019	0.037	0.972	-0.998	0.26
$\Delta z_{t_2}^{big}$	0.0035	0.033	0.835	-1.120	0.30
$\Delta z_{t_2}^{med}$	-0.0035	0.000	2.493	-2.013	0.49
$\Delta z_{t_2}^{sma}$	0.0000	-0.015	1.926	-2.221	0.55
Δq_{t_2}	0.0118	0.001	2.996	-1.351	0.38

Note: The sample period is 15 January 1999 - 7 January 2005 (313 observations) and the Z_{t_2} variables in the four upper rows are in thousands of Norwegian kroner.

Table 5.4: Sample correlations between variables based on volume and quote data

	q_{t_2}	$\hat{q}_{t_2}$	q_{t_2-1}	$\bar{q}^{5}_{t_2-1}$	$\bar{q}^{15}_{t_2-1}$
q_{t_2}	1.00				
$\hat{q}_{t_2}$	0.86	1.00			
q_{t_2-1}	0.83	0.96	1.00		
$\bar{q}^{5}_{t_2-1}$	0.83	0.95	0.89	1.00	
$\bar{q}^{15}_{t_2-1}$	0.80	0.94	0.83	0.94	1.00

	$z^{tot}_{t_2}$	$\hat{z}^{tot}_{t_2}$	$z^{tot}_{t_2-1}$	$\bar{z}^{tot/2}_{t_2-1}$	$\bar{z}^{tot/3}_{t_2-1}$
$z^{tot}_{t_2}$	1.00				
$\hat{z}^{tot}_{t_2}$	0.60	1.00			
$z^{tot}_{t_2-1}$	0.50	0.82	1.00		
$\bar{z}^{tot/2}_{t_2-1}$	0.52	0.88	0.87	1.00	
$\bar{z}^{tot/3}_{t_2-1}$	0.55	0.91	0.79	0.94	1.00

	$z^{big}_{t_2}$	$\hat{z}^{big}_{t_2}$	$z^{big}_{t_2-1}$	$\bar{z}^{big/2}_{t_2-1}$	$\bar{z}^{big/3}_{t_2-1}$
$z^{big}_{t_2}$	1.00				
$\hat{z}^{big}_{t_2}$	0.73	1.00			
$z^{big}_{t_2-1}$	0.62	0.84	1.00		
$\bar{z}^{big/2}_{t_2-1}$	0.64	0.87	0.90	1.00	
$\bar{z}^{big/3}_{t_2-1}$	0.67	0.91	0.84	0.95	1.00

	$z^{med}_{t_2}$	$\hat{z}^{med}_{t_2}$	$z^{med}_{t_2-1}$	$\bar{z}^{med/2}_{t_2-1}$	$\bar{z}^{med/3}_{t_2-1}$
$z^{med}_{t_2}$	1.00				
$\hat{z}^{med}_{t_2}$	0.80	1.00			
$z^{med}_{t_2-1}$	0.73	0.91	1.00		
$\bar{z}^{med/2}_{t_2-1}$	0.76	0.95	0.93	1.00	
$\bar{z}^{med/3}_{t_2-1}$	0.78	0.97	0.89	0.97	1.00

	$z^{sma}_{t_2}$	$\hat{z}^{sma}_{t_2}$	$z^{sma}_{t_2-1}$	$\bar{z}^{sma/2}_{t_2-1}$	$\bar{z}^{sma/3}_{t_2-1}$
$z^{sma}_{t_2}$	1.00				
$\hat{z}^{sma}_{t_2}$	0.72	1.00			
$z^{sma}_{t_2-1}$	0.64	0.87	1.00		
$\bar{z}^{sma/2}_{t_2-1}$	0.67	0.91	0.91	1.00	
$\bar{z}^{sma/3}_{t_2-1}$	0.67	0.92	0.85	0.96	1.00

Note: The sample period is 15 January 1999 - 7 January 2005 (313 weekly observations).

Table 5.5: Regressions of log of period variability $v_{t_2}^{se}$ on a constant and market activity variables

	(5.5)		(5.6)	
	Est.	*Pval.*	*Est.*	*Pval.*
$const.$	-12.383	0.09	-5.789	0.00
$\Delta z_{t_2}^{tot}$	0.583	0.26		
$\bar{z}_{t_2-1}^{tot/2}$	0.769	0.18		
Δq_{t_2}			0.448	0.11
$\bar{q}_{t_2-1}^{15}$			0.451	0.04
$id_{t_2}^2$	-10.985	0.00	-10.518	0.00
$id_{t_2}^3$	-11.082	0.00	-11.212	0.00
R^2	0.17		0.18	
AR_{1-10}	5.31	0.87	3.57	0.96
$ARCH_{1-10}$	10.81	0.37	14.40	0.16
$Het.$	2.15	0.91	3.62	0.73
$Hetero.$	4.43	0.73	4.09	0.77
JB	53.91	0.00	54.14	0.00
$Obs.$	311		298	

Note: The sample period is 15 January 1999 - 7 January 2005 (313 weekly observations), computations are in EViews 5.1 with OLS estimation and standard errors are of the White (1980) type. *Pval* stands for *p*-value and corresponds to a two-sided test with zero as null, AR_{1-10} is the χ^2 version of the Lagrange-multiplier test for serially correlated residuals up to lag 10, $ARCH_{1-10}$ is the χ^2 version of the Lagrange-multiplier test for serially correlated squared residuals up to lag 10, $Het.$ and $Hetero.$ are White's (1980) heteroscedasticity tests without and with cross products, respectively, and JB is the Jarque and Bera (1980) test for non-normality.

Table 5.6: Regressions of log of range variability $v_{t_2}^{hl}$ on a constant and market activity variables

	(5.7)		(5.8)	
	$Est.$	$Pval.$	$Est.$	$Pval.$
$const.$	-10.678	0.00	-2.852	0.00
$\Delta z_{t_2}^{tot}$	0.213	0.31		
$\bar{z}_{t_2-1}^{tot/2}$	0.798	0.00		
Δq_{t_2}			0.284	0.15
$\bar{q}_{t_2-1}^{15}$			0.324	0.00
R^2	0.05		0.06	
AR_{1-10}	29.33	0.00	26.28	0.00
$ARCH_{1-10}$	4.86	0.90	10.65	0.39
$Het.$	5.90	0.21	7.62	0.11
$Hetero.$	27.82	0.00	7.94	0.16
JB	1.69	0.43	4.80	0.09
$Obs.$	311		298	

Note: Standard errors are of the Newey and West (1987) type, otherwise see table 5.5.

Table 5.7: Regressions of log of variabilities $v_{t_2}^{se}$ and $v_{t_2}^{hl}$, respectively, on a constant and persistence

	(5.9)		(5.10)	
	Est.	*Pval.*	*Est.*	*Pval.*
const.	-2.640	0.00	-0.432	0.00
persistence	0.100	0.07	0.027	0.00
$id_{t_2}^2$	-10.728	0.00		
$id_{t_2}^3$	-10.739	0.00		
R^2	0.16		0.11	
AR_{1-10}	4.98	0.89	11.03	0.36
$ARCH_{1-10}$	8.01	0.63	6.35	0.78
Het.	1.41	0.84	1.02	0.60
Hetero.	1.41	0.84	1.02	0.60
JB	54.32	0.00	2.10	0.35
Obs.	313		312	

Note: See table 5.5.

Table 5.8: Regressions of log of period variability $v_{t_2}^{se}$ on a constant, persistence and market activity variables

	(5.11)		(5.12)	
	$Est.$	$Pval.$	$Est.$	$Pval.$
$const.$	-11.523	0.12	-5.467	0.00
$persistence$	0.087	0.12	0.111	0.03
$\Delta z_{t_2}^{tot}$	0.579	0.26		
$\bar{z}_{t_2-1}^{tot/2}$	0.701	0.23		
Δq_{t_2}			0.376	0.19
$\bar{q}_{t_2-1}^{15}$			0.406	0.06
$id_{t_2}^2$	-10.916	0.00	-10.478	0.00
$id_{t_2}^3$	-11.016	0.00	-11.114	0.00
R^2	0.17		0.19	
AR_{1-10}	4.98	0.89	3.34	0.97
$ARCH_{1-10}$	11.44	0.32	15.00	0.13
$Het.$	2.85	0.94	4.92	0.77
$Hetero.$	6.20	0.86	6.01	0.87
JB	52.48	0.00	52.88	0.00
$Obs.$	311		298	

Note: See table 5.5.

Table 5.9: Regressions of log of range variability $v_{t_2}^{hl}$ on a constant, persistence and market activity variables

	(5.13)		(5.14)	
	$Est.$	$Pval.$	$Est.$	$Pval.$
$const.$	-7.248	0.02	-1.960	0.00
$persistence$	0.024	0.00	0.025	0.00
$\Delta z_{t_2}^{tot}$	0.141	0.50		
$\bar{z}_{t_2-1}^{tot/2}$	0.538	0.03		
Δq_{t_2}			0.287	0.07
$\bar{q}_{t_2-1}^{15}$			0.216	0.01
R^2	0.13		0.14	
AR_{1-10}	10.30	0.41	7.85	0.64
$ARCH_{1-10}$	4.42	0.93	5.34	0.87
$Het.$	4.43	0.62	7.61	0.27
$Hetero.$	31.28	0.00	10.37	0.32
JB	1.94	0.38	2.88	0.24
$Obs.$	311		298	

Note: See table 5.5.

Table 5.10: Parsimonious regressions of log of period variability $v_{t_2}^{se}$ on a constant, persistence and market activity controlling for other variables

	(5.15)		(5.16)	
	$Est.$	$Pval.$	$Est.$	$Pval.$
$const.$	-9.116	0.20	-7.401	0.34
$persistence$	0.096	0.08	0.105	0.04
$\Delta z_{t_2}^{tot}$	0.446	0.37	0.398	0.45
$\bar{z}_{t_2-1}^{tot/2}$	0.512	0.36	0.179	0.78
Δq_{t_2}			0.334	0.25
$\bar{q}_{t_2-1}^{15}$			0.352	0.13
$x_{t_2}^{w}$	0.111	0.06		
$ir_{t_2}^{no,b\Delta}$			0.006	0.00
$ir_{t_2}^{no,c\Delta}$	0.015	0.00	0.015	0.00
id_1	-10.866	0.00	-10.549	0.00
id_2	-10.462	0.00	-11.134	0.00
R^2	0.20		0.21	
AR_{1-10}	5.35	0.87	3.40	0.97
$ARCH_{1-10}$	12.20	0.27	17.56	0.06
$Het.$	11.23	0.51	9.90	0.87
$Hetero.$	19.99	0.58	17.43	0.99
JB	44.25	0.00	54.98	0.00
$Obs.$	311		298	

Note: See table 5.5.

Table 5.11: Parsimonious regressions of log of range variability $v_{t_2}^{hl}$ on a constant, persistence and market activity controlling for other variables

	(5.17)		(5.18)	
	Est.	*Pval.*	*Est.*	*Pval.*
$const.$	-6.863	0.02	-6.996	0.03
$persistence$	0.023	0.00	0.021	0.00
$\Delta z_{t_2}^{tot}$	0.103	0.62	0.020	0.93
$\bar{z}_{t_2-1}^{tot/2}$	0.504	0.04	0.431	0.11
Δq_{t_2}			0.250	0.08
$\bar{q}_{t_2-1}^{15}$			0.147	0.12
$ir_{t_2}^{no,b\Delta}$	0.001	0.00	0.002	0.00
$ir_{t_2}^{no,c}$	0.001	0.03	0.001	0.02
$ir_{t_2}^{no,c\Delta}$	0.007	0.00	0.007	0.00
r_{t_1}	0.137	0.02	0.103	0.07
R^2	0.17		0.20	
AR_{1-10}	11.55	0.32	9.95	0.45
$ARCH_{1-10}$	4.31	0.93	4.52	0.92
$Het.$	21.62	0.09	27.63	0.07
$Hetero.$	73.07	0.00	101.31	0.00
JB	1.23	0.54	2.44	0.29
$Obs.$	311		298	

Note: See table 5.5.

Table 5.12: Parsimonious regressions of log of period $v^{se}_{t_2}$ and range $v^{hl}_{t_2}$ variability on a constant, persistence and disaggregated volume variables controlling for other other variables

	(5.19)		(5.20)	
	Est.	*Pval.*	*Est.*	*Pval.*
const.	-2.886	0.70	-4.804	0.13
persistence	0.091	0.10	0.023	0.00
$\Delta z^{big}_{t_2}$	0.027	0.95	-0.093	0.60
$\bar{z}^{big/2}_{t_2-1}$	0.265	0.56	0.378	0.05
$\Delta z^{med}_{t_2}$	-0.110	0.65	0.014	0.90
$\bar{z}^{med/2}_{t_2-1}$	-0.067	0.73	0.043	0.58
$\Delta z^{sma}_{t_2}$	0.387	0.18	0.090	0.37
$\bar{z}^{sma/2}_{t_2-1}$	-0.253	0.29	-0.093	0.29
$x^{w}_{t_2}$	0.115	0.05		
$ir^{no,b}_{t_2}$			0.001	0.07
$ir^{no,b\Delta}_{t_2}$			0.001	0.00
$ir^{no,c}_{t_2}$			0.001	0.05
$ir^{no,c\Delta}_{t_2}$	0.014	0.00	0.007	0.00
r_{t_1}			0.122	0.04
id_1	-10.636	0.00		
id_2	-10.146	0.00		
R^2	0.21		0.19	
AR_{1-10}	4.72	0.91	12.12	0.28
$ARCH_{1-10}$	12.62	0.25	3.37	0.97
Het.	23.71	0.26	27.59	0.28
Hetero.	43.89	0.88	128.36	0.00
JB	38.61	0.00	1.08	0.58
Obs.	311		311	

Note: See table 5.5.

Table 5.13: Parsimonious regressions of log of period $v_{t_2}^{se}$ and range variability $v_{t_2}^{hl}$ on a constant, persistence and disaggregated volume variables controlling for other variables

	(5.19)		(5.20)	
	Est.	*Pval.*	*Est.*	*Pval.*
$const.$	1.077	0.54	-5.847	0.00
$persistence$	0.098	0.07	0.023	0.00
$\bar{z}_{t_2-1}^{big/2}$			0.434	0.00
$\bar{z}_{t_2-1}^{sma/2}$	-0.414	0.04		
$x_{t_2}^{w}$	0.113	0.06		
$ir_{t_2}^{no,b}$			0.001	0.07
$ir_{t_2}^{no,b\Delta}$			0.001	0.00
$ir_{t_2}^{no,c}$			0.001	0.02
$ir_{t_2}^{no,c\Delta}$	0.015	0.00	0.007	0.00
r_{t_1}			0.134	0.02
id_1	-10.666	0.00		
id_2	-10.529	0.00		
R^2	0.20		0.18	
AR_{1-10}	5.02	0.89	12.57	0.25
$ARCH_{1-10}$	12.18	0.27	3.93	0.95
$Het.$	14.84	0.14	16.88	0.26
$Hetero.$	18.88	0.27	48.07	0.01
JB	39.71	0.00	1.53	0.47
$Obs.$	311		311	

Note: See table 5.5.

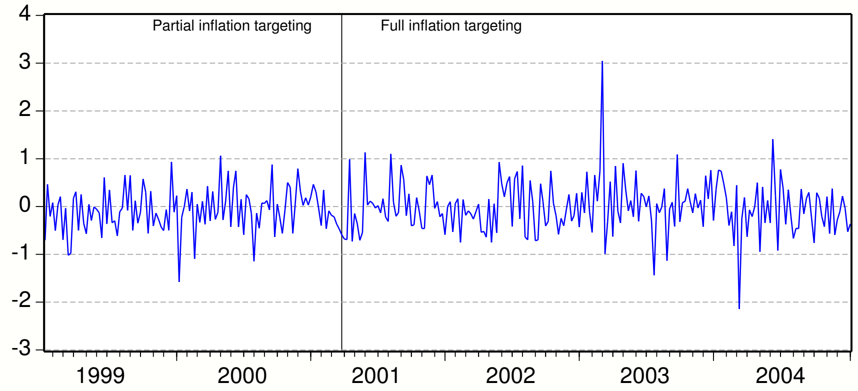

Figure 5.1: Period return $r_{t_2}^{se}$ over the last two trading days of week t from 15 January 1999 to 7 January 2005

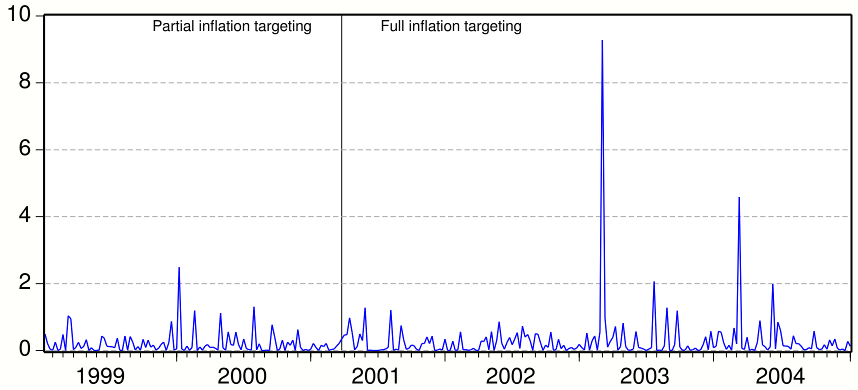

Figure 5.2: Period variability $V_{t_2}^{se}$ from 15 January 1999 to 7 January 2005

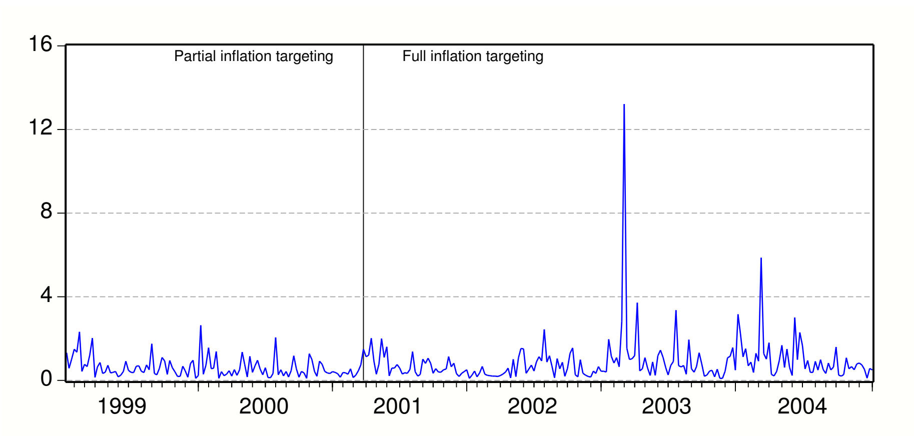

Figure 5.3: Range variability $V_{t_2}^{hl}$ from 15 January 1999 to 7 January 2005

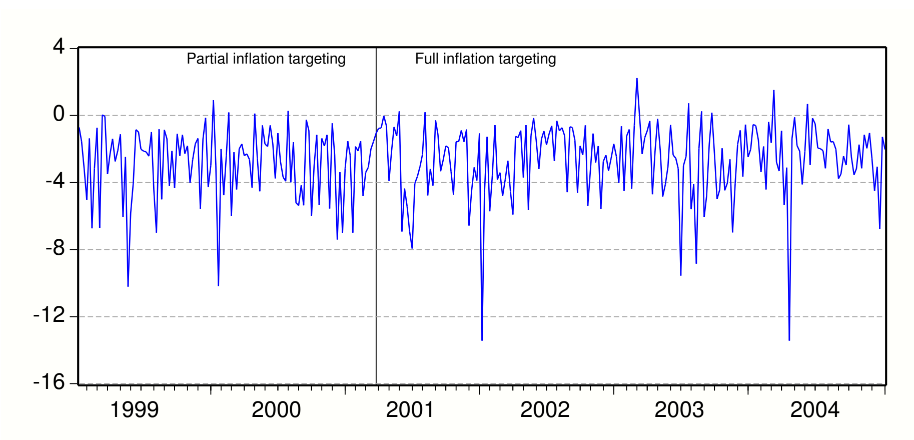

Figure 5.4: Log of period variability $v_{t_2}^{se}$ from 15 January 1999 to 7 January 2005

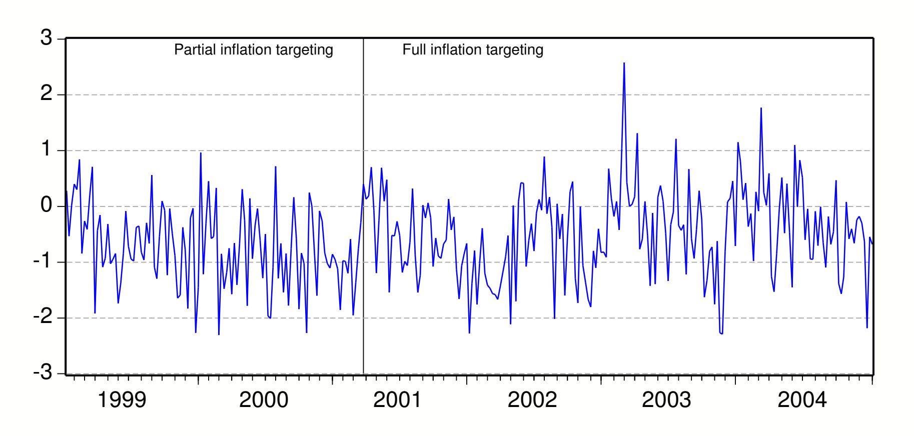

Figure 5.5: Log of range variability $v_{t_2}^{hl}$ from 15 January 1999 to 7 January 2005

Chapter 6

The first stage in Hendry's reduction theory revisited

6.1 Introduction

When Trygve Haavelmo suggested that the n observations in a dataset could "be considered as *one* observation of n variables...following an n-dimensional *joint* probability law" (1944, p. iii), his main aim was to convert more economists to the praxis of evaluating economic theories against empirical economic data using statistical techniques. The deeper question about how the joint n-dimensional probability distribution was related to reality, however, he remained agnostic about. In his own words, the existence of such a joint probability distribution "may be purely hypothetical" (same place, p. iii).[33] Although Haavelmo's ideas had a profound and immediate impact on contemporary economic analysis it nevertheless took until the 1970s and 1980s before a systematic approach to the study of the relation between reality and models thereof in terms of probability concepts developed in econometrics. At the centre of several important contributions during these years, including Florens and Mouchart (1980, 1985), Hendry and Richard (1990), and Florens et al. (1990), was the notion of a "probabilistic reduction", that is, the idea of replacing a complex probabilistic structure with a simpler one through marginalisation and/or conditioning, and led to the development of important econometric concepts like weak exogeneity, strong exogeneity and super exogeneity (see Engle et al. (1983)).

The reduction theory of David F. Hendry (1995, chapter 9), where the term reduction is used in a broader way than originally, provides an overall probabilistic framework for the analysis and classification of the simplifications associated with empirical models.[34] In Hendry's own words "it seeks to explain the origin of empirical models in terms of reduction operations conducted implicitly on the DGP [data generating process] to induce the relevant empirical model" (1995, p. 344). Starting with the joint probability distribution of the "complete set" of theory "variables relevant to the economy under investigation" (same place, p. 345), the reduction theory distinguishes between twelve reduction operations which ultimately leads to the empirical model. Although Hendry's theory is a powerful and comprehensive framework for the analysis of the relation between models and reality, it is nevertheless unable to provide an analysis on the same probability space of the first stage (and hence of the subsequent stages) of reduction given a commonplace theory of social reality. The commonplace theory consists of the joint hypotheses that, literally, a) the course of history is indeterministic, b) history does not repeat itself, and c) the future depends on the past (historical inheritance). In philosophical jargon, that the human world is made up of indeterministic, historically supervenient particulars.[35] According to Hendry the economic mechanism under study, that is, the joint distribution of the complete set of relevant theory variables defined on the underlying probability space, is an entity that *can* change but does not necessarily do so. In other words, the economic mechanism is a probabilistic statement about how variables are related, which may change but does not do so necessarily. According to the commonplace theory of social reality on the other hand, everything is changing all the time—including how variables are related—in a way that is indeterministically dependent on the past. Conceptually this is not necessarily incompatible with Hendry's setup, but it would nevertheless imply that the economic mechanism as defined by Hendry is constantly changing. As a consequence, Hendry's theory is unable to provide a probabilistic analysis on the same underlying probability space of how the theory variables are related. Moreover, according to Hendry the underlying probability space is transformed—again—when data are collected, so the theory is unable to provide a probabilistic analysis on the same underlying probability space of the relation between the theory and data variables.

In this chapter I propose a certain structure on the underlying outcome space which implies that the associated probability space does not change even though the human world is changing all the time, and which implies that the probability space does not change when data are collected. This is enabled by devising the outcome set as consisting of possible worlds made up of indeterministic and historically inherited particulars and provides several gains and possibilities of which only a few are explored in this chapter. First, the formulation of theory variables can be seen as a simplification and can thus be interpreted as the "perspective" from which we study an issue. Second, a probabilistic definition of measurement validity, that is, the absence of measurement error, is enabled. Third, a history based probabilistic definition of indeterministic causality that nests discrete, continuous and "interval" versions of probabilistic causality is proposed. Finally, mathematical expectation conditional on an information set is re-interpreted.

More generally devising the outcome set as consisting of indeterministic worlds made up of historically inherited particulars also provides a bridge between econometric (/probabilistic) reduction analysis and metaphysics, that is, the part of philosophy that deals with what there exists and its nature. There is already a voluminous philosophical literature that employs the idea of possible worlds to shed light on various metaphysical issues, and by providing a bridge between these two literatures econometrics might eventually benefit from these insights.

The rest of this chapter is organised into five sections. In the next, section 6.2, the most relevant parts of Hendry's reduction theory is detailed. Section 6.3 describes and motivates the structure of the outcome space that is proposed. Section 6.4 details in what sense the formulation of theory variables can be seen as a simplification, and contains the proposed definition of measurement validity, that is, absence of measurement error. Section 6.5 formulates the proposed definition of causality and re-interprets mathematical expectation conditional on an information set. Finally, section 6.6 concludes.

6.2 The first stage in Hendry's reduction theory

The purpose of Hendry's reduction framework is "to explain the origin of empirical models in terms of reduction operations conducted implicitly on the DGP" (1995, p. 344), and his framework details twelve reductions whose order is not unique.[36] Since the focus is on the first stage I concentrate on this in what follows.

The most informative account of the first stage of reduction is given in a single paragraph in chapter 9 of *Dynamic Econometrics* (1995), which is an adaptation of Hendry and Cook (1994). Most of the paragraph is about the concepts and actions involved in the first stage, so it seems useful to reproduce it here almost in its entirety. Note however that I have modified Hendry's notation in order to retain a consistent notation throughout this chapter. Most importantly, random variables and vectors appear in capitals. This is to distinguish them from their realisations, which I denote in small letters later in the chapter. The passage is:

> "The analysis begins with the complete set of random variables $\{\mathbf{U}_t^*\}$ relevant to the economy under investigation over a time span $t = 1, \ldots, T$, where the superscript * denotes a perfectly measured variable $\mathbf{U}^* = (\mathbf{U}_1^*, \ldots, \mathbf{U}_T^*)$, defined on the probability space $(\Omega, \mathcal{F}, P)$ ... The $\{\mathbf{U}_t^*\}$ comprise all the potential variables from the economic mechanism under study which operates at the level of $\mathbf{U}^*$, and hence the vector $\mathbf{U}_t^*$ comprises details of every economic action of every agent at time t in all the regions of the geographical space relevant to the analysis. However, many of the $\{U_{ti}^*\}$ variables are either unobserved or badly measured, so the term data is not strictly applicable to $\mathbf{U}_t^*$. The mapping from the economic mechanism to the data-generation process through the measurement system is the first reduction, which can lose a vast amount of information, and introduce inaccuracy but leads

> to a data-set which is denoted by $\{\mathbf{U}_t\}$. At a conceptual level, all variables $\{U^*_{ti}\}$ are assumed to be measured as $\{U_{ti}\}$ although for some variables, the level of quantification may be low, possibly even an artificial entry of zero. The probability space $(\Omega, \mathcal{F}, P)$ is transformed by the measurement process (usually markedly) ..."—Hendry (1995, p. 345)

Thus the starting point of Hendry's reduction framework is all the variables relevant to an economy under investigation. This set is here denoted $\mathbf{U}^*$ and is defined on the probability space $(\Omega, \mathcal{F}, P)$. Together, $\mathbf{U}^*$ and $(\Omega, \mathcal{F}, P)$ constitute the "economic mechanism" on which the first reduction is performed. A change in the relation between the theory variables would entail a change in the underlying probability space $(\Omega, \mathcal{F}, P)$. The actions of collecting and recording the data, that is, the measurement process, eventually produces the first reduction, a dataset $\mathbf{U}$ defined on an altered probability space $(\Omega', \mathcal{F}', P')$, and is called the "data generating process" (DGP). Schematically the first stage of reduction is summarised in table 1.

6.3 The outcome set as consisting of possible worlds

If $(\Omega, \mathcal{F}, P)$ denotes a probability space with Ω, $\mathcal{F}$ and P being the outcome space, the event space and the probability measure, respectively, then in what follows the elements $\omega \in \Omega$ will be referred to as "worlds" or "possible worlds". The purpose of this section is to formulate and motivate the proposed structure of the worlds ω. The proposed structure serves as some sort of social ontology, that is, a theory of the nature of social reality, and is contained in definition 4 in the last subsection of this section. It may be a good idea to have a look at it before reading the preceding subsections that provide the details. The first subsection 6.3.1 presents the idea of a possible world which in philosophy has proved very useful in analysing, discussing and communicating many philosophical ideas and theories, and shows that there is no loss of generality in interpreting the ω as worlds. Then, subsections 6.3.2 and 6.3.3 formalise the ideas of contingent particularism and historically inherited particulars, respectively. Finally, subsection 6.3.4 contains the definition of outcomes sets consisting of indeterministic worlds made up of historically inherited particulars.

6.3.1 Possible worlds

The idea of a world is normally credited to the German philosopher and mathematician Gottfried Wilhelm Leibniz (1646 - 1716) (Crane 1995).[37] Intuitively a world contains everything in the past, everything in the present and everything in the future, or in Leibniz' own words "the entire sequence and the entire collection of all existing things" (*Theodicy*, par. 8, G VI 107. Quoted in Parkinson 1995, p. 213). In contemporary philosophy the notion is often associated with David Lewis (1941-2001), who describes worlds as consisting of

> "the planet Earth, the solar system, the entire Milky Way, the remote galaxies we see through telescopes... Anything at any distance at all is to be included. Likewise the world is inclusive in time. No long-gone ancient Romans, no long-gone pterodactyls, no long-gone primordial clouds of plasma are too far in the past, nor are the dead dark stars too far in the future, to be part of this same world"—Lewis (1986b, p.

1)

Where Lewis differ from Leibniz is with respect to how things are connected and with respect to the existence of other worlds. Whereas Leibniz was a determinist and believed in a single world, Lewis was a non-determinist and believed in the rather unusual thesis that non-actual, possible worlds exist objectively and independent of thought because "philosophy [his own?] goes more easily" if we believe so (1986b, p. vii).[38] Although I take side with Lewis in the determinism *vs.* indeterminism debate my view differs most certainly from Lewis' regarding the existence of non-actual worlds, since I only see them as useful mind-constructs not existing independent of thought.[39]

But do we really need the whole world for the purpose of econometric reduction analysis? Spatially, yes, if we want to ensure a complete analysis, but it is not necessary to be all-including backwards and forward in time. Differently put, the worlds must contain everything between a start point and an end point, but the portions outside this interval are not really necessary although including them changes little. So henceforth I will devise a world ω as a continuous time process $\{w(t) : t \in [0, \infty)\}$ of worldly states-of-affairs $w(t)$, where $[0, \infty)$ is contained in the set of real numbers $\mathbb{R}$. The number 0 denotes an arbitrary starting point, say, yesterday at midnight or four million years ago, and is not restrictive. However, bounding worlds temporally backwards in time entails an implicit conditioning on the realised history preceding 0. Backwards bounding thus means probabilities acquire an interpretation of special interest, but apart from this the only function bounding serves is to simplify the exposition.

Interpreting ω as worlds retains the intuitive use of probability algebra. For example, if we want to say that $A \in \mathcal{F}$ denotes the event that (say) 10% of the labour force of an economy is unemployed at t, then the only change in interpreting the ω as a world is that A now denotes the set of all worlds in which 10% of the labour force of a certain economy is unemployed at t. More formally, $A = \{\omega : 10\% \text{ unemployed at } t\}$. If the worlds are bounded backwards, then the interpretation becomes that A denotes the set of all worlds in which 10% of an economy is unemployed at t given the history of the world up to $t = 0$. Another common practice is to interpret the outcome set Ω as a set of possible "states-of-affairs" or "facts". In possible worlds terminology a state-of-affair or fact at t is now the set of all worlds in which a certain state-of-affairs or fact attains at t. Finally, the possible worlds interpretation also accommodates "interval" events. With respect to the unemployment example, the event A now becomes the set of worlds in which 10% of the labour force of an economy is registered as unemployed over the time interval, say, $[t_0, t_1]$.

6.3.2 Contingent particularism

"I am inclined", in the words of Geoffrey Hawthorn, "to the view that the human world consists of contingent particulars" (1995, p. 10). Contingency refers to the thesis that social events are not connected in a deterministic manner, a question that has occupied philosophers for thousands of years. There are at least two philosophical literatures of relevance for this issue. The first is concerned with whether human being is endowed with a socalled "free will" and if so what kind of free will. The second literature is the socalled "philosophy of mind" literature and starts from two seemingly contradictory views. On the one hand that human being presumably is made up of a finite number of indivisible objects, usually referred to as particles, and on the other hand that human being is capable of a presumably infinite number of mental states (imagination, thought, etc.).[40] Depending on one's views on free will and on the relationship between mind and matter, a variety of possible views on how social events are connected is possible. Since I am unlikely to convince the reader of my belief in the indeterminism thesis unless she or he is already a believer I merely state the thesis as some sort of axiom that I start from. Formally, with respect to the probability space $(\Omega, \mathcal{F}, P)$, indeterminism is simply characterised by Ω containing more than one element, that is, part a) in definition 4.

The meaning of the philosophical idea of a "particular" is best understood when contrasted with its opposite, a "universal". In brief, something is said to be of particular nature if there exists only one of its kind, whereas something is said to be of universal nature if it is one out of several of its kind or type. Another way to put it is that a particular refers to the unique and non-repeatable, whereas a universal refers to the repeatable. In the current context particularism is the thesis that, literally, history does not repeat itself (no two points in time are exactly equal in all respects).[41] Formally this may be stated as follows.

Definition 1. Worldly particularism. A world $\omega = \{w(t) : t \in [0, \infty)\} \in \Omega$ is said to be made up of particulars if for all pairs t, $t' \in [0, \infty)$ such that $t \neq t'$ and $w(t)$, $w(t') \in \omega$, then $w(t) \neq w(t')$.

6.3.3 Historically inherited particulars

A further thesis I start from is that the current and the future depends on and inherit the characteristics of the past. Differently put, every turn history takes contributes in one or another way to the characteristics of the worldly state-of-affairs of the future. This thesis I shall call "historical inheritance", but before providing a formal formulation of this property we need the idea of a worldly state-of-affairs process up to t.

> **Definition 2. Worldly state-of-affairs process.** The process $\omega_t = \{w(a) : a < t, t \in (0, \infty)\} \subsetneq \omega$ is said to be a worldly states-of-affairs process up to but not including t.

So intuitively ω_t is a history up to t and note that the number 0 is not included in the interval $(0, \infty)$ in order to ensure that ω_t is non-empty. We can now define historical inheritance.

> **Definition 3. Historical inheritance.** The outcome space Ω is said to consist of worlds made up of historically inherited particulars if:
>
> a) All $\omega \in \Omega$ are made up of particulars.
>
> b) For all pairs of unequal worlds ω^1, $\omega^2 \in \Omega$, that is, $\omega^1 \neq \omega^2$: If $\omega_t^1 \neq \omega_t^2$, then $w^1(t') \neq w^2(t')$ for all $t' \in [t, \infty)$, where $w^1(t') \in \omega^1$ and $w^2(t') \in \omega^2$.

In words, if two worlds contains the same history up to t (but not at t), then the two worlds differ from each other in at least one respect at every point in the future, that is, from t and onwards.

6.3.4 Outcome sets consisting of indeterministic worlds made up of historically inherited particulars

The proposed structure of the worlds ω is contained in definition 4. The definition summarises the ideas of this section and provides the starting point for what follows.

> **Definition 4. Outcome set consisting of indeterministic worlds made up of historically inherited particulars.** Let $(\Omega, \mathcal{F}, P)$ be a probability space and let each $\omega \in \Omega$ be equal to a non-stochastic continuous time process $\{w(t) : t \in [0, \infty)\}$ with $[0, \infty) \subset \mathbb{R}$. The outcome space Ω is said to consist of possible worlds made up of indeterministic and historically inherited particulars if:
>
> a) There exists more than one element in Ω (indeterminism).
>
> b) For each $\omega \in \Omega$: For all pairs t, $t' \in [0, \infty)$ such that $t \neq t'$ and $w(t)$, $w(t') \in \omega$, then $w(t) \neq w(t')$ (particularism).
>
> c) For each pair of unequal worlds $\omega^1, \omega^2 \in \Omega$, that is, $\omega^1 \neq \omega^2$: If $\omega^1_t \neq \omega^2_t$ then $w^1(t') \neq w^2(t')$ for all $t' \in [t, \infty)$, where $w^1(t') \in \omega^1$ and $w^2(t') \in \omega^2$ (historical inheritance).

The first property a) essentially states that the course of history is indeterministic. If Ω contained only a single world, then this would imply that no other worlds are possible and therefore that the course of history is deterministic. The second property b) makes use of the notion "states-of-affairs" at t which is denoted $w(t)$, and essentially states that history does not repeat itself. The third and final property c) imposes a certain structure on the history-does-not-repeat itself property. Specifically, it ensures that future properties are shaped by the past, thus the terminology "historical inheritance".

6.4 The first stage in Hendry's reduction theory revisited

Definition 4 in subsection 6.3.4 provides the starting point of this section. The economic mechanism under study is defined as the joint distribution of the theory variables, together with the probability space upon which they are defined. It follows in a straightforward manner that an economic mechanism is non-changing when the theory variables are defined on the probability space, since all the change takes place at the level of the worlds in the outcome set rather than at the level of the joint distribution that relates the variables. The formulation of theory variables can therefore be given a specific interpretation of practical use in econometrics. This interpretation is outlined in subsection 6.4.1. It also follows in a straightforward manner that the collection of data variables does not alter the underlying probability space. To see this recall that any realisation of the data variable $\mathbf{U}$ corresponds to the worlds in which the data were collected or could have been collected. For example, for any realisation $\mathbf{u}_t$ of $\mathbf{U}_t$ there is an associated set of possible worlds $\{\omega : \mathbf{U}_t(\omega) = \mathbf{u}_t\}$ in which these data realisations can be obtained. Also, if we would like to restrict ourselves to the worlds enabled by history, then we can restrict ourselves to the intersection of $\{\omega : \mathbf{U}_t(\omega) = \mathbf{u}_t\}$ and the set of possible worlds enabled by the course of history preceding t. The second and final subsection 6.4.2 proposes a formal definition of measurement validity, that is, the absence of measurement error, which is enabled by the property that the probability space does not change due to the measurement process.

6.4.1 Formulation of theoretical variables as a reduction

Normative analysis is about how things should be, it is said, whereas positive analysis is value-independent and "objective" investigation of how things are. But is positive analysis entirely objective? Do we not, in any investigation, choose which questions to address, which portions of social reality to study, and which categorical schemes, concepts, techniques and language to employ? The idea that these choices are non-objective in some sense is old and in my view not controversial. Examples of economists who held this view are Max Weber (1994), Joseph Schumpeter (1949) and Gunnar Myrdal (1953, pp. vii-viii; 1969). Since a world contains everything and since the outcome set contains all the possible worlds, the formulation of theoretical variables defined on the probability space can be seen as reflecting some of these choices. In particular, the formulation of theoretical variables can be seen to reflect which portions of reality are studied as opposed to others, that is, as some sort of "marginalisation" of the variables that are not studied. Differently put, the formulation of theoretical variables can be seen as the "conceptual lenses" we view reality with. For example, in delineating and defining theoretical price and theoretical quantity, then other aspects of the transaction process are not included in the analysis. This is clearly an abstraction, since an anthropologist or an institutional economist might be interested in whether the parties engaged in any form of negotiation, whether there were implicit power-relations governing the transaction process, or what the means of transactions were. All this and many other aspects of the transaction are excluded from the analysis when the only theoretical variables delineated are price and quantity. With the modified probability space the selection of which portions of reality to analyse and the way they are depicted in terms of variables can be treated as a simplification.

6.4.2 A definition of measurement validity

The idea of measurement error is familiar to econometricians. In the methodological literature of the social sciences, discussions of measurement error are often couched in terms of theoretical or nominal or concept definition *vs.* measure or indicator or operational definition—see for example de Vaus (2001, pp. 24-33), Punch (1998, pp. 47-48) and Crano and Brewer (2002, pp. 5-12). That is, to what extent a measure (say, the number of people receiving unemployment benefits) is capable of providing information about a theoretical definition (say, the number of unemployed). An operational definition that satisfactorily provides the information sought is thus said to be measurement valid or concept valid.

To recall, random variables are denoted in capitals and their realisation in small letters. For example, a realisation of the theoretical vector of variables $\mathbf{U}^*$ is denoted $\mathbf{u}^* = (\mathbf{u}_1^*, \mathbf{u}_2^*, \ldots, \mathbf{u}_t^*, \ldots, \mathbf{u}_T^*)$, with $\mathbf{u}_t^* = (u_{t1}^*, u_{t2}^*, \ldots, u_{ti}^*, \ldots, u_{tI(t)}^*)$ for each t, where the symbolism $I(t)$ means the number of theoretical variables can vary with t. Similarly, a realisation of the vector of data variables $\mathbf{U}$ is denoted $\mathbf{u} = (\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_t, \ldots, \mathbf{u}_T)$, with $\mathbf{u}_t = (u_{t1}, u_{t2}, \ldots, u_{tj}, \ldots, u_{tJ(t)})$ for each t, where the symbolism $J(t)$ means the number of data variables can vary with t. $J(t)$ may of course differ from $I(t)$. Ideally a definition of measurement validity of $\mathbf{U}^*$ should be sequential and formulated for a sequence of pairs $(\mathbf{U}_1^*, \mathbf{U}_1), (\mathbf{U}_2^*, \mathbf{U}_2), \ldots, (\mathbf{U}_t^*, \mathbf{U}_t), \ldots, (\mathbf{U}_T^*, \mathbf{U}_T)$, where at each t one may (or may not) condition on history and/or on data realisations preceding t. However, such a definition complicates notation considerably so I only provide the definition for a generic t only, $(\mathbf{U}_t^*, \mathbf{U}_t)$, since the extension to $t = 1, 2, \ldots, T$ is straightforward. Now, recall the definition of a measurable variable.

> **Definition 5. Measurable variable.** Let $(\Omega, \mathcal{F})$ and $(\Omega^*, \mathcal{G}^*)$ denote two measurable spaces, that is, $\mathcal{F}$ and $\mathcal{G}^*$ are σ-fields on Ω and Ω^*, respectively, and denote the elements of $\mathcal{F}$ and $\mathcal{G}^*$ for F and G^*, respectively. A function $f : \Omega \longrightarrow \Omega^*$ is said to be $\mathcal{F}$-measurable if for all $G^* \in \mathcal{G}^*$ we have $\{\omega : f(\omega) \in G^*\} \in \mathcal{F}$.

In the case where Ω^* is Euclidean space then f is a random vector. For notational convenience I will use the symbolism $f : (\Omega, \mathcal{F}) \longrightarrow (\Omega^*, \mathcal{G}^*)$ to mean that f is a $\mathcal{F}$-measurable function from Ω to Ω^*, with $\mathcal{F}$ and $\mathcal{G}^*$ being the associated σ-fields. Now, consider the two measurable variables

$$\mathbf{U}_t^* : (\Omega, \mathcal{F}) \longrightarrow (\mathbf{X}_t^*, \mathcal{G}_t^*) \text{ and } \mathbf{U}_t : (\Omega, \mathcal{F}) \longrightarrow (\mathbf{X}_t, \mathcal{G}_t),$$

where $\mathbf{X}_t^* = X_{t1}^* \times X_{t2}^* \times \cdots \times X_{tI(t)}^*$ and $\mathbf{X}_t = X_{t1} \times X_{t2} \times \cdots \times X_{tJ(t)}$, and think of the first as the theory variable and the second as the data variable. The elements of $\mathcal{F}$, $\mathcal{G}_t^*$ and $\mathcal{G}_t$ will be referred to as worldly events at t, theory events at t and data events at t, respectively. Measurement validity of the data event $G_t \in \mathcal{G}_t$ with respect to the theoretical event $G_t^* \in \mathcal{G}_t^*$ can now be defined in terms of equality between the worldly events $\{\omega : \mathbf{U}_t^*(\omega) \in G_t^*\} \in \mathcal{F}$ and $\{\omega : \mathbf{U}_t(\omega) \in G_t\} \in \mathcal{F}$. In words, to what extent the set of possible worlds associated with a certain data realisation equals the set of worlds associated with the theory event it purports to measure. Generalised the idea can be summarised in the following definition.

Definition 6. Measurement validity of data events. A data event $G_t \in \mathcal{G}_t$ is said to be:

a) *measurement valid* with respect to a theory event $G_t^* \in \mathcal{G}_t^*$ if $\{\omega : \mathbf{U}_t(\omega) \in G_t\} = \{\omega : \mathbf{U}_t^*(\omega) \in G_t^*\}$.

b) *measurement invalid* with respect to a theory event $G_t^* \in \mathcal{G}_t^*$ if $\{\omega : \mathbf{U}_t(\omega) \in G_t\} \cap \{\omega : \mathbf{U}_t^*(\omega) \in G_t^*\} = \emptyset$.

c) *partially measurement valid* with respect to a theory event $G_t^* \in \mathcal{G}_t^*$ if $\{\omega : \mathbf{U}_t(\omega) \in G_t\} \neq \{\omega : \mathbf{U}_t^*(\omega) \in G_t^*\}$ and $\{\omega : \mathbf{U}_t(\omega) \in G_t\} \cap \{\omega : \mathbf{U}_t^*(\omega) \in G_t^*\} \neq \emptyset$.

For convenience we may say that a data event is measurement valid, invalid or partially valid, respectively, since it is implicitly understood that the validity is with respect to a certain theory event. The extensions to theoretical variables is more or less straightforward, but for convenience I only provide the definition for measurement validity.

Definition 7. Measurement validity of a data variable. A data variable $\mathbf{U}_t : (\Omega, \mathcal{F}) \longrightarrow (\mathbf{X}_t, \mathcal{G}_t)$ is said to be measurement valid if each $G_t \in \mathcal{G}_t$ is measurement valid.

Implicitly the definition thus assumes there is a theory variable $\mathbf{U}_t^* : (\Omega, \mathcal{F}) \longrightarrow (\mathbf{X}_t^*, \mathcal{G}_t^*)$ defined on the probability space $(\Omega, \mathcal{F}, P)$. Finally, a definition of almost sure measurement validity can be formulated.

Definition 8. Almost sure measurement validity of a data variable. Consider a data variable $\mathbf{U}_t : (\Omega, \mathcal{F}) \longrightarrow (\mathbf{X}_t, \mathcal{G}_t)$ and denote the set containing measurement valid data events for $\mathcal{G}_t^1 = \{G_t \in \mathcal{G}_t : G_t \text{ is measurement valid}\} \subset \mathcal{G}_t$. If $P[\bigcup_{i=1}^{\infty} G_t^1(i)] = 1$ where $\{G_t^1(1), G_t^1(2), \ldots, G_t^1(i), \ldots\} = \mathcal{G}_t^1$, then $\mathbf{U}_t$ said to be *measurement valid almost surely.*

The modified framework is summarised in table 6.2.

6.5 A history based probabilistic definition of indeterministic causality

Discussions over what the appropriate definition of causality is for econometrics enjoys a reasonably long history, for overviews and references see amongst others Geweke (1984), Aigner and Zellner (1988), and Bauwens et al. (forthcoming). Most of the suggested definitions have put more weight on empirical implementability rather than philosophical justification, which is understandable given econometrics' nature. In this sense the definition outlined here distinguishes itself by explicitly giving more weight to philosophical considerations rather than empirical implementation. Indeed, I believe that the principal use of the notion of causality proposed here is conceptual analysis rather than empirical analysis. The main characteristic of the definition is that it conceives causality as having two aspects, historical possibility and causal efficiency, and the section proceeds in four steps. In the first subsection the idea of historical possibility is introduced and discussed, and in the second causal efficiency. Subsection 6.5.3 brings out the most important similarities and differences between the proposed definition and David Lewis' ideas, whereas the final subsection relates the ideas in this section to a common definition of causality in econometrics, namely mathematical expectation conditional on an information set.

6.5.1 Historical possibility

The first aspect of causality is in a sense obvious. How can an event C be considered as a cause of another event E if the second event is not even possible given the first? The approach to possibility pursued here is that of historical possibility, that is, the idea that what is possible tomorrow depends crucially on where we stand today. In other words, the course of history up to t determines what is possible at and after t. Before we can define this idea formally we need a definition of history.

> **Definition 9. History up to t.** Let ω_t be a state-of-affairs process up to t. The event $H_t = \{\omega : \omega_t \subsetneq \omega\} \in \mathcal{F}$ is said to be a history up to t.

In words H_t is the set of all possible worlds that contain the state-of-affairs process ω_t and intuitively H_t is exactly what its name suggests, namely history up to t. Now, a possible or historically possible event is defined as follows.

> **Definition 10. A historically possible event.** Let $H_{t_1}, E_{t_2} \in \mathcal{F}$ where $t_1 \leq t_2$ and where H_{t_1} is a history up to t. E_{t_2} is said to be a possible event with respect to the history H_{t_1} if $H_{t_1} \cap E_{t_2} \neq \emptyset$.

In words, the event E_{t_2} at t_2 is said to be historically possible or possible for short if at least one of its worlds is contained in history. Similarly, an event is impossible if $E_{t_2} \cap H_{t_1} = \emptyset$, since H_t by construction contains the set of all possible worlds containing the course of history up to and including t. A consequence of definition 10 is that situations where $E_{t_2} \cap H_{t_1} \neq \emptyset$ and $P(E_{t_2} \cap H_{t_1}) = 0$, that is, that E_{t_2} is possible but probabilistically impossible, are not excluded from the outset. Situations where the effect precedes its cause are on the other hand excluded from the outset by the condition $t_1 \leq t_2$. Another characteristic of the definition is that it allows for events being causal for some t but not necessarily at all t. In particular, the definition allows for socalled "single case" causality with "many case" causality being obtained as a probabilistic reduction. Finally, the definition resolves a problem discussed by Salmon (1993b, 1993a). In his view the definitions of probabilistic causality put forward by Reichenbach (1956), Good (1961, 1962 and 1963) and Suppes (1970) all suffer from the fact "that they attempt to carry out the construction of causal relations on the basis of probabilistic relations among discrete events…" (1993b, p. 151). In other words, they fail to take into account the continuous processes that connect events. Salmon (1993a) himself proposed a solution that takes "processes rather than events as basic entities" (same place, p. 155). The current approach follows in the same vein and thus constitutes an alternative to Salmon's approach. Specifically the current approach takes continuous states-of-affairs processes (that is, worlds) as basic entities with the consequence that discrete and continuous accounts—indeed, even "interval" accounts—of causality are reconciled in a neat manner.

6.5.2 Causal efficiency

Defining possibility in this way means conditional probability suggests itself as a measure of causal efficiency. Heuristically a cause is said to be more efficient than another if the first is more likely to bring about the event in question, and formally we may define this as follows:

Definition 11. Causal efficiency. Let $H_{t_1}^a$, $H_{t_1}^b \in \mathcal{F}$ denote two different histories, that is, $\omega_{t_1}^a \neq \omega_{t_1}^b$, and consider the event $E_{t_2} \in \mathcal{F}$ where $t_1 \leq t_2$. Further let $P(E_{t_2} \cap H_{t_1}^a)$, $P(E_{t_2} \cap H_{t_1}^b) \neq \emptyset$ and $P(H_{t_1}^a)$, $P(H_{t_1}^b) > 0$:

a) If $P(E_{t_2}|H_{t_1}^a) > P(E_{t_2}|H_{t_1}^b)$, then $H_{t_1}^a$ is said to be causally more efficient than $H_{t_1}^b$ in bringing about E_{t_2}.

b) If $P(E_{t_2}|H_{t_1}^a) = P(E_{t_2}|H_{t_1}^b)$, then $H_{t_1}^a$ and $H_{t_1}^b$ are said to be causally equally efficient in bringing about E_{t_2}.

As an example, let $H_{t_1}^a$ and $H_{t_1}^b$ denote two different policy choices, say, increasing the interest rate with 0.5%-point and no-change, respectively, and let E_{t_2} denote the desired policy objective, say, a yearly inflation of 2.5% two years into the future. The conditions $E_{t_2} \cap H_{t_1}^a \neq \emptyset$ and $E_{t_2} \cap H_{t_1}^b \neq \emptyset$ essentially state that the policy choices in question both are capable of bringing about the desired objective E_{t_2}. So if, say, $P(E_{t_2}|H_{t_1}^a) > P(E_{t_2}|H_{t_1}^b)$, then alternative a is more likely to attain E_{t_2} than alternative b.

How does all this relate to the more common idea of an event C_{t_1} being the cause of an effect E_{t_2}, most often expressed in terms of $P(E_{t_2}|C_{t_1})$? The answer lies in the structure of the underlying outcome space. Recall that each element in the outcome space is devised as a continuous time process of states-of-affairs from 0 and onwards. This means $P(E_{t_2}|C_{t_1})$ can be interpreted as the probability of the event E_{t_2} given the event C_{t_1}, *and* given history up to but not including $t = 0$. This is why bounding worlds backwards produces a particularly interesting interpretation of conditional probabilities. Moreover, it is not necessary for worlds to be unbounded forward for this interpretation to obtain. Worlds may just as well be devised as finite non-stochastic continuous time processes starting at $t = 0$ and ending when $t = T$. Similarly, the probability $P(E)$ where E is an arbitrary event in the event-set can be interpreted as the probability of E coming about *conditional* on history up to but not including $t = 0$.

6.5.3 David Lewis' ideas compared

The account of causality outlined here is similar in so many ways to Lewis' account that one may ask where they actually differ. They do differ in many ways but the most important are three. The first was alluded to in subsection 6.3.1 and concerns the existence of possible worlds. Whereas Lewis believed other worlds exist objectively and independent of thought, I believe they originate in our imagination. Second, Lewis aims to provide a framework that "can serve alike under indeterminism or determinism" (1986d, p. 179). The account outlined here on the other hand has been formulated with indeterminism in mind, and I am not ready to say yet how related they are in the case when the outcome space only contains a single world, which can be interpreted as a version of determinism. Third, Lewis' account "is in terms of counterfactual conditionals about probability; not in terms of conditional probabilities" (same place, p. 178). Here, conditional probability is one of two aspects of causality and (formal) counterfactual conditionals play no role.

With respect to similarities the most important is how close my ideas regarding causality are to Lewis' (1986a) view on causal explanation—in particular sections I and II. Another similarity concerns the interpretation of probability. Events, that is, elements of $\mathcal{F}$, are sets of possible worlds, and the conditional probability (say) $P(E|C)$ is the objective propensity of the event C to bring about the event E. In other words, conditional probability may apply to single instances of cases. The propensity (probability) is interpreted in the objective sense as opposed to the subjective, but this should not be interpreted as a critique against subjective accounts of probability. Indeed, in the words of Lewis:

> "We subjectivists conceive of probability as the measure of reasonable partial belief. But we need not make war against other conceptions of probability, declaring that where subjective credence leaves off, there nonsense begins. Along with subjective credence we should believe also in objective chance. The practice and analysis of science require both concepts"—(1986e, p. 83).[42]

6.5.4 Conditional expectations re-interpreted

It is common in econometrics to model the impact of one set of variables on another by means of conditional expectations. In its general form such conditional expectations may be denoted $E(X_t|\mathcal{I} = I)$, where X_t is the random variable or variables in question, $\mathcal{I}$ is a σ-field contained in $\mathcal{F}$ and $I \in \mathcal{I}$.[43] A common example are socalled "filtrations". If $\mathbf{X}_t = \{X_0, X_1, \ldots, X_t\}$ denotes as sequence of the random variable up to and including t, then the sequence of σ-fields generated by $\mathbf{X}_0, \mathbf{X}_1, \ldots, \mathbf{X}_t$, commonly denoted $\mathcal{I}_0, \mathcal{I}_1, \ldots, \mathcal{I}_t$ and called "information-sets", is the filtration of $\mathbf{X}_t$ if it is the case that $\mathcal{I}_{t-1} \subset \mathcal{I}_t$ and $\mathcal{I}_t \subset \mathcal{F}$ for all t. As a consequence, the conditional expectation $E(X_t|\mathcal{I}_{t-1} = I_{t-1})$ is often referred to as the conditional expectation of X_t on all the information available up to t, and sometimes even *the* true conditional expectation. This is a peculiar practice if the information-set is interpreted to contain what its name suggest, namely information.

The current structure of the outcome space permits us to distinguish between two distinct but compatible and complementary ideas, history and information. Let $I_t \in \mathcal{F}$ denote the event that an entity (a person, a group of persons, or whatever) possesses or uses a certain collection of information at t, and let $H_t \in \mathcal{F}$ denote history up to and including t as defined above. Two useful distinctions can be made. Between correct and incorrect information of the past on the one hand, and between complete and incomplete information of the past on the other. More formally, sets of correct and incorrect information are characterised by $I_t \cap H_t \neq \emptyset$ and $I_t \cap H_t = \emptyset$, respectively, and sets of complete and incomplete correct information by $I_t \cap H_t = H_t$ and $I_t \cap H_t \subsetneq H_t$, respectively. This gives three cases. The first case is when the information in the information-set is both correct and complete, and is of course entirely unrealistic. Formally, $I_t = H_t$. The second case is when I_t contains some correct information, but not all the correct information that exists. Formally, $I_t \cap H_t \neq \emptyset$ and $I_t \subsetneq H_t$. Finally, the third case is when I_t contains incorrect information only. Formally, $I_t \neq \emptyset$ and $I_t \cap H_t = \emptyset$.

The point I am driving at is intuitively obvious, namely that in practical econometrics our information is both incomplete and possibly incorrect, and that we use this suboptimal information in estimating conditional expectations. An attempt to formalise this idea could be the following. The "correct" or true expectation conditional on history is given by $E(X_t|\mathcal{F} = H_t)$, whereas what the econometrician in practice estimates is $E(X_t|\mathcal{I}_t = I_t)$ where I_t is an incomplete and possibly partially incorrect information set. Denoting this estimate by $\hat{E}(X_t|\mathcal{I}_t = I_t)$, we may say that one of the key concerns of econometrics is that of efficiently choosing and making use of information such that $\hat{E}(X_t|\mathcal{I}_t = I_t)$ is as close to $E(X_t|\mathcal{F} = H_t)$ as possible.

6.6 Conclusions

In this chapter I have argued that the underlying outcome space in Hendry's (1995) reduction framework can usefully be interpreted as consisting of possible worlds made up of indeterministic and historically inherited particulars. First, although the human world is changing all the time in indeterministic ways, the interpretation means the relationship between random variables can still be analysed on the same underlying probability space, since all the change takes place at the level of worlds. Moreover, as an additional interpretation of practical interest the formulation of theoretical variables can be seen as the "perspective" from which an issue is studied. Second, a probabilistic analysis on the same underlying probability space on the relation between theory variables and data variables is enabled. Third, a history based probabilistic definition of indeterministic causality that nests discrete, continuous and "interval" versions is proposed. Fourth, mathematical expectation conditional on an information set is re-interpreted. Finally and more generally, a bridge between econometric (/probabilistic) reduction analysis and metaphysics is provided.

This suggest several possible lines of research, both within the theory and practice of econometrics, of which only the one with greatest relevance to the rest of the thesis will be outlined here. The notion of weak stationarity plays a central role in dynamic econometrics and is cast in terms of marginal entities. To recall, a series $\{Y_t\}_{t=1}^{T}$ is defined as weakly stationary if $E(Y_t) = \mu$ for all t, and given any t we have that $E(Y_t - \mu)(Y_j - \mu) = \sigma_j$ for all j (that is, σ_j does not depend on t).

However, as argued in this chapter, everything in the human world is conditional and so it seems natural that this is taken into account in the definition of weak stationarity. Specifically, a straightforward extension is to reformulate the definition of weak stationarity as conditional on a collection of events. Indeed, I would be very surprised if such a definition has not already been proposed. Denote a collection of such conditioning events as $\mathcal{I} = \{I_1, \ldots, I_N\}$, where $\mathcal{I} \subset \mathcal{F}$ but where $\mathcal{I}$ is not necessarily a σ-field. An example of a "conditional" definition of weak stationarity with respect to $\mathcal{I}$ would then be that, for all $I \in \mathcal{I}$, $E(Y_t|\mathcal{F} = I) = \mu$ for all t, and given any t we have that $E[(Y_t - \mu)(Y_j - \mu)|\mathcal{F} = I] = \sigma_j$ for all j. Breaks or change in stationarity could then be sought explained in terms of a change from one conditioning event $I_1 \in \mathcal{I}$ to another $I_2 \in \mathcal{I}$. For example, a "break" from μ to μ' could be sought explained in terms of a change in circumstances I_1, say, the existence of the Bretton Woods Order of international finance, to I_2, say, the breakdown of the Bretton Woods Order. This would provide a conceptual solution and a unifying framework to study such ideas as co-breaking (Hendry and Massmann 2005), changing "unconditional" volatilities (Engle and Gonzalo 2005; how can changes in unconditional volatiliy be explained unless the change is due to a change in economic or other circumstances?) and "common features" (Engle and Kozicki 1993).

Table 6.1: Starting point, action and the resulting reduction in Hendry's framework associated with the first stage of reduction.

Reduction no.	Starting point and resulting reduction	Action
	The economic mechanism under study: The theory variables $\mathbf{U}^* = (\mathbf{U}_1^*, ..., \mathbf{U}_T^*)$ defined on the probability space $(\Omega, \mathcal{F}, P)$	
		Data collection and recording of $\mathbf{U}_t \in \mathbf{U}$, that is, the process of trying to measure the $\mathbf{U}_t^* \in \mathbf{U}^*$ variables
1.	The data generation process (DGP): The data set $\mathbf{U} = (\mathbf{U}_1, ..., \mathbf{U}_T)$ defined on the transformed probability space $(\Omega', \mathcal{F}', P')$	

Table 6.2: Starting points, actions and resulting reductions associated with Hendry's framework when the outcome set consists of possible worlds made up of indeterministic and historically inherited particulars.

Reduction no.	Starting points and resulting reductions	Action
	A probability space $(\Omega, \mathcal{F}, P)$, where the outcome-space Ω consists of possible worlds made up of indeterministic and historically inherited particulars	
		The delineation and definition of a set of theory variables $\mathbf{U}^* = (\mathbf{U}_1^*, \ldots, \mathbf{U}_T^*)$
1.	The economic mechanism under study: The theory variables $\mathbf{U}^* = (\mathbf{U}_1^*, ..., \mathbf{U}_T^*)$ defined on the probability space $(\Omega, \mathcal{F}, P)$	
		Data collection and recording of $\mathbf{U}_t \in \mathbf{U}$, that is, the process of trying to measure the $\mathbf{U}_t^* \in \mathbf{U}^*$ variables
2.	The data generation process (DGP): A data realisation $\mathbf{u} = (\mathbf{u}_1, \ldots, \mathbf{u}_T)$ of the data variables $\mathbf{U} = (\mathbf{U}_1, \ldots, \mathbf{U}_T)$ defined on the probability space $(\Omega, \mathcal{F}, P)$	

Chapter 7

Conclusions

This thesis has sought contributing to the study and modelling of exchange rate volatility in several ways. Here, the results of the thesis are summarised and suggestions for further research are proposed.

7.1 Summary of thesis

In chapter 2 a distinction was made between period and within-period variability, a distinction of interest when studying variability across different exchange rate regimes. We also proposed the exponential model of variability (EMOV) as a particularly convenient framework for explanatory exchange rate volatility modelling.

Chapter 3 made full fuse of these ideas in studying the impact of market activity on exchange rate variability in the case of Norway. The main findings of this study are that the impact of short-term change in market activity, as measured by relative week-to-week changes in quoting frequency, is positive and statistically significant for both our definitions of variability, and that the impact is relatively stable across three different exchange rate regimes. One might have expected that the effect would increase with a shift in regime from exchange rate stabilisation to partial inflation targeting, and then to full inflation targeting, since the Norwegian central bank actively sought to stabilise the exchange rate before the full inflation targeting regime. In our data however there are no clear breaks, shifts upwards nor trends following the points of regime change. Indeed, the instability there is suggest the opposite, namely that the impact was higher in the first regime when the Norwegian central bank actively sought to stabilise the exchange rate. With respect to the hypothesis that changes in long-term market activity—as measured by the average level of quoting frequency in the previous six weeks—increases variability, our results support to some extent that this is the case for weekly range variability. However, our results do not support the hypothesis in the case of weekly period variability. We also find some evidence that impact of long-term market activity on range variability depends on exchange rate regime. In particular, that the impact is higher in the first regime and lower (and possibly insignificant) in the two subsequent regimes. Finally, our results do not suggest that the persistence in variability can be explained by persistence in the level of volume.

Chapter 4 undertook an out-of-sample evaluation of general to specific (GETS) modelling of exchange rate volatility. The GETS methodology has proved particularly useful in explanatory econometric modelling of many economic time series, but can be difficult to implement in computationally complex models—as is often the case for financial volatility models—when many variables are involved—as is typically the case in GETS-modelling. We proposed and showed that computational complexity can be sidestepped by working with the EMOV, and our out-of-sample forecast evaluation suggests that GETS-derived models are particularly useful in conditional forecasting.

Chapter 5 studied the relation between exchange rate variability, market activity and heterogeneity using spot NOK/EUR transaction data from banks within Norway's regulatory borders. Whereas an increase in short term global interbank market activity (as measured by the relative increase in quoting frequency) increases range variability, our results do not support the hypothesis that increases in short term Norwegian market activity (as measured by the relative increase in spot NOK/EUR trading volume) has a statistically significant impact on neither period nor range variability. Moreover, we do not find support for the hypothesis that some groups of banks, for example big banks, have an impact on variability through their short term market activity. With respect to the impact of long term market activity, however, our results do suggest that Norwegian NOK/EUR trading has an impact. In particular, we find some support of the hypothesis that increased long term activity by banks in Norway increases range variability through their long term spot NOK/EUR trading, and that groups of similarly sized banks have different impacts. The group of small banks' long term market activity has a negative impact on period variability, whereas the group of large banks' long term market activity has a positive impact on range variability.

Chapter 6 argued that the underlying outcome space in Hendry's (1995) reduction framework can usefully be interpreted as consisting of possible worlds made up of indeterministic and historically inherited particulars. First, although the human world is changing all the time in indeterministic ways, the interpretation means the relationship between random variables can still be analysed on the same underlying probability space, since all the change takes place at the level of worlds.

Moreover, as an additional interpretation of practical interest the formulation of theoretical variables can be seen as the "perspective" from which an issue is studied. Second, a probabilistic analysis on the same underlying probability space on the relation between theory variables and data variables is enabled. Third, a history based probabilistic definition of indeterministic causality that nests discrete, continuous and "interval" versions is proposed. Fourth, mathematical expectation conditional on an information set is re-interpreted. Finally and more generally, a bridge between econometric (/probabilistic) reduction analysis and metaphysics is provided.

7.2 Suggestions for further research

The results of the thesis suggests many areas and questions for further investigation. In particular, Norwegian exchange rate variability, but also the variability of other exchange rates, exhibit what seems to be a structural break, that is, a shift upwards, around the end of 1996 and/or beginning of 1997. The exact nature and timing of this event is not well understood. According to van Dijk et al. (2005) several non-Euro exchange rates against the USD experienced a break in unconditional volatilities (their study was conducted by means of a dynamic conditional correlation framework), and the NOK/USD exchange rate is the one that exhibits the largest shift upwards (50%). They attribute the break to a European Council meeting in December 1996 in which a decision regarding the EMU was taken, and that this was pronounced in the Norwegian case because of a change in the intervention policy of the Norwegian Central Bank. According to Bjønnes et al. (2005) on the other hand the events at the end of 1996/beginning of 1997 were due to a speculative attack by foreign speculants. These two explanations are not necessarily incompatible, but to understand the exact nature, timing and reasons for the shift upwards in variability around the end of 1996/beginning of 1997 is therefore a further area of research.

In parts of this thesis we have used the Norwegian policy interest rate rather than Norwegian money market interest rates, because the impact of the variables constructed with the former is much more stable. Although the policy interest rate affect interest rate bearing securities, the actors in foreign exchange markets are mainly concerned with the money market interest rates. So it would be desirable to use money market interest rates instead. Further understanding on the relation between policy interest rate changes, money market interest rates and exchange rate variability is therefore necessitated.

Another area for further research that our results suggest should be pursued is the heterogeneous impact Norwegian long term market activity has on variability. Our results suggest a negative impact on period variability from the long term market activity of small banks, and a positive impact on range variability from the long term market activity of large banks. It is not evident why this is the case, so further investigation—possibly approaching the issue in different ways—could shed further light on the issue.

The study that undertook an out-of-sample evaluation of GETS-modelling in chapter 4 shows promise, but the generality of the results must be established. To what extent is GETS-modelling of financial volatility useful on higher/lower frequencies than the weekly? On other exchange rates and for other financial assets? Moreover, contrary to McAleer's (2005) assertion, automated GETS-modelling of financial volatility can be readily implemented and should be investigated more fully. A drawback with the EMOV is that the conditional mean is restricted to zero, which means that predictability in the direction of exchange rate changes can not be exploited. One interesting line for further research is therefore to make use of two-step OLS estimators of (say) ARCH-like models so that on the one hand all the numerical issues and problems associated with GETS modelling of volatility are avoided, and on the other that conditional means also can be modelled.

Finally, the results of chapter 6 suggest numerous possible lines of research, both within the theory and practice of econometrics, of which only the one with greatest relevance to the thesis content is outlined here. The notion of weak stationarity plays a central role in dynamic econometrics and is cast in terms of marginal entities. However, according to the commonplace social ontology that underpinned chapter 6 everything in the human world is conditional, and so it seems natural that this is taken into account in the definition of weak stationarity. Specifically, a straightforward extension is to reformulate the definition of weak stationarity as conditional on a collection of events. Breaks or changes in stationarity could then be sought explained in terms of a change from one conditioning event to another. This would provide a conceptual solution and a unifying framework to study such ideas as co-breaking (Hendry and Massmann 2005), changing "unconditional" volatilities (Engle and Gonzalo 2005; how can changes in unconditional volatiliy be explained unless the change is due to a change in economic or other circumstances?) and "common features" (Engle and Kozicki 1993).

Notes

[1]See for example Karl Polanyi's (2002 [1944]) classic *The Great Transformation* for an argument along these lines. 730 delegates from 45 nations attended the United Nations Monetary and Financial Conference, commonly known as the Bretton Woods conference, from 1 to 22 July 1944 (Wikipaedia 2006)

[2]The term "supervene" is borrowed from metaphysics. Here, it loosely means that one event cannot "cause" another unless there are processes that either connect them "sequentially" or "simultaneously".

[3]This section draws to a notable extent on Mestad (2002) and the relevant web-pages of the Central Bank of Norway (Norges Bank): http://www.norges-bank.no.

[4]This was made more precise in another Government resolution dated 6 May 1994.

[5]As an anecdote, Svein Gjerdrem, the Governor that assumed the position in January 1999 and which still occupies the post, was a key representative of the Ministry of Finance in the letter exchange of May 1998.

[6]According to a recent estimate (Meyer and Skjelvik 2006, p. 36) spot NOK/EUR trading accounts for 71% of total spot NOK-volume during the period October 2005 to January 2006, whereas spot NOK/USD trading accounts for only 14% of total spot volume. The estimate is based on daily data collected by Norges Bank, and comprises all NOK-trading in Norway and a substantial part of NOK-trading outside Norwegian regulatory borders.

[7]This series is denoted S_t in the data appendix.

[8]This series is denoted r_t in the data appendix.

[9]We make no claim to originality in suggesting this specification. Indeed, both Epps and Epps (1976, p. 311) and Tauchen and Pitts (1983, p. 494) estimated linear versions in their studies.

[10]It should be noted that since conditioning occurs within a different statistical setup the $\{\mathcal{I}_t\}$ may differ from above in the EMOV setup. For example, in the SV case $\mathcal{I}_t$ may contain past values of the stochastic term in the volatility specification.

[11]No generality is lost by only considering the ARCH family since the same type of argument applies with respect to the SV family under standard assumptions.

[12]Jorion noted that, on daily data, the factor is typically 700 to 1. In our case the median of the fitted values of σ_t^2 is between 500 and 600 times greater than the median of the fitted values of μ_t^2 in the GARCH(1,1) and EGARCH(1,1) specifications of chapter 4.

[13]An alternative measure is Brousseau and Scacciavillani (1999), see Bernhardsen and Røisland (2000) for an application of this measure in a Norwegian context.

[14]The awareness of financial market linkages and spill-overs in crisis periods is old, see for example Kindleberger (1993), Kindleberger (2005) and **?**, chapter X for historical studies. For more recent references, see amongst other Hamao et al. (1990), King and Wadhwani (1990), Lin et al. (1994), and Baele (2005). There is not much literature on the impact of stock market variability on exchange rate variability, but for studies of the opposite effect see Bodart and Reding (1999), and **?**.

[15]Prior to 1999 central bank interest rates were very stable, at least from late 1993 until late 1996, and it was less clear to the market what role the interest rate actually had.

[16]Another possibility could be that the oilprice variable is insignificant because its impact works through the stock market variables. However, our data do not suggest this since excluding the stock market variables from the analysis still renders the oilprice variable insignificant.

[17]GETS modelling is also sometimes referred to as the "LSE methodology", after the institution in which the methodology to a large extent originated in, and sometimes even "British econometrics", see Gilbert (1989), Mizon (1995) and Hendry (2003).

[18]The term "congruent" is borrowed from geometry: By "analogy with one triangle which matches another in all respects, the model matches the evidence in all measured respects." (Hendry 1995, p. 365)

[19]Other expositions of the GETS methodology and its foundations are Hendry and Richard (1990), Gilbert (1990), Mizon (1995) and Jansen (2002).

[20]See Hendry (1995, pp. 362-367) and Mizon (1995) for further discussion.

[21]See Campos et al. (2005) for a more complete discussion.

[22]The single-path specification encompasses the PcGets specification in the sense that v^{w}_{t-2} is significant. In the settings for PcGets we use comparable significance levels and search strategies to the single-path simplification, and we have not been able to verify why PcGets produces a different specification. A possible reason is that PcGets employs information criteria in the case where multiple terminal models result.

[23]Our sample of 573 observations is considerably larger than those investigated by Lovell (1983), Hoover and Perez (1999) and Hendry and Krolzig (1999), the sequence

of studies that resulted in PcGets. Whereas Lovell (1983) used only 23 observations, the other two studies employed a maximum of 140 observations.

[24]This sample was chosen because the volatility of r_t looks relatively stable over this period. Specifically, the values of $\bar{x}^w$, $\bar{u}^w$ and $\bar{ir}^{emu}$ are 0.633, 0.412 and 0.006.

[25]Specifically, the values of $\dot{x}^w$, $\dot{u}^w$ and $\dot{ir}^{emu}$ are 1.090, 0.727 and 0.001.

[26]If t denotes the sample size, k the number of parameters in $\mathbf{b}$ and M the observation at which recursive estimation starts, then for $t = M, \ldots, T$ the 1-step, breakpoint and forecast tests are computed in PcGive as $F(1, t-k-1)$, $F(T-t+1, t-k-1)$ and $F(t-M+1, M-k-1)$, respectively, see Hendry and Doornik (2001).

[27]The number 473 is due to the fact that the recursive estimation was initialised at observation number 100.

[28]To be more precise, the parameter values are those suggested by the 1995 version of RiskMetrics for daily data, which then was part of the merchant bank J.P. Morgan. RiskMetrics is now an independent company and two versions of RiskMetrics have superseded the 1995 May edition, see http://www.riskmetrics.com/techdoc.html. Note also that the parameter values are obtained with a definition of volatility that differs slightly from the one employed here.

[29]Patton (2005) has recently argued in favour of MSE in volatility forecast comparison. It should be noted however that his argument applies (under certain assumptions) when the problem to be solved is to choose, in our notation, an $\hat{V}_t$ such that expected $L(\sigma_t^2, \hat{V}_t)$ is minimised, where L is a loss function. As we argued in the previous subsection 4.4.1, however, the problem to be solved is to choose an $\hat{V}_t$ such that expected $L(V_t, \hat{V}_t)$ is minimised. This is a qualitatively important difference and it is not clear that Patton's conclusions hold when the problem is formulated in our way. Nevertheless, MSE is one of the most commonly applied statistic and there are few disadvantages associated with it, so we follow suit in using it.

[30]Several other approaches to out-of-sample forecast comparison have been proposed. One consists of adding other ingredients to the evaluation scheme, see for example West et al. (1993) where the expected utility of a risk averse investor serves as the ranking criterion. Similarly, Engle et al. (1993) provide a methodology in which the profitability of a certain trading strategy ranks the forecasts. Yet another approach takes densities as the object of interest, see Diebold et al. (1998), whereas Lopez (2001) has proposed a framework that provides probability forecasts of the event of interest.

[31]The underlying data collection methodology and definitions were changed again in January 2005. For this reason we do not use the data after 7 January 2005, see appendix 3 for more details.

[32]These expressions were obtained through simplification of $b_2v_{t_1}^{se}+b_3v_{t_1}^{hl}+b_4v_{t_2-1}^{se}+b_5v_{t_2-1}^{hl}$ in the period variability case, and $b_3v_{t_1}^{hl}+b_5v_{t_2-1}^{hl}+b_7v_{t_1-1}^{hl}+b_9v_{t_2-2}^{hl}$ in the range variability case, respectively. In the range case the specific form of persistence is required for residuals to be serially uncorrelated.

[33]Later Haavelmo quotes Pareto: "Il n'y a pas de proposition qu'on ne puisse certifier vraie sous certaines conditions, à determiner" (same place, p. 1). This might be taken as an indication on why he was agnostic about the existence or "truthfulness" of such a joint distribution. Differently put, agreeing with Pareto that truth is a complex issue and highly dependent one tests for it, Haavelmo chose to defer the topic rather than engaging into a detailed and possibly futile discussion.

[34]Chapter 9 in Hendry (1995) is a revised version of Cook and Hendry (1994), which is based on Hendry and Richard (1990).

[35]Throughout the chapter I will employ philosophical terms which I explain only briefly. Readers interested in fuller explanations or further reading are referred to (say) Honderich (1995) and Craig (2000).

[36]The "important point", he says, "is that empirical relationships must arise from these reductions of the DGP" (same place, p. 345).

[37]Leibniz was religious and originally he used the idea to argue that the world is perfect because among all the possible worlds God must have chosen the most perfect one, an idea that was ridiculed by Voltaire in his play *Candide* (Crane 1995). In today's philosophical usage however the term usually carries no religious connotation.

[38]The whole book is a defence of this thesis but see in particular pp. vii-ix and pp. 133-135. For a brief and amusing summary of other philosophers' reactions to Lewis's thesis, see Hawthorn (1995, footnote 24 pp. 23-24).

[39]For further philosophical issues and references regarding the idea of a possible world useful starting points are Forbes (1995) and Moravcsik (1995). For an alternative but related use of the idea of a possible world by an economist, see Kluve (2004).

[40]An entry on "free will" is contained in virtually any philosophy or metaphysics dictionary, see for example Honderich (1995) or Kim and Sosa (1995), and usually contain further reading. A very accessible introduction to these issues, which is based

on the author's BBC lectures, is Searle (1991). Useful introductions to the philosophy of mind are Kim (1996) and Heil (1998), the second being more advanced than the first. A good text on the relation between mind and recent biological currents is Ruse (1988). Texts that consider themselves to specifically address issues of social ontology are Ruben (1985) and Pettit (1993). A useful introduction to metaphysics as it is often conceived, a form of category theory, is Loux (1998).

[41]A further interpretation of the thesis that the human world is made up of particulars is that, literally, people differ from each other: No two persons are equal in all respects at any point in time. In the current context, however, we only need the first interpretation.

[42]Essentially this essay is Lewis' account of the relation between subjective and objective versions of probability. See also Lewis (1986c).

[43]Little is lost by restricting our attention to conditional expectations, since the conditional probability of an event A given and event B is obtained by taking the conditional expectation of the indicator function of A. For instance, $E(I_A|\mathcal{F} = B) = P(A|B)$.

NOTES

References

Aigner, D. J. and A. Zellner (1988). Causality. *Supplement to Journal of Econometrics 39*.

Aït-Sahalia, Y. (2006). Estimating Continuous-Time Models with Discretely Sampled Data. Invited lecture at the 2005 World Congress of the Econometric Society. To appear as chapter 9 in volume 3 of the conference proceedings, and currently available via http://cemmap.ifs.org.uk/.

Akram, F. (2004). Oil prices and exchange rates: Norwegian evidence. *Econometrics Journal 7*, 476–504.

Andersen, T. G. and T. Bollerslev (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review 39*, 885–905.

Andersen, T. G., T. Bollerslev, P. F. Christoffersen, and F. S. Diebold (2005). Volatility Forecasting. Center for Financial Studies Discussion Paper 2005/08, Frankfurt.

Andersen, T. G., T. Bollerslev, F. S. Diebold, and P. Labys (2001). The Distribution of Realized Exchange Rate Volatility. *Journal of the American Statistical Association 96*, 42–55. Correction published in 2003, volume 98, page 501.

Andersen, T. G., T. Bollerslev, and S. Lange (1999). Forecasting Financial Market Volatility: Sample Frequency vis-à-vis Forecast Horizon. *Journal of Empirical Finance 6*, 457–477.

Baele, L. (2005). Volatility spillover effects in european equity markets. *Journal of Financial and Quantitative Analysis 40*, 373–402.

Bauwens, L., W. Ben Omrane, and P. Giot (2005). News announcements, market activity and volatility in the euro/dollar foreign exchange market. *Journal of International Money and Finance*. Forthcoming.

Bauwens, L., H. P. Boswijk, and J.-P. Urbain (forthcoming). Causality and exogeneity in econometrics. *Journal of Econometrics*.

Bauwens, L., D. Rime, and G. Sucarrat (2006). Exchange Rate Volatility and the Mixture of Distribution Hypothesis. *Empirical Economics 30*, 889–911.

Bauwens, L. and G. Sucarrat (2006). General to Specific Modelling of Exchange Rate Volatility: A Forecast Evaluation. UCL - Louvain la Neuve, Belgium: CORE Discussion Paper 2006/21.

Baxter, M. W. and A. J. Rennie (1996). *Financial Calculus: An Introduction to Derivative Pricing*. Cambridge: Cambridge University Press.

Beltran, H. (2005). *Dynamics of liquidity in order book markets*. Louvain la Neuve, Belgium: Université catolique de Louvain.

Bernhardsen, T. and Ø. Røisland (2000). Factors that influence the exchange rate. *Economic Bulletin 4*, 143–151.

Bertsimas, D., L. Kogan, and A. Lo (2000). When is time continuous? *Journal of Financial Economics 55*, 173–204.

Bessembinder, H. and P. Seguin (1992). Futures-trading activity and stock price volatility. *Journal of Finance 47*, 2015–2034.

Bjønnes, G., S. Holden, D. Rime, and H. O. A. Solheim (2005). 'Large' vs. 'small' players: A closer look at the dynamics of speculative attacks. Working Paper 2005/13, Oslo: The Central Bank of Norway.

Bjønnes, G., D. Rime, and H. Solheim (2003). Volume and volatility in the FX market: Does it matter who you are? Working Paper 2003/3, Oslo: The Central Bank of Norway.

Bodart, V. and P. Reding (1999). Exchange rate regime, volatility and international correlations on bond and stock markets. *Journal of International Money and Finance 18*, 133–151.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroscedasticity. *Journal of Econometrics 31*, 307–327.

Bollerslev, T. and I. Domowitz (1993). Trading Patterns and Prices in the Interbank Foreign Exchange Market. *Journal of Finance 4*, 1421–1443.

Bollerslev, T. and J. Wooldridge (1992). Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances. *Econometric Reviews 11*, 143–172.

Brousseau, V. and F. Scacciavillani (1999). A Global Hazard Index for the World Foreign Exchange Markets. Working Paper Series No. 1, Frankfurt: European Central Bank.

Campos, J., N. R. Ericsson, and D. F. Hendry (2005). General-to-Specific Modeling: An Overview and Selected Bibliography. In J. Campos, D. F. Hendry, and N. R. Ericsson (Eds.), *General-to-Specific Modeling*. Cheltenham: Edward Elgar Publishing.

Clark, P. (1973). A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices. *Econometrica 41*, 135–155.

Cook, S. and D. F. Hendry (1994). The Theory of Reduction in Econometrics. In B. Hamminga and N. B. De Marchi (Eds.), *Idealization IV: Idealization in Economics*, Poznan Studies in the Philosophy of the Sciences and the Humanities 38. Amsterdam: Rodopi B.V.

Craig, E. (Ed.) (2000). *Concise Routledge Encyclopedia of Philosphy*. London: Routledge.

Crane, T. (1995). Possible worlds. In T. Honderich (Ed.), *The Oxford Companion to Philosophy*. Oxford: Oxford University Press.

Crano, W. D. and M. B. Brewer (2002). *Principles and Methods of Social Research*. London: Sage Publications.

de Vaus, D. (2001). *Research Design in Social Research*. London: Sage Publications.

Demos, A. and C. A. Goodhart (1996). The Interaction between the Frequency of Market Quotations, Spreads and Volatility in the Foreign Exchange Market. *Applied Economics 28*, 377–386.

Diebold, F., T. Gunther, and A. Tay (1998). Evaluating Density Forecasts with Applications to Financial Risk Management. *International Economic Review 39*, 863–883.

Engle, R. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflations. *Econometrica 50*, 987–1008.

Engle, R., C.-H. Hong, A. Kane, and J. Noh (1993). Arbitrage Valuation of Variance Forecasts with Simulated Options. In D. Chance and R. Trippi (Eds.), *Advances in Futures and Options Research*. Greenwich, Connecticut: JAI Press.

Engle, R. F. and J. Gonzalo Rangel (2005). The Spline GARCH Model for Unconditional Volatility and its Global Macroeconomic Causes. Available as http://pages.stern.nyu.edu/˜rengle/SplineGarch.pdf.

Engle, R. F., D. F. Hendry, and J.-F. Richard (1983). Exogeneity. *Econometrica 51*, 277–304.

Engle, R. F. and S. Kozicki (1993). Testing for Common Features. *Journal of Business and Economic Statistics 11*, 369–380.

Epps, T. W. and M. L. Epps (1976). The Stochastic Dependence of Security Price Changes and Transaction Volumes: Implications for the Mixture-of-Distribution Hypothesis. *Econometrica 44*, 305–321.

Florens, J.-P. and M. Mouchart (1980). Initial and Sequential Reduction of Bayesian Experiments. CORE Discussion Paper 15/1980, Louvain la Neuve (Belgium).

Florens, J.-P. and M. Mouchart (1985). Conditioning in Dynamic Models. *Journal of Time Series Analysis 6*, 15–34.

Florens, J.-P., M. Mouchart, and J.-F. Richard (1990). *Elements of Bayesian Statistics*. New York: Marcel Dekker.

Forbes, G. (1995). Possible Worlds. In J. Kim and E. Sosa (Eds.), *A Companion to Metaphysics*. Oxford: Blackwell Publishers Ltd.

Galati, G. (2003). Trading volume, volatility and spreads in foreign exchange markets: evidence from emerging market countries. BIS Working Paper.

Gardner, R. (1980). *Sterling-Dollar Diplomacy in Current Perspective: The Origins and the Prospects of Our International Economic Order*. New York: Columbia University Press.

Garman, M. B. and M. J. Klass (1980). On the Estimation of Securities Price Volatilities from Historical Data. *Journal of Business 53*, 67–78.

Geweke, J. (1984). Inference and Causality in Economic Time Series Models. In Z. Griliches and M. D. Intriligator (Eds.), *Handbook of Econometrics: Volume II*. Amsterdam: North Holland.

Geweke, J. (1986). Modelling the Persistence of Conditional Variance: A Comment. *Econometric Reviews 5*, 57–61.

Gilbert, C. L. (1989). LSE and the British Approach to Time Series Econometrics. *Oxford Economic Papers 41*, 108–128.

Gilbert, C. L. (1990). Professor Hendry's Econometric Methodology. In C. W. Granger (Ed.), *Modelling Economic Series*. Oxford: Oxford University Press. Earlier publised in Oxford Bulletin of Economics and Statistics 48 (1986), pp. 283-307.

Giot, P. (1999). *Econometric Studies on the Formation of Prices in the Financial and Commodity Markets*. Louvain la Neuve, Belgium: UCL Faculty of Economic, Social and Economic Sciences. Department of Economics. Nouvelle série-No. 341.

Giot, P. (2003). The Asian financial crisis: the start of a regime switch in volatility. UCL - Louvain la Neuve, Belgium: CORE Discussion Paper 2003/78.

Good, I. (1961). A Causal Calculus I. *British Journal for the Philosophy of Science 11*, 305–318.

Good, I. (1962). A Causal Calculus II. *British Journal for the Philosophy of Science 12*, 43–51.

Good, I. (1963). Errata and Corrigenda. *British Journal for the Philosophy of Science 13*, 88.

Goodhart, C. (1991). Every Minute Counts in Financial Markets. *Journal of International Money and Financial Markets 10*, 23–52.

Goodhart, C. (2000). News and the Foreign Exchange Market. In C. Goodhart (Ed.), *The Foreign Exchange Market*. London: MacMillan Press Ltd.

Grammatikos, T. and A. Saunders (1986). Futures Price Variability: A Test of Maturity and Volume Effects. *Journal of Business 59*, 319–330.

Haavelmo, T. (1944). The Probability Approach in Econometrics. *Econometrica 12*, iii–vi+1–115. Supplement.

Hamao, Y., R. Masulis, and N. Ng (1990). Correlations in price changes and volatility across international stock markets. *Review of financial studies 3*, 281–307.

Hartmann, P. (1999). Trading volumes and transaction costs in the foreign exchange market. evidence from daily dollar-yen spot data. *Journal of Banking and Finance 23*, 801–824.

Harvey, D., S. Leybourne, and P. Newbold (1997). Testing the equality of prediction mean squared errors. *International Journal of Forecasting 23*, 801–824.

Hawthorn, G. (1995). *Plausible Worlds.* Cambridge: Cambridge University Press.

Heil, J. (1998). *Philosophy of Mind. A Contemporary Introduction.* London: Routledge.

Helleiner, E. (1994). *States and the Reemergence of Global Finance.* London: Cornell University Press.

Hendry, D. F. (1987). Econometric Methodology: A Personal Perspective. In T. F. Bewley (Ed.), *Fifth World Congress of the Econometric Society.* Oxford: Oxford University Press.

Hendry, D. F. (1995). *Dynamic Econometrics.* Oxford: Oxford University Press.

Hendry, D. F. (1997). On Congruent Econometric Relations: A Comment. *Carnegie-Rochester Conference Series on Public Policy 47*, 163–190.

Hendry, D. F. (2003). J. Denis Sargan and the Origins of LSE Econometric Methodology. *Econometric Theory 19*, 457–480.

Hendry, D. F. and S. Cook (1994). The Theory of Reduction in Econometrics. *Poznan Studies in the Philosophy of the Sciences and the Humanities 38*, 71–100.

Hendry, D. F. and J. A. Doornik (2001). *Empirical Econometric Modelling Selection Using PcGive 10. Volume I.* London: Timberlake Consultants Press.

Hendry, D. F. and H.-M. Krolzig (1999). Improving on 'Data Mining Reconsidered' by K.D. Hoover and S.J. Perez. *Econometrics Journal 2*, 202–219.

Hendry, D. F. and H.-M. Krolzig (2001). *Automatic Econometric Model Selection using PcGets.* London: Timberlake Consultants Press.

Hendry, D. F. and H.-M. Krolzig (2004). We Ran One Regression. Research in Progress. Http://www.nuff.ox.ac.uk/Economics/Papers/2004/w17/OneReg.pdf.

Hendry, D. F. and M. Massmann (2005). Co-Breaking: Recent Advances and a Synopsis of the Literature. Research in Progress, available as http://www.nuff.ox.ac.uk/Users/Hendry/Papers/cobreakreview05.pdf.

Hendry, D. F. and J.-F. Richard (1990). On the Formulation of Empirical Models in Dynamic Econometrics. In C. W. Granger (Ed.), *Modelling Economic Series*. Oxford: Oxford University Press. Earlier published in Journal of Econometrics 20 (1982), pp. 3-33.

Honderich, T. (Ed.) (1995). *The Oxford Companion to Philosphy*. Oxford: Oxford University Press.

Hoover, K. D. and S. J. Perez (1999). Data Mining Reconsidered: Encompassing and the General-to-Specific Approach to Specification Search. *Econometrics Journal 2*, 167–191.

Horsefield, J. (1969). *The International Monetary Fund 1945 - 1965*, Volume III. Washington D.C.: IMF.

Hull, J. C. (2000). *Options, Futures and Other Derivatives*. London: Prentice Hall International. Fourth Edition.

Jansen, E. (2002). Statistical Issues in Macroeconomic Modelling. *Scandinavian Journal of Statistics 29*, pp. 193–217.

Jarque, C. and A. Bera (1980). Efficient Tests for Normality, Homoskedasticity, and Serial Independence of Regression Residuals. *Economics Letters 6*, 255259.

Jorion, P. (1995). Predicting Volatility in the Foreign Exchange Markets. *The Journal of Finance 50*, 507–528.

Jorion, P. (1996). Risk and Turnover in the Foreign Exchange Market. In J. Frankel et al. (Ed.), *The Microstructure of Foreign Exchange Markets*. Chicago: University of Chicago Press.

Karpoff, J. (1987). The relation between price changes and trading volume: A survey. *Journal of Financial and Quantitative Analysis 22*, 109–126.

Keynes, J. M. (1980a). Activities 1940-1944: Shaping the Post-war World, the Clearing Union. In D. Moggridge (Ed.), *The Collected*

Writings of John Maynard Keynes Vol. 25. Cambridge University Press.

Keynes, J. M. (1980b). Activities 1941-1946: Shaping the Post-war World, Bretton Woods and Reparations. In D. Moggridge (Ed.), *The Collected Writings of John Maynard Keynes Vol. 26.* Cambridge University Press.

Kim, J. (1996). *Philosophy of Mind.* Boulder CO: Westview Press.

Kim, J. and E. Sosa (Eds.) (1995). *A Companion to Metaphysics.* Oxford: Blackwell Publishers Ltd.

Kindleberger, C. P. (1993). *A Financial History of Western Europe.* Oxford: Oxford University Press.

Kindleberger, C. P. (2005). *Maniacs, Panics, and Crashes: A History of Financial Crisis.* New Jersey: John Wiley and Sons, Inc.

King, M. and S. Wadhwani (1990). Transmission of volatility between stock markets. *Review of financial studies 3*, 5–33.

Kluve, J. (2004). On the Role of Counterfactuals in Inferring Causal Effects. *Foundations of Science 9*, 65–101.

Lamoureux, C. G. and W. D. Lastrapes (1990). Heteroscedasticity in Stock Return Data: Volume versus GARCH Effects. *Journal of Finance*, 221–229.

Lewis, D. (1986a). Causal Explanation. In *Philosophical Papers. Volume II.* Oxford: Oxford University Press.

Lewis, D. (1986b). *On the Plurality of Worlds.* Oxford: Basil Blackwell.

Lewis, D. (1986c). Postscript to 'A Subjectivist's Guide to Objective Chance'. In *Philosophical Papers. Volume II.* Oxford: Oxford University Press.

Lewis, D. (1986d). Postscript to 'Causation'. In *Philosophical Papers. Volume II.* Oxford: Oxford University Press.

Lewis, D. (1986e). A Subjectivist's Guide to Objective Chance. In *Philosophical Papers. Volume II.* Oxford: Oxford University Press.

Lin, W., R. Engle, and T. Ito (1994). Do bulls and bears move across borders? international transmission of stock returns and volatility. *Review of financial studies 7*, 507–538.

Ljung, G. and G. Box (1979). On a Measure of Lack of Fit in Time Series Models. *Biometrika 66*, 265–270.

Lopez, J. A. (2001). Evaluating the Predictive Accuracy of Volatility Models. *Journal of Forecasting 20*, 87–109.

Loux, M. J. (1998). *Metaphysics. A Contemporary Introduction.* London: Routledge.

Lovell, M. C. (1983). Data Mining. *The Review of Economics and Statistics 65*, 1–12.

McAleer, M. (2005). Automated inference and learning in modeling financial volatility. *Econometric Theory 21*, 232–261.

Melvin, M. and Y. Xixi (2000). Public Information Arrival, Exchange Rate Volatility, and Quote Frequency. *The Economic Journal 110*, 644–661.

Mestad, V. (2002). Frå fot til feste - norsk valutarett og valutapolitikk 1873-2001. Supplement to *Penger og Kreditt* 3/2002. Occasional Papers No. 30. Oslo: Norges Bank.

Meyer, E. and J. Skjelvik (2006). Valutahandelsstatistikken - ny innsikt i valutamarkedet. *Penger og Kreditt 1*, 30–38. Oslo: Norges Bank.

Mincer, J. and V. Zarnowitz (1969). The Evaluation of Economic Forecasts. In J. Zarnowitz (Ed.), *Economic Forecasts and Expectations.* New York: National Bureau of Economic Research.

Mizon, G. (1995). Progressive Modeling of Macroeconomic Time Series: The LSE Methodology. In K. D. Hoover (Ed.), *Macroeconometrics. Developments, Tensions and Prospects.* Kluwer Academic Publishers.

Moravcsik, J. (1995). Potentiality/Actuality. In J. Kim and E. Sosa (Eds.), *A Companion to Metaphysics.* Oxford: Blackwell Publishers Ltd.

Myrdal, G. (1953). *The political element in the development of economic theory.* London: Routledge.

Myrdal, G. (1969). *Objectivity in Social Research.* London: Duckworth.

Nelson, D. B. (1991). Conditional Heteroscedasticity in Asset Returns: A New Approach. *Econometrica 51*, 485–505.

Newey, W. and K. West (1987). A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica 55*, 703–708.

Pagan, A. (1984). Econometric Issues in the Analysis of Regressions with Generated Regressors. *International Economic Review 25*, 221–247.

Pantula, S. (1986). Modelling the Persistence of Conditional Variance: A Comment. *Econometric Reviews 5*, 71–73.

Parkinson, G. (1995). Philosophy and Logic. In N. Jolley (Ed.), *The Cambridge Companion to Leibniz*. Cambridge: Cambridge University Press.

Parkinson, M. (1980). The Extreme Value Method for Estimating the Variance of the Rate of Return. *Journal of Business 53*, 61–65.

Patton, A. J. (2005). Volatility Forecast Evaluation and Comparison Using Imperfect Volatility Proxies. Downloadable from http://fmg.lse.ac.uk/˜patton/research.html.

Pettit, P. (1993). *The Common Mind. An Essay on the Psychology, Society and Politics*. New York.

Polanyi, K. (2002). *The Great Transformation. The Political and Economic Origins of Our Time*. London: Beacon Press. [Originally published in 1944]. This edition contains a foreword by Joseph Stiglitz.

Punch, K. F. (1998). *Introduction to Social Research*. London: Sage Publications.

Reichenbach, H. (1956). *The Direction of Time*. Berkeley and Los Angeles.

Ruben, D.-H. (1985). *The Metaphysics of the Social World*. London: Routledge.

Ruse, M. (1988). *Philosophy of Biology Today*. Albany, NY: State University of New York Press.

Salmon, W. C. (1993a). Causality: Production and Propagation. In E. Sosa and M. Tooley (Eds.), *Causality*. Oxford: Oxford University Press.

Salmon, W. C. (1993b). Probabilistic Causality. In E. Sosa and M. Tooley (Eds.), *Causality*. Oxford: Oxford University Press.

Schumpeter, J. (1949). Science and Ideology. *The American Economic Review 39*, 345–359.

Searle, J. (1991). *Minds, Brains and Science.* London: Penguin Books.

Starica, C. and T. Mikosch (2004). Non-stationarities in financial time series, the long range dependence and the IGARCH effects. *Review of Economics and Statistics 86*, 378–390.

Suppes, P. (1970). *A Probabilistic Theory of Causality.* Amsterdam.

Tauchen, G. and M. Pitts (1983). The Price Variability-Volume Relationship on Speculative Markets. *Econometrica 51*, 485–505.

Tiles, M. (1989). *The Philosophy of Set Theory.* New York: Dover Publications, Inc.

van Dijk, D., H. Munandar, and C. M. Häfner (2005). The euro introduction and non-euro currencies. Rotterdam: Tinbergen Institute Discussion Paper TI 2005-044/4.

Veredas, D. (2002). *Econometric Modeling of Financial Durations.* Louvain la Neuve, Belgium: UCL Faculty of Economic, Social and Economic Sciences. Department of Economics. Thesis Number 403/2002.

Weber, M. (1994). Objectivity and Understanding in Economics. In D. M. Hausman (Ed.), *The Philosophy of Economics.* Cambridge: Cambridge University Press.

West, K., H. Edison, and D. Cho (1993). A Utility Based Comparison of Some Models for Exchange Rate Volatility. *Journal of International Economics 35*, 23–45.

White, H. (1980). A Heteroskedasticity-Consistent Covariance Matrix and a Direct Test for Heteroskedasticity. *Econometrica 48*, 817–838.

White, H. (1990). A consistent model selection procedure based on m-testing. In C. W. Granger (Ed.), *Modelling Economic Series: Readings in Econometric Methodology.* Oxford: Clarendo Press.

Wikipaedia (2006, April). United Nations Monetary and Financial Conference. Internet On-Line Encyclopaedia.

Data appendix

This data appendix contains three parts. Part A describes the sources and transformations underlying the weekly data used in chapters 1 to 4. Part B describes the sources and transformations underlying the bi-weekly data used in chapter 5, whereas part C provides the details of the origin of the spot NOK/EUR volume data used in the same chapter. All data transformations are undertaken in Ox 3.4 and EViews 5.1.

A. Sources and transformations of weekly data

S_t	BID NOK/1EUR closing value of the last trading day of week t. Before 1.1.1999 the BID NOK/1EUR rate is obtained by the formula BID NOK/100DEM $\times$ 0.0195583, where 0.0195583 is the official DEM/1EUR conversion rate 1.95583 DEM = 1 EUR divided by 100. The source of the BID NOK/100DEM series is Olsen Financial Technologies and the source of the BID NOK/1EUR series is Reuters.
r_t	$(\log S_t - \log S_{t-1}) \times 100$
V_t^w	$\{\{\log[S_t + I(S_t = S_{t-1}) \times 0.0009] - \log(S_{t-1})\} \times 100\}^2$. $I(S_t = S_{t-1})$ is an indicator function equal to 1 if $S_t = S_{t-1}$ and 0 otherwise, and $S_t = S_{t-1}$ occurs for $t = 10/6/1994$, $t = 19/8/1994$ and $t = 17/2/2000$.
v_t^w	$\log V_t^w$
V_t^r	$\sum_n [\log(S_n/S_{n-1}) \times 100]^2$, where $n = 1(t), 2(t), ..., N(t)$ and $1(t) - 1 = N(t-1)$. $S_{1(t)}$ is the first BID NOK/1EUR opening exchange rate of week t, $S_{2(t)}$ is the first closing rate, $S_{3(t)}$ is the second opening rate, and so on, with $S_{N(t)}$ denoting the last closing rate of week t, that is, $S_{N(t)} = S_t$.

v_t^r $\log V_t^r$

V_t^{hl} $[\log(S_t^h/S_t^l) \times 100]^2$, where S_t^h and S_t^l are the maximum and minimum values of bid NOK/EUR in week t.

v_t^{hl} $\log V_t^{hl}$

M_t BID USD/EUR closing value of the last trading day of week t. Before 1.1.1999 the BID USD/EUR rate is obtained with the formula 1.95583/(BID DEM/USD). The source of the BID DEM/USD and BID USD/EUR series is Reuters.

m_t $\log M_t$

M_t^w $\{\{\log[M_t + I(M_t = M_{t-1}) \times 0.0009] - \log(M_{t-1})\} \times 100\}^2$. $I(M_t = M_{t-1})$ is an indicator function equal to 1 if $M_t = M_{t-1}$ and 0 otherwise.

m_t^w $\log M_t^w$

Q_t Weekly number of NOK/EUR quotes (NOK/100DEM before 1.1.1999). The underlying data is a daily series from Olsen Financial Technologies, and the weekly values are obtained by summing the values of the week.

q_t $\log Q_t$. This series is "synthetic" in that it has been adjusted for changes in the underlying quote collection methodology at Olsen Financial Technologies. More precisely q_t has been generated under the assumption that Δq_t is equal to zero in the weeks containing Friday 17 August 2001 and Friday 5 September 2003, respectively. In the first week the underlying feed was changed from Reuters to Tenfore, and on the second a feed from Oanda was added.

Δq_t $q_t - q_{t-1}$. The values of this series has been set to zero in the weeks containing Friday 24 August 2001 and Friday 5 September 2003, respectively.

O_t Closing value of the Brent Blend spot oilprice in USD per barrel in the last trading day of week t. The untransformed series is Norges Bank database series D2001712.

o_t $\log O_t$

O_t^w $\{\{\log[O_t + I(O_t = O_{t-1}) \times 0.009] - \log(O_{t-1})\} \times 100\}^2$. $I(O_t = O_{t-1})$ is an indicator function equal to 1 if $O_t = O_{t-1}$ and 0 otherwise, and $O_t = O_{t-1}$ occurs three times, for t = 1/7/1994, t = 13/10/1995 and t = 25/7/1997.

o_t^w $\log O_t^w$

X_t Closing value of the main index of the Norwegian Stock Exchange (TOTX) in the last trading day of week t. The source of the daily untransformed series is EcoWin series ew:nor15565.

x_t $\log X_t$

X_t^w $\{[\log(X_t/X_{t-1})] \times 100\}^2$. $X_t = X_{t-1}$ does not occur for this series.

x_t^w $\log X_t^w$

U_t Closing value of the composite index of the New York Stock Exchange (the NYSE index) in the last trading day of week t. The source of the daily untransformed series is EcoWin series ew:usa15540.

U_t^w $\{[\log(U_t/U_{t-1})] \times 100\}^2$. $U_t = U_{t-1}$ does not occur for this series.

u_t^w $\log U_t^w$

IR_t^{emu} Average of closing values of the 3-month market interest rates of the European Monetary Union (EMU) countries in the last trading day of week t. The source of the daily untransformed series is EcoWin series ew:emu36103.

ir_t^{emu} $(\Delta IR_t^{emu})^2$.

F_t The Norwegian central bank's main policy interest-rate, the so-called "folio", at the end of the last trading day of week t. The source of the untransformed daily series is Norges Bank's web-pages.

f_t^a $|\Delta F_t| \times I_a$, where I_a is an indicator function equal to 1 in the period 1 January 1999 - Friday 30 March 2001 and 0 elsewhere

f_t^b $|\Delta F_t| \times I_b$, where I_b is an indicator function equal to 1 after Friday 30 March 2001 and 0 before

id_t Russian moratorium impulse dummy, equal to 1 in the week containing Friday 28 August 1998 and 0 elsewhere.

sd_t Step dummy, equal to 0 before 1997 and 1 thereafter.

ia_t Skewness variable, equal to 1 when $r_t > 0$ and 0 otherwise.

h_{lt} $l = 1, 2, \ldots, 8$. Holiday variables with values equal to the number of official Norwegian holidays that fall on weekdays. For example, if 1 January falls on a Saturday then h_{1t} is equal to 0, whereas if 1 January falls on a Monday, then h_{1t} is equal to 1. h_{2t} is associated with Maundy Thursday and Good Friday and thus always equal to 2, h_{3t} with Easter Monday and thus always equal to 1, h_{4t} with Labour Day (1 May), h_{5t} with the Norwegian national day (17 May), h_{6t} with Ascension Day, h_{7t} with Whit Monday and h_{8t} with Christmas (Christmas Day and Boxing Day). Source: Http://www.timeanddate.com.

B. Sources and transformations of bi-weekly data

t, t_2, t_1 — Time indices. t denotes the week in question, t_2 stands for the period that comprises the last two trading days in Norway of week t, that is, typically Thursday and Friday when neither is a holiday, and t_1 stands for the period that comprises the other trading days in Norway of week t, that is, typically Monday, Tuesday and Wednesday. The symbolism $t_2 - 1$ denotes the second part of week $t-1$, $t_1 - 1$ denotes the first part of week $t-1$, $t_2 - 2$ denotes the second part of week $t-2$, and so on.

s_{t_2}, s_{t_1} — $\log S_{t_2}$, $\log S_{t_1}$. S variables denote BID NOK/1EUR exchange rates: Open (07:00 GMT), close (21:50 GMT), high and low. $S^C_{t_2}$ stands for the closing value in the last trading day of t_2, $S^C_{t_1}$ stands for the closing value in the last trading day of t_1, $S^O_{t_2}$ stands for the opening value in the first trading day of t_2, $S^O_{t_1}$ stands for the opening value in the first trading day of t_1, $S^H_{t_2}$ stands for the highest value in t_2, $S^H_{t_1}$ stands for the highest value in t_1, $S^L_{t_2}$ stands for the lowest value in t_2 and $S^L_{t_1}$ stands for the lowest value in t_1. The corresponding log-transformed exchange rates are denoted in small letters, that is, $s^c_{t_2}$, $s^c_{t_1}$, $s^o_{t_2}$, $s^o_{t_1}$, $s^h_{t_2}$, $s^h_{t_1}$, $s^l_{t_2}$ and $s^l_{t_1}$. The source of the daily untransformed data is Reuters.

$r^{se}_{t_2}$, $r^{se}_{t_1}$ — Period or "start-end" log-returns in percent. Specifically, $r^{se}_{t_2} = (s^c_{t_2} - s^o_{t_2}) \times 100$, and $r^{se}_{t_1} = (s^c_{t_1} - s^c_{t_2-1}) \times 100$.

$r^{hl}_{t_2}$, $r^{hl}_{t_1}$ — Range or "high-low" log-returns in percent. Specifically, $r^{hl}_{t_2} = (s^h_{t_2} - s^l_{t_2}) \times 100$, and $r^{hl}_{t_1} = (s^h_{t_1} - s^l_{t_1}) \times 100$.

$v^{se}_{t_2}$, $v^{se}_{t_1}$ — $\log V^{se}_{t_2}$, $\log V^{se}_{t_1}$. V^{se} variables denote period volatility in basis points. Specifically, $V^{se}_{t_2} = (r^{se}_{t_2} \times 100)^2$ and $V^{se}_{t_1} = (r^{se}_{t_1} \times 100)^2$. In order to avoid the log-transformation being applied on zero-values, $r^{se}_{t_2}$ is replaced by min $|r^{se}_{t_2}|$ when $r^{se}_{t_2} = 0$ where the minimum is taken over the set of non-zero values of $r^{se}_{t_2}$. Similarly $r^{se}_{t_1}$ is replaced by min $|r^{se}_{t_1}|$ when $r^{se}_{t_1} = 0$ where the minimum is taken over the set of non-zero values of $r^{se}_{t_1}$.

$v^{hl}_{t_2}$, $v^{hl}_{t_1}$ — $\log V^{hl}_{t_2}$, $\log V^{hl}_{t_1}$. V^{hl} variables denote range volatility in basis points. Specifically, $V^{hl}_{t_2} = (r^{hl}_{t_2} \times 100)^2$ and $V^{hl}_{t_1} = (r^{hl}_{t_1} \times 100)^2$. There is no need to handle zero-values since neither $r^{hl}_{t_2} = 0$ nor $r^{hl}_{t_1} = 0$ occur.

q_{t_2}, q_{t_1} — $\log Q_{t_2}$, $\log Q_{t_1}$. Q_{t_2} is the number of NOK/EUR quotes in t_2, whereas Q_{t_1} is the number of NOK/EUR quotes in t_1. The source of the untransformed data is Olsen Financial Technologies (OFT) and the variables have been adjusted for changes in the underlying quote collection methodology at OFT. More precisely q_{t_2} (q_{t_2}) has been generated under the assumption that Δq_{t_2} (Δq_{t_2}) is equal to zero in the weeks containing Friday 17 August 2001 and Friday 5 September 2003, respectively. In the first week the underlying data feed was changed from Reuters to Tenfore, and on the second a feed from Oanda was added to the Tenfore feed.

$\bar{q}_{t_2-1}$, $\bar{q}_{t_1-1}$ — Lagged averages of q_{t_2} and q_{t_1}, respectively, where a superscript indicates the number of terms in the average. For example, $\bar{q}^2_{t_2-1} = \frac{1}{2}(q_{t_2-1} + q_{t_2-2})$, $\bar{q}^2_{t_2-1} = \frac{1}{3}(q_{t_2-1} + q_{t_2-2} + q_{t_2-3})$, and so on.

z_{t_2} — $\log Z_{t_2}$. Z_{t_2} variables denote measures of spot NOK/EUR transaction volumes by banks in Norway in t_2: Total volume, the volume of big banks, the volume of medium-sized banks and the volume of small banks. The four variables are denoted $Z^{tot}_{t_2}$, $Z^{big}_{t_2}$, $Z^{med}_{t_2}$ and $Z^{sma}_{t_2}$, and by definition $Z^{tot}_{t_2} = Z^{big}_{t_2} + Z^{med}_{t_2} + Z^{sma}_{t_2}$. The source of the untransformed data is Norges Bank, see appendix 2 for more details.

$\bar{z}_{t_2-1}$ — Lagged averages of z_{t_2} where a superscript indicates the volume category in question and the number of terms in the average. For example, $\bar{z}^{tot/2}_{t_2-1} = \frac{1}{2}(z^{tot}_{t_2-1}+z^{tot}_{t_2-2})$, $\bar{z}^{tot/3}_{t_2-1} = \frac{1}{3}(z^{tot}_{t_2-1}+z^{tot}_{t_2-2}+z^{tot}_{t_2-3})$, and so on.

$m^{se}_{t_2}$ — $\log M^{se}_{t_2}$, where $M^{se}_{t_2}$ is USD/EUR volatility in basis points constructed in the same way as $V^{se}_{t_2}$. The source of the untransformed daily BID USD/EUR series is Reuters.

$o^{se}_{t_2}$ — $\log O^{se}_{t_2}$, where $O^{se}_{t_2}$ is oilprice volatility in basis points at t_2 . If $o^c_{t_2}$ and $o^c_{t_1}$ denote the log of the Brent Blend spot oilprice in USD per barrel in the last trading day of t_2 and t_1, respectively, then $O^{se}_{t_2} = [(o^c_{t_2} - o^c_{t_1}) \times 100^2]^2$, where $(o^c_{t_2} - o^c_{t_1})$ has been zero-adjusted in the same way as $r^{se}_{t_2}$ so that the log is not applied on zero values. The underlying untransformed daily series consists of Norges Bank database series D2001712.

$x^{se}_{t_2}$ — $\log X^{se}_{t_2}$, where $X^{se}_{t_2}$ is the volatility in basis points of the main index (TOTX) of the Norwegian stock exchange at t_2. If $x^c_{t_2}$ and $x^c_{t_1}$ denote the log of the closing values in the last trading day of t_2 and t_1, respectively, then $X^{se}_{t_2} = [(x^c_{t_2} - x^c_{t_1}) \times 100^2]^2$, where $(x^c_{t_2} - x^c_{t_1})$ has been zero-adjusted in the same way as $r^{se}_{t_2}$ so that the log is not applied on zero values. The underlying untransformed daily series consists of EcoWin database series ew:nor15565.

$u^{se}_{t_2}$ — $\log U^{se}_{t_2}$, where $U^{se}_{t_2}$ is the volatility in basis points of the New York Stock Exchange (NYSE) index at t_2. If $u^c_{t_2}$ and $u^c_{t_1}$ denote the log of the closing values in the last trading day of t_2 and t_1, respectively, then $U^{se}_{t_2} = [(u^c_{t_2} - u^c_{t_1}) \times 100^2]^2$, where $(x^c_{t_2} - x^c_{t_1})$ has been zero-adjusted in the same way as $r^{se}_{t_2}$ so that the log is not applied on zero values. The underlying untransformed daily series consists of EcoWin database series ew:usa15540.

$ir^{emu}_{t_2}$ — A measure of EU short-term market interest rate volatility in basis points. If $IR^{emu}_{t_2}$ and $IR^{emu}_{t_1}$ denote the averages of the EMU countries' 3-month money market interest rates in percent (closing values) in the last trading day of t_2 and t_1, respectively, then $ir^{emu}_{t_2} = [(IR^{emu}_{t_2} - IR^{emu}_{t_1}) \times 100]^2$. The underlying untransformed daily series consists of EcoWin database series ew:emu36103.

$ir^{no}_{t_2}$ — A measure of Norwegian short-term market interest rate volatility in basis points. If $IR^{no}_{t_2}$ and $IR^{no}_{t_1}$ denote Norwegian 3-month money market interest rates in percent (closing values) in the last trading day of t_2 and t_1, respectively, then $ir^{no}_{t_2} = [(IR^{no}_{t_2} - IR^{no}_{t_1}) \times 100]^2$. The variable $ir^{no,b}_{t_2}$ is the short term interest rate volatility in the partial inflation targeting regime, and $ir^{no,c}_{t_2}$ in the full inflation targeting regime. Specifically, $ir^{no,b}_{t_2} = ir^{no}_{t_2}$ until 30 March 2001 and zero afterwards, and $ir^{no,c}_{t_2} = ir^{no}_{t_2}$ after 30 March 2001 and zero before. The $ir^{no,b}_{t_2}$ and $ir^{no,c}_{t_2}$ variables are further decomposed according to whether the Norwegian central bank (Norges Bank) changes its policy rate (the socalled "Folio") or not, and these variables appear with the additional superscripts 0 or Δ. For example, $ir^{no,c0}_{t_2}$ is equal to $ir^{no,c}_{t_2}$ when Norges Bank does not change its policy rate in the full inflation period and zero when it does, whereas $ir^{no,c\Delta}_{t_2}$ is equal to $ir^{no,c}_{t_2}$ when Norges Bank changes its policy interest rate in the full inflation period and zero otherwise. Similarly for $ir^{no,b0}_{t_2}$ and $ir^{no,b\Delta}_{t_2}$ in the partial inflation targeting period. The underlying untransformed daily series consists of EcoWin database series ew:nor14103.

$id^2_{t_2}$, $id^3_{t_2}$ — Impulse dummies. $id^2_{t_2}$ is equal to 1 in the week containing Friday 11 January 2002 and 0 elsewhere, and $id^3_{t_2}$ is equal to 1 in the week containing Friday 23 April 2004 and 0 elsewhere.

ia_{t_2} — Skewness variable equal to 1 when $r^{se}_{t_2} > 0$ and 0 otherwise.

C. Sources and transformation of spot NOK/EUR volume data

The Z_{t_2} variables are constructed using information obtained from a form that the most important currency banks within Norwegian regulatory borders fill out and send to Norges Bank every week, and we are indebted to Erik Meyer at the statistics department of Norges Bank for confirming that the data essentially consists of NOK-trading (the documentation of the data suggests otherwise). The data have been collected since the beginning of the 1990s, but have undergone significant changes with respect to data definitions, data collection methodology and data correction methodology, all only partially documented. Table 7.1 contains a form that is similar to the one which each reporting bank submitted electronically over the period 1 January 1999 - 7 January 2005. Collection of the data discontinued after 7 January 2005 in order to prepare for an entirely new, more detailed and comprehensive data methodology. The new methodology was implemented in October 2005, see Meyer and Skjelvik (2006). An "asset" refers to a purchase contract, that is, a purchase of non-Norwegian currency paying with Norwegian kroner, and a "liability" refers to a sales contract, that is, a sales of non-Norwegian currency paid with Norwegian kroner. For all fields the amount reported is in Norwegian kroner even if the contract is denominated in Euros. If the contract is denominated in Euros then the value of the contract is transformed to Norwegian kroner using the official *daily* exchange rate of Norges Bank of the day in which the contract is made, that is, not the exchange rate corresponding to the day in which the contract is cleared. Accordingly, the trading volume of reporting banks comprise not only NOK/EUR trading, but also NOK/USD trading, NOK/GBP trading, and so on. For this reason we use only fields

2 and 7, that is, uncleared spot assets and liabilities, since the spot market on Norwegian currency is dominated by NOK/EUR trading. In the interbank spot market currency purchases and sales are made with actual delivery typically taking place two trading days later. The category "uncleared" thus refers to transactions which took place in the last two trading days of the week. This explains our focus on exchange rate volatility over the last two trading days of the week.

Let $Z_{it_2}^{j}$ to denote the value of field j for bank i at t_2, and for convenience we will refer to Z_{it_2} variables as NOK/EUR trading although it strictly speaking may comprise some non-Euro trading against the NOK. Total spot NOK/EUR transaction volume is then defined as

$$Z_{t_2}^{tot} = \sum_i (Z_{it_2}^2 + Z_{it_2}^5).$$

In words, the sum of all banks' purchase and sales volumes of spot NOK/EUR in the last two trading days of week t. The volume of big, medium sized and small banks are all sub sums of this expression. If *big* refers to the set of big banks in terms of currency volume, *med* to the set of medium sized banks and *sma* to the set of small sized banks, then the variables are defined as

$$Z_{t_2}^{big} = \sum_{i \in big} (Z_{it_2}^2 + Z_{it_2}^5)$$

$$Z_{t_2}^{med} = \sum_{i \in med} (Z_{it_2}^2 + Z_{it_2}^5)$$

$$Z_{t_2}^{sma} = \sum_{i \in sma} (Z_{it_2}^2 + Z_{it_2}^5),$$

where by definition $Z_{t_2}^{tot} = Z_{t_2}^{big} + Z_{t_2}^{med} + Z_{t_2}^{sma}$. When comparing the volumes of the individual banks, which are relatively stable over the sample, the banks classified as big are substantially bigger than the others in terms of spot NOK/EUR volume, and the banks classified as small are substantially smaller than than the others in terms of spot NOK/EUR volume. For confidentiality reasons we cannot disclose which banks enter in which category.

Table 7.1: Norges Bank's currency volume form

		Field	Amount (in NOK)
Spot assets	Total	1	
	Uncleared	2	
	Norwegian banks	3	
	Norwegian customers	4	
	Foreign customers	5	
Spot liabilities	Total	6	
	Uncleared	7	
	Norwegian banks	8	
	Norwegian customers	9	
	Foreign customers	10	
Forward assets	Total	11	
	Norges Bank	12	
	Other Norwegian banks	13	
	Norwegian customers	14	
	Foreign customers	15	
Forward liabilities	Total	16	
	Norges Bank	17	
	Other Norwegian banks	18	
	Norwegian customers	19	
	Foreign customers	20	