

EXCHANGE RATE VOLATILITY AND THE MIXTURE OF DISTRIBUTION HYPOTHESIS

Luc Bauwens *
Dagfinn Rime[†]
Genaro Sucarrat[‡]

15 July 2005

Abstract

This paper sheds new light on the mixture of distribution hypothesis by means of a study of the weekly exchange rate volatility of the Norwegian krone. In line with other studies we find that the impact of information arrival on exchange rate volatility is positive and statistically significant, and that the hypothesis that an increase in the number of traders reduces exchange rate volatility is not supported. The novelties of our study consist in documenting that the positive impact of information arrival on volatility is relatively stable across three different exchange rate regimes, and in that the impact is relatively similar for both weekly volatility and weekly realised volatility. It is not given that the former should be the case since exchange rate stabilisation was actively pursued by the central bank in parts of the study period. We also report a case in which undesirable residual properties attained within traditional frameworks are easily removed by applying the log-transformation on volatilities.

JEL Classification: F31

Keywords: Exchange rate volatility, log-linear analysis, mixture of distribution hypothesis

*CORE and Department of Economics, Université catholique de Louvain, Belgium. Email: bauwens@core.ucl.ac.be.

[†]Norges Bank, Norway. Email: dagfinn.rime@norges-bank.no.

[‡]Corresponding author. CORE and Department of Economics, Université catholique de Louvain, Belgium. Email: sucarrat@core.ucl.ac.be. Homepage: <http://www.core.ucl.ac.be/~sucarrat/index.html>.

1 Introduction

If exchange rates walk randomly and if the number of steps depends positively on the number of information events, then exchange rate volatility over a given period should increase with the number of information events in that period. This chain of reasoning is the essence of the so-called "mixture of distribution hypothesis" (MDH) associated with Clark (1973) and others. Several versions of the MDH have been put forward, including one that suggests the size of the steps depends negatively on the number of traders, see for example Tauchen and Pitts (1983). In other words, an increase in the number of traders—a measure of liquidity—should decrease the size of the steps and thus volatility. Exchange rate volatility may of course depend on other factors too, including country-specific institutional factors, market conditions and economic fundamentals. Bringing such factors together in a general framework and trying to disentangle their distinct effects on exchange rate volatility leads to economic or explanatory volatility modelling as opposed to "pure" forecast modelling, which may remain silent about the economic reasons for variation in volatility.

When Karpoff (1987) surveyed the relationship between financial volatility and trading volume—a measure of information intensity—during the mid-eighties, only one out of the nineteen studies he cited was on exchange rates. The increased availability of data brought by the nineties has changed this, and currently we are aware of ten studies that directly or indirectly investigate the relationship between exchange rate volatility and information intensity. The ten studies are summarised in table 1 and our study of Norwegian weekly exchange rate volatility from 1993 to 2003 adds to this 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 changes in the number of information events on exchange rate volatility is positive and statistically significant, recursive parameter 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 number of traders reduces exchange rate volatility.

Another contribution of our study concerns the economic modelling of exchange rate volatility as such. We report a case in which undesirable residual properties are easily removed by applying the logarithmic transformation on volatilities. In particular, we show that OLS-regressions of the *logarithm* of volatility on its own lags and on several economic variables can produce uncorrelated and homoscedastic residuals. Moreover, in the log of realised volatility case the residuals are also normal. When Geweke (1986), Pantula (1986) and Nelson (1991) proposed that volatilities should be analysed in logs it was first and foremost in order to ensure non-negativity. In our case the motivation stems from unsatisfactory residual properties and fragile inference results. Without the log-transformation we do not generally produce uncorrelated residuals, and when we do the results are very sensitive to small changes in specification.

The rest of this paper contains three sections. In section 2, we review the link between exchange rate volatility and the MDH hypothesis, and discuss measurement issues. We also present our data and other economic variables that we believe may impact on the

volatility of the Norwegian exchange rate. In section 3, we present the models we use and the empirical results. We conclude in the last section, whereas an appendix provides the details of the data sources and transformations.

2 Exchange rate volatility and economic determinants

The purpose of this section is to motivate and describe our exchange rate volatility measures, and the economic determinants that we use in our empirical study. In subsection 2.1, we define our volatility measures and present the Norwegian exchange rate data. We make a distinction between period volatility on the one hand and within or intra-period volatility on the other, arguing that analysis of both is desirable since level-expectations may have an impact. In subsection 2.2, we review the link between volatility and the MDH, and after presenting our quote frequency data we explain how we use them to construct the explanatory variables we include in our volatility equations. In subsection 2.3, we motivate and describe the other economic determinants of volatility which we include as explanatory variables in the empirical part.

2.1 Period vs. intra-period volatility measures

Conceptually we may distinguish between period volatility on the one hand and within or intra-period volatility on the other. If $\{S_0, S_1, \dots, S_n, \dots, S_{N-1}, S_N\}$ denotes a sequence of exchange rates between two currencies at times $0, 1, \dots, N$, then the squared (period) return $[\log(S_N/S_0)]^2$ is an example of a period measure of volatility, and realised volatility $\sum_{n=1}^N [\log(S_n/S_{n-1})]^2$ is an example of a within-period measure of volatility. (Another example of a within-period measure of volatility is high - low.) It has been showed that realised volatility is an unbiased and consistent measure of integrated volatility under certain assumptions, see Andersen, Bollerslev, Diebold, and Labys (2001). The reader should be aware though that nowhere do we rely on such assumptions. Rather, our focus is on the *formula* of realised volatility. The main difference between period volatility and realised volatility is that 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 measures may produce substantially different results. Essentially this can be due to one of two reasons. If the random walk model provides a decent description of how exchange rates behave, then it is due to chance. On the other hand, if there are strong level-effects present among market participants, then the return back to the level of S_0 might be due to market expectations rather than chance. Although market participants' views on exchange rate level clearly matter, we believe most observers would agree that such level-effects are relatively small or infrequent on a day-to-day basis for most exchange rates. Differently put, at very short horizons the random walk model provides a reasonably good description of exchange rate increments. However, the two measures are still qualitatively different, so that any eventual differences in their

relation with (say) the rate of information arrival should be investigated—in particular for weekly data where level-expectations is more likely to play a role.

Our period measure will be referred to as "weekly volatility" whereas our within-period measure will be referred to as "within-weekly volatility" or "realised volatility". Weekly volatility is just the squared return from the end of one week to the end of the subsequent week. More precisely, if $S_{N(t)}$ denotes the closing value in the last day of trading in week t and $S_{N(t-1)}$ denotes the closing value in the last day of trading in the previous week, then weekly volatility recorded in week t is denoted by V_t^w and defined as

$$V_t^w = [\log(S_{N(t)}/S_{N(t-1)})]^2. \quad (1)$$

On the other hand, realised volatility in week t , denoted by V_t^r , is the sum of squared returns of the sequence $\{S_{N(t-1)}, S_{1(t)}, S_{2(t)}, \dots, S_{N(t)}\}$, that is,

$$V_t^r = \sum_{n=1(t)}^{N(t)} [\log(S_n/S_{n-1})]^2, \quad (2)$$

where $1(t) - 1 = N(t - 1)$. It should be noted though that we use only a small subset of the within-week observations in the construction of realised volatility (typically ten observations per week).

In order to distinguish between volatilities and logs of volatilities we use lower and upper case letters. So $v_t^w = \log V_t^w$ and $v_t^r = \log V_t^r$. Our data set span the period from 8 January 1993 to 26 December 2003, a total of 573 observations, and before 1 January 1999 we use the BID NOK/DEM exchange rate converted to euro-equivalents with the official conversion rate 1.95583 DEM = 1 EURO. After 1 January 1999 we use the BID NOK/EUR rate.

The main characteristics of the two measures are contained in table 2 and in figure 1. At least three attributes of the graphs should be noted. First, although the two measures of volatility are similar level-wise, that is, if plotted in the same diagram they would be "on top of each other", the sample correlation between the log of weekly volatility and the log of realised volatility is only 0.55. In other words, the two measures differ considerably and one of the differences is that the realised volatility measure is less variable. Second, sustained increases in volatility around 1 January 1999 and 29 March 2001 are absent—or at least seemingly so. On the first date the current central bank governor assumed the job and reinterpreted the guidelines, which in practice entailed a switch from exchange rate stabilisation to "partial" inflation targeting. On the second date the Norwegian central bank was instructed by the Ministry of Finance to pursue an inflation target of 2.5% as main policy objective. One might have expected that both of these changes would have resulted in shifts upwards in volatility. However, if this is the case then this is not evident by just looking at the graphs. Alternatively, the apparent absence of shifts in volatility might be due to the fact that the markets had expected these changes and already adapted to them. A third interesting feature is that there is a marked and lasting increase in volatility 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 volatility of financial markets in general. In the case of Norwegian exchange rate volatility, 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.

2.2 MDH and quote frequency

If exchange rates follow a random walk and if the number of steps depends positively on the number of information events, then exchange rate volatility over a given period should increase with the number of information events in that period. This chain of reasoning is the essence of the MDH, an acronym which is due to the statistical setup used by Clark (1973). Formally, focusing on the economic content of the hypothesis, the MDH can also be formulated as

$$\Delta s_t = \sum_{n=1}^{N(t)} \Delta s_n, \quad n = 1, \dots, N(t), \quad s_0 = s_{N(t-1)}, \quad (3)$$

$$\{\Delta s_n\} \text{ IID}, \quad \Delta s_n \sim N(0, 1), \quad (4)$$

$$\frac{\partial E[N(t)|\nu_t]}{\partial \nu_t} > 0. \quad (5)$$

where $s_t = \log S_t$. The first line (3) states that the price increment of period t is equal to the sum of the intra-period increments, (4) is a random walk hypothesis (any "random walk" hypothesis would do), and (5) states that the mean of the number of intra-period increments $N(t)$ conditioned on the number of information events ν_t in period t is strictly increasing in ν_t . Several variations of the MDH have been formulated, but for our purposes it is the economic content of Tauchen and Pitts (1983) that is of most relevance. In a nutshell, they argue that an increase in the number of traders reduces the size of the intra-period increments. Here this is akin to replacing (4) with (say)

$$\Delta s_n = \sigma_n(\eta_n)z_n, \quad \sigma'_n < 0, \quad \{z_n\} \text{ IID}, \quad z_n \sim N(0, 1), \quad (6)$$

where η_n denotes the number of traders at time n and where σ'_n is the derivative. But markets differ and theoretical models thus have to be adjusted accordingly. In particular, in a comparatively small currency market like the Norwegian an increase in the number of currency traders is also likely to increase substantially the number of increments per period, that is, $N(t)$, resulting in two counteracting effects. One effect would tend to reduce period-volatility through the negative impact on the size of the intra-period increments, whereas the other effect would tend to increase period-volatility by increasing the number of increments. So it is not known beforehand what the overall effect will be. Replacing (5) with

$$\frac{\partial E[N(t)|\nu_t, \eta_t]}{\partial \nu_t} > 0, \quad \frac{\partial E[N(t)|\nu_t, \eta_t]}{\partial \eta_t} > 0. \quad (7)$$

means the conditional mean of the number of increments $N(t)$ is strictly increasing in both the number of information events ν_t and the number of traders η_t . Taking (7) together with (3) and (6) as our starting point we may formulate our null hypotheses a

$$\frac{\partial \text{Var}(\Delta s_t | \nu_t, \eta_t)}{\partial \nu_t} > 0 \quad (8)$$

$$\frac{\partial \text{Var}(\Delta s_t | \nu_t, \eta_t)}{\partial \eta_t} < 0. \quad (9)$$

In words, the first hypothesis states that an increase in the number of information events given the number of traders increases period volatility, whereas the second holds that an increase in the number of traders without changes in the information intensity reduces volatility. That (8) is the case is generally suggested by table 1, whereas (9) is suggested by Tauchen and Pitts (1983). However, it should be noted that the empirical results of Jorion (1996) and Bjønnes, Rime, and Solheim (2005) do not support the hypothesis that an increase in the number of traders reduces volatility.

The most commonly used indicators of information arrival are selected samples from the news-screens of Reuters or Telerate, quoting frequency, the number of transacted contracts and transaction volume. The former is laborious to construct and at any rate not exhaustive with respect to the range of information events that might induce price revision, and the latter two are not readily available in foreign exchange markets. So quote frequency is our indicator of information arrival. 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. We denote the log of the number of quotes in week t by q_t , but 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 2. In empirical analysis it is common to distinguish between "expected" and "unexpected" activity, see amongst others Bessembinder and Seguin (1992), Jorion (1996) and Bjønnes, Rime, and Solheim (2005). Expected activity is supposed to reflect "normal" or "everyday" quoting or trading activity by traders, and should thus be negatively associated with volatility according to (9) since this essentially reflects the number of active traders. Unexpected activity on the other hand refers to changes in the rate at which relevant information arrives to the market and should increase volatility. The strategy that is used in order to obtain the expected and unexpected components is to interpret the fitted values of an ARMA-GARCH model as the expected component and the residual as the unexpected. In our case an ARMA(1,1) specification of Δq_t with a GARCH(1,1) structure on the error terms suffices in order to obtain uncorrelated standardised residuals and uncorrelated squared standardised residuals. The model and estimation output is contained in table 3. The expected values are then computed by generating fitted values of q_t (not of Δq_t) and are denoted \hat{q}_t . The unexpected values are

defined as $q_t - \hat{q}_t$. It has been argued that such a strategy might result in a so-called "generated regressor bias"—see for example Pagan (1984), so we opt for an alternative strategy which yields virtually identical results. As it turns out using q_t directly instead of \hat{q}_t , and Δq_t instead of the residual, has virtually no effect on the estimates in section 3. The reason can be deduced by looking at the bottom graph of figure 2. For statistical purposes q_t is virtually identical to \hat{q}_t , and Δq_t is virtually identical to the residual (the sample correlations are 0.85 and 0.94, respectively). Summarised, then, we use q_t as our measure of the number of active traders and Δq_t as our measure of changes in the rate at which information arrives to the market. Both variables serve as explanatory variables in the modelling of volatility in section 3.

2.3 Other impact variables

Other economic variables may also influence the level of volatility and should be controlled for in empirical models. In line with the conventions introduced above lower-case means the log-transformation is applied, and upper-case means it is not. The only exceptions are the interest-rate variables, 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, and a step dummy sd_t equal to 0 before 1997 and 1 after.

The first economic variable is a measure of general currency market turbulence and is measured through EUR/USD-volatility. If $m_t = \log(\text{EUR/USD})_t$, then Δm_t denotes the weekly return of EUR/USD, M_t^w stands for weekly volatility, m_t^w is its log-counterpart, M_t^r is realised volatility and m_t^r is 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 volatility. If the log of the oilprice is denoted o_t , then the weekly return is Δo_t , weekly volatility is O_t^w with o_t^w as its log-counterpart, and realised volatilities are denoted O_t^r and o_t^r , respectively. We proceed similarly for the 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_t^w , x_t^w , X_t^r and x_t^r . 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_t^w , u_t^w , U_t^r and u_t^r .

The interest-rate variables that are included are constructed using the main policy interest rate variable of the Norwegian central bank. We do not use market interest-rates because this produces interest-rate based measures that are substantially intercorrelated with q_t and sd_t , with the consequence that inference results are affected. The interest-rate variables reflect two important regime changes that took place over the period in question. As the current central bank governor assumed the position in 1999, the bank switched from exchange rate stabilisation to "partial" inflation targeting. However, a full mandate to target inflation was not given before 29 march 2001, when the Ministry of Finance instructed the bank to target an inflation of 2.5%. So an interesting question is whether policy interest rate changes contributed differently to exchange rate volatility in the partial and full inflation targeting periods, respectively.¹ This motivates the construction of our

¹Prior to 1999 central bank interest rates were very stable, at least from late 1993 until late 1996, and

interest rate variables. 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.

3 Models and empirical results

In this section, we present the econometric models of volatility and their estimated versions, together with interpretations. In subsection 3.1 we use linear regression models for the log of our volatility measures defined in subsection 2.1, hence the expression "log-linear analysis". In subsection 3.2 we use EGARCH models. Of these two our main focus is on the results of the log-linear analysis, and the motivation for the EGARCH analysis is that it serves as a point of comparison since both frameworks model volatility in logs.

3.1 Log-linear analysis

In this part we report the estimates of six specifications:

$$v_t^w = b_0 + b_1 v_{t-1}^w + b_2 v_{t-2}^w + b_3 v_{t-3}^w + b_{14} id_t + b_{15} sd_t + e_t \quad (10)$$

$$v_t^w = b_0 + b_1 v_{t-1}^w + b_2 v_{t-2}^w + b_3 v_{t-3}^w + b_6 q_t + b_7 \Delta q_t + b_{14} id_t + b_{15} sd_t + e_t \quad (11)$$

$$v_t^w = b_0 + b_1 v_{t-1}^w + b_2 v_{t-2}^w + b_3 v_{t-3}^w + b_6 q_t + b_7 \Delta q_t \\ + b_8 m_t^w + b_9 o_t^w + b_{10} x_t^w + b_{11} u_t^w + b_{12} f_t^a + b_{13} f_t^b + b_{14} id_t + b_{15} sd_t + e_t \quad (12)$$

$$v_t^r = b_0 + b_1 v_{t-1}^r + b_2 v_{t-2}^r + b_3 v_{t-3}^r + b_4 v_{t-4}^r + b_5 v_{t-5}^r + b_{14} id_t + b_{15} sd_t + b_{16} e_{t-1} + e_t \quad (13)$$

$$v_t^r = b_0 + b_1 v_{t-1}^r + b_2 v_{t-2}^r + b_3 v_{t-3}^r + b_4 v_{t-4}^r + b_5 v_{t-5}^r + b_6 q_t + b_7 \Delta q_t \\ + b_{14} id_t + b_{15} sd_t + b_{16} e_{t-1} + e_t \quad (14)$$

$$v_t^r = b_0 + b_1 v_{t-1}^r + b_2 v_{t-2}^r + b_3 v_{t-3}^r + b_4 v_{t-4}^r + b_5 v_{t-5}^r + b_6 q_t + b_7 \Delta q_t \\ + b_8 m_t^r + b_9 o_t^r + b_{10} x_t^r + b_{11} u_t^r + b_{12} f_t^a + b_{13} f_t^b + b_{14} id_t + b_{15} sd_t + b_{16} e_{t-1} + e_t. \quad (15)$$

The first three have log of weekly volatility v_t^w as left-side variable and the latter three have log of realised volatility v_t^r as left-side variable. In each triple the first specification

it was less clear to the market what role the interest rate actually had.

consists of an autoregression augmented with the Russian moratorium dummy id_t and the step dummy sd_t for the lasting shift upwards in financial volatility in 1997. In the realised case a moving average (MA) term e_{t-1} is also added for reasons to be explained below. The second specification in each triple consists of the first together with the quote variables, and the third specification is an autoregression augmented by all the economic variables. The estimates of the first triple is contained in table 4, whereas the estimates of the second triple is contained in table 5. The results can be summarised in five points.

1. *Information arrival.* The estimated impacts of changes in the rate at which information arrives to the market Δq_t carry the hypothesised positive sign and are significant at all conventional levels. In the weekly case the estimates are virtually identical and equal to about 1, whereas in the realised case the coefficient drops from 0.88 to 0.73 as other variables are added. Summarised, then, the results support the idea that exchange rate variability increases with the number of information events, and the results suggest the impact is higher for weekly than for realised volatility. There might be a small caveat in the realised case though. The MA(1) term e_{t-1} is needed in (14) and (15) in order to account for residual serial correlation at lag 1 induced by the inclusion of Δq_t . We have been unsuccessful sofar in identifying why Δq_t induces this serial correlation, and excluding Δq_t from (15) also removes the signs of heteroscedasticity indicated by White's (1980) test with cross products in the sense that the p -value increases from 10% to 24%.

2. *Number of traders.* The hypothesised effect of an increase in the number of traders as measured by q_t is negative, but in all the four specifications in which it is included it does come out positive. Moreover, it is significantly positive at 5% in both realised specifications. Figure 3 aims at throwing light on why we obtain these unanticipated results and contains recursive OLS estimates of the impact of q_t with approximate 95% confidence bands. In the weekly case the value starts out negative, but then turns positive and stays so for the rest of the sample. However, it descends steadily towards the end. In the realised case, the value is positive all the time but for a short interval in the beginning, and exhibits the same downwards tendency towards the end as in the weekly case. The recursive estimates are more stable here though than in the weekly case. All in all, then, the recursive graphs suggest the impact of q_t over the sample is positive rather than negative, and this may be explained in one of two ways: Either our measure of number of traders is faulty, or the impact of number of traders is positive rather than negative.

3. *Volatility persistence.* The autoregressions (10) and (13) were constructed according to a simple-to-general philosophy. The starting equation was volatility regressed on a constant, volatility lagged once, the step dummy sd_t and the impulse dummy id_t , and then lags of volatility were added until two properties were satisfied in the following order of importance: (i) Residuals and squared residuals were serially uncorrelated, and (ii) the coefficient in question was significantly different from zero at 5%. Interestingly such simple autoregressions are capable of producing uncorrelated and almost homoscedastic residuals in the weekly case, and uncorrelated, homoscedastic and normal residuals in the realised

case. One might suggest that normality in the log-realised specifications comes as no surprise since Andersen, Bollerslev, Diebold, and Labys (2001) have shown that taking the log of realised exchange rate volatility produces variables close to the normal. In our data, however, the Russian moratorium dummy id_t is necessary for residual normality. The step dummy sd_t is necessary for uncorrelatedness in all six specifications, but not the impulse dummy id_t . The MA(1) term in (13) is not needed for any of the residual properties but is included for comparison with (14) and (15). However, it does influence the coefficient estimates and the inference results of the lag-structure in all three specifications. Most importantly v_{t-2}^r would be significant if the MA(1) term were not included. Finally, when the lag coefficients are significant at the 10% level, then they are relatively similar across the specifications in both the weekly and realised cases. The only possible exception is the coefficient of the first lag in the realised case, which ranges from 0.41 to 0.64 across the three specifications.

4. *Policy interest rate changes.* One would expect that policy interest rate changes in the full inflation targeting period—as measured by f_t^b —increase contemporaneous volatility, whereas the hypothesised contemporaneous effect in the partial inflation period—as measured by f_t^a —is lower or at least uncertain. The results in both (12) and (15) support this since they suggest a negative but insignificant contemporaneous impact in the partial inflation targeting period, and a positive, significant and substantially larger contemporaneous impact (in absolute value) in the full inflation targeting period.

5. *Other.* The effect of general currency market volatility, as measured by m_t^w and m_t^r , is positive as expected, significant in both (12) and (15), but a little bit higher in the latter specification. The effect of oilprice volatility, as measured by o_t^w and o_t^r , is estimated to be positive in the first case and negative in the second, but the coefficients are 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 volatility is non-linear in ways not captured by our measure, see Akram (2000). With respect to the effects of Norwegian and US stock market volatility the two equations differ noteworthy. In the weekly case both x_t^w and u_t^w are estimated to have an almost identical, positive impact on volatility, and both are significant at 1%. In the realised case on the other hand everything differs. Norwegian stock market volatility x_t^r is estimated to have a positive and significant (at 10%) impact—albeit somewhat smaller than in the weekly case, whereas US stock market volatility u_t^r is estimated to have an insignificant negative impact.

In order to study the evolution of the impact of Δq_t free from any influence of (statistically) redundant regressors, we employ a general-to-specific (GETS) approach to derive more parsimonious specifications. In this way we reduce the possible reasons for changes in the evolution of the estimates. In a nutshell GETS proceeds in three steps. First, formulate a general model. Second, simplify the general model sequentially while tracking

the residual properties at each step. Finally, test the resulting model against the general starting model. See Hendry (1995), Hendry and Krolzig (2001), Mizon (1995) and Gilbert (1986) for more extensive and rigorous expositions of the GETS approach. In our case we posited (12) and (15) without the MA(1) term as general models, and it should be noted that a GETS "purist" would probably oppose to the use of the second specification as a starting model, since it exhibits residual serial correlation. Then we tested hypotheses regarding the parameters sequentially with a Wald-test (these tests are not reported), where at each step the simpler model was posited as null. In the weekly case we used heteroscedasticity consistent standard errors of the White (1980) type, and in the realised case we used heteroscedasticity and autocorrelation consistent standard errors of the Newey and West (1987) type. Our final models are not rejected in favour of the general starting models when all the restrictions are tested jointly, their estimates are contained in table 6, and their specifications are

$$\hat{v}_t^w = b_2(v_{t-2}^w + v_{t-3}^w) + b_7\Delta q_t + b_8m_t^w + b_{10}(x_t^w + u_t^w) + b_{13}f_t^b + b_{14}id_t + b_{15}sd_t \quad (16)$$

$$\hat{v}_t^r = b_1v_{t-1}^r + b_2(v_{t-2}^r + v_{t-3}^r + v_{t-5}^r) + b_7\Delta q_t + b_8m_t^r + b_{13}f_t^b + b_{14}id_t + b_{15}sd_t. \quad (17)$$

In both cases the estimates of the impact of Δq_t in the parsimonious specifications are close to those of the general starting specifications. In the weekly case the estimates are equal to 0.99 in the general specification and 1.00 in the specific, whereas in the realised case the estimate changes from 0.57 in the general specification (15) without the MA(1) term (not reported) to 0.56 in the parsimonious specification (17). Figure 4 contains recursive OLS estimates of the coefficients of Δq_t in the parsimonious specifications. They are relatively stable over the sample, but admittedly we do not test this formally. Also, the estimates seems to be more stable in the realised case than in the weekly, in the sense that the difference between the maximum and minimum values is larger in the weekly case ($1.66 - 0.67 = 0.99$) than in the realised ($0.98 - 0.47 = 0.51$). Both graphs appear to be trending downward for most of the sample, the exception being towards the end in the weekly case, and in both graphs there seems to be a distinct shift downwards as the change to partial inflation targeting takes place in the beginning of 1999. One should be careful however in attributing the shift to the change in regime without further investigation. Indeed, another possible reason is the transition to the euro, since Δq_t attains both its maximum and minimum in the first weeks of 1999.

3.2 EGARCH analysis

The estimates of the three EGARCH specifications which we report have all equal mean-specification $r_t/\hat{\sigma}_r = \mu + e_t = \mu + \sigma_t z_t$, where $r_t = \log(S_t/S_{t-1})$ is the weekly return, $\hat{\sigma}_r = 0.007615$ is the sample standard deviation of the returns, and where $\{z_t\}_{t=1,572}$ is an IID sequence of $N(0, 1)$ variables. For exchange rates it is also common to include an AR(1) term in the mean-equation in order to account for the possibility of negative serial correlation in the returns. In our data however there are signs that this term induces serial correlation in either the standardised residuals or in the squared standardised

residuals or in both. So we do not include it in the specifications reported here. The three EGARCH specifications can be considered as the ARCH counterparts of the weekly log-linear equations, that is, equations (10) - (12), and their log-variance specifications are

$$\log \sigma_t^2 = \alpha_0 + \alpha_1 \left| \frac{e_{t-1}}{\sigma_{t-1}} \right| + \gamma_1 \frac{e_{t-1}}{\sigma_{t-1}} + \beta_1 \log \sigma_{t-1}^2 + c_{11} id_t + c_{12} sd_t \quad (18)$$

$$\log \sigma_t^2 = \alpha_0 + \alpha_1 \left| \frac{e_{t-1}}{\sigma_{t-1}} \right| + \gamma_1 \frac{e_{t-1}}{\sigma_{t-1}} + \beta_1 \log \sigma_{t-1}^2 + c_1 q_t^* + c_2 \Delta q_t^* + c_{11} id_t + c_{12} sd_t \quad (19)$$

$$\begin{aligned} \log \sigma_t^2 = \alpha_0 + \alpha_1 \left| \frac{e_{t-1}}{\sigma_{t-1}} \right| + \gamma_1 \frac{e_{t-1}}{\sigma_{t-1}} + \beta_1 \log \sigma_{t-1}^2 + c_1 q_t^* + c_2 \Delta q_t^* + c_3 m_t^{f*} + c_4 o_t^{f*} \\ + c_5 x_t^{f*} + c_6 u_t^{f*} + c_7 f_t^a + c_8 f_{t-1}^a + c_9 f_t^b + c_{10} f_{t-1}^b + c_{11} id_t + c_{12} sd_t. \end{aligned} \quad (20)$$

Specification (18) is an EGARCH(1,1) with the Russian moratorium dummy id_t and the step dummy sd_t as only regressors, (19) is an EGARCH(1,1) augmented with the quote variables and the dummies, and (20) is an EGARCH(1,1) with all the economic variables as regressors. Note that * as superscript means the variable has been divided by its sample standard deviation. Specifications (18) - (20) are analogous to the ARCH-specifications in Lamoureux and Lastrapes (1990), but note that our results are not directly comparable to theirs since our measure of information intensity Δq_t does not exhibit strong positive serial correlation (in fact, our measure Δq_t exhibits weak negative serial correlation). Strong positive serial correlation is an important assumption for their conclusions.

The estimates of (18) - (20) are contained in table 7 and are relatively similar significance-wise to the results of the weekly log-linear analysis above, that is, to the estimates of (10) - (12). Note however that the magnitudes of the coefficient estimates are not directly comparable since the variables are scaled differently. The most important similarity is that the coefficient of Δq_t^* is positive and significant in both (19) and (20), and that the coefficient estimates are almost identical in (19) and (20). Another important similarity is that the measure of number of traders q_t^* is insignificant in the two EGARCH specifications in which it is included. There are three minor differences in the inference results compared with the weekly log-linear analysis. The first is that the measure of US stock market volatility u_t^* is significant at 9% in the EGARCH specification (20) containing all the variables, whereas it is significant at 1% in the weekly log-linear counterpart (12). The second minor difference is that in the EGARCH case the impacts of x_t^{w*} and u_t^{w*} respectively are not so similar as in the weekly case. Finally, the step dummy sd_t is not significant in the EGARCH specification that only contains the dummies as economic variables, whereas it is in its weekly counterpart.

There are also some parameters particular to the EGARCH setup that merit attention. The news term $\left| \frac{e_{t-1}}{\sigma_{t-1}} \right|$ is estimated to be positive as expected and reasonably similar across the three specifications, but its significance is at the borderline since the two-sided p -values range from 7% to 11%. The impact of the asymmetry term $\frac{e_{t-1}}{\sigma_{t-1}}$ is not significant in any

of the equations at conventional significance levels, which suggest no (detectable) leverage nor asymmetry as is usually found for exchange rate data. Persistence is high as suggested by the estimated impact of the autoregressive term $\log \sigma_{t-1}^2$ since it is 0.91 in (18), but it drops to 0.79 when the quote variables are included, and then to 0.59 when the rest of the economic variables are included, though it remains quite significant in all cases. Finally, the standardised residuals are substantially closer to the normal distribution in (20) compared with the other two EGARCH specifications.

4 Conclusions

Our study of weekly Norwegian exchange rate volatility sheds new light on the mixture of distribution hypothesis in several ways. We find that the impact of changes in the number of information events is positive and statistically significant within two different frameworks, that the impact is relatively stable across three different exchange rate regimes for both weekly and realised volatility, and that the estimated impacts are relatively similar in both cases. One might have expected that the effect of changes in the number of information events 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 previous to the full inflation targeting regime. In our data however there are no clear breaks, shifts upwards nor trends following the points of regime change. Moreover, our results do not support the hypothesis that an increase in the number of traders reduces volatility. Finally, we have shown that simply applying the log of volatility can improve inference and remove undesirable residual properties. In particular, OLS-estimated autoregressions of the log of volatility are capable of producing uncorrelated and (almost) homoscedastic residuals, and the in the log of realised volatility case the residuals are also Gaussian.

Our study suggests at least two avenues for future research. First, our results suggest there is no impact of the number of traders on exchange rate volatility, but this might be due to our measure being unsatisfactory. So the first avenue of research is to reconsider the hypothesis with a different approach. The second avenue of future research is to uncover why applying the log works so well. Pantula (1986), Geweke (1986) and Nelson (1991) proposed that volatility should be analysed in logs in order to ensure nonnegativity. In our case the motivation stems from unsatisfactory residual properties and fragile inference-results. Before we switched to the log-linear framework we struggled only to obtain uncorrelated residuals within the ARCH, ARMA and linear frameworks, and when we did attain satisfactory residual properties the results turned out to be very sensitive to small changes in the specification. With the log-transformation, however, results are robust across a number of specifications. So the second avenue of further research consists of understanding better why the log works. Is it due to particularities in our data? For example, is it due to our—in financial contexts—relatively small sample of 573 observations? Is it due to influential observations? Is it due to both? Or is it just due to the simple fact that applying the log is believed to lead to faster convergence towards the asymptotic

theory which our residual tests rely upon? Further application of log-linear analysis is necessary in order to answer these questions, and to verify the possible usefulness of the log-linear framework more generally.

Acknowledgements

We are indebted to various people for useful comments and suggestions at different stages, including Farooq Akram, Sébastien Laurent, an anonymous referee, participants at the JAE conference in Venice June 2005, participants at the poster session following the joint CORE-ECARES-KUL seminar at Brussels in April 2005, participants at the MICFINMA summer school in Konstanz in June 2004, and participants at the bi-annual doctoral workshop in economics at Université catholique de Louvain (Louvain la Neuve) in May 2004. The usual disclaimer about remaining errors and interpretations being our own applies of course. This work was supported by the European Community's Human Potential Programme under contract HPRN-CT-2002-00232, Microstructure of Financial Markets in Europe, and by the Belgian Program on Interuniversity Poles of Attraction initiated by the Belgian State, Prime Minister's Office, Science Policy Programming. The third author would like to thank Finansmarkedsfondet (the Norwegian Financial Market Fund) and Lånkassen (the Norwegian government's student funding scheme) for financial support at different stages, and the hospitality of the Department of Economics at the University of Oslo and the Norwegian Central Bank in which part of the research was carried out.

References

- AKRAM, Q. F. (2000): "When does the oil price affect the Norwegian exchange rate?," Working Paper 2000/8, Oslo: The Central Bank of Norway.
- 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.
- 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.
- BESSEMBINDER, H., AND P. SEGUIN (1992): "Futures-trading activity and stock price volatility," *Journal of Finance*, 47, 2015–2034.
- BJØNNES, G., D. RIME, AND H. SOLHEIM (2005): "Volume and volatility in the FX market: Does it matter who you are?," in *Exchange Rate Modelling: Where Do We Stand?*, ed. by P. De Grauwe. MIT Press, Cambridge, MA.
- BOLLERSLEV, T., AND I. DOMOWITZ (1993): "Trading Patterns and Prices in the Inter-bank 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.
- CLARK, P. (1973): “A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices,” *Econometrica*, 41, 135–155.
- 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.
- GALATI, G. (2003): “Trading volume, volatility and spreads in foreign exchange markets: evidence from emerging market countries,” BIS Working Paper.
- GEWEKE, J. (1986): “Modelling the Persistence of Conditional Variance: A Comment,” *Econometric Reviews*, 5, 57–61.
- GILBERT, C. L. (1986): “Professor Hendry’s Econometric Methodology,” *Oxford Bulletin of Economics and Statistics*, 48, 283–307.
- GIOT, P. (2003): “The Asian financial crisis: the start of a regime switch in volatility,” CORE Discussion Paper 2003/78.
- 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 *The Foreign Exchange Market*, ed. by C. Goodhart. MacMillan Press Ltd., London.
- GRAMMATIKOS, T., AND A. SAUNDERS (1986): “Futures Price Variability: A Test of Maturity and Volume Effects,” *Journal of Business*, 59, 319–330.
- HENDRY, D. F. (1995): *Dynamic Econometrics*. Oxford University Press, Oxford.
- HENDRY, D. F., AND H.-M. KROLZIG (2001): *Automatic Econometric Model Selection using PcGets*. Timberlake Consultants Press, London.
- JARQUE, C., AND A. BERA (1980): “Efficient Tests for Normality, Homoskedasticity, and Serial Independence of Regression Residuals,” *Economics Letters*, 6, 255–259.
- JORION, P. (1996): “Risk and Turnover in the Foreign Exchange Market,” in *The Microstructure of Foreign Exchange Markets*, ed. by J. Frankel et al. University of Chicago Press, Chicago.
- KARPOFF, J. (1987): “The relation between price changes and trading volume: A survey,” *Journal of Financial and Quantitative Analysis*, 22, 109–126.

- LAMOUREUX, C. G., AND W. D. LASTRAPES (1990): "Heteroscedasticity in Stock Return Data: Volume versus GARCH Effects," *Journal of Finance*, pp. 221–229.
- LJUNG, G., AND G. BOX (1979): "On a Measure of Lack of Fit in Time Series Models," *Biometrika*, 66, 265–270.
- MELVIN, M., AND Y. XIXI (2000): "Public Information Arrival, Exchange Rate Volatility, and Quote Frequency," *The Economic Journal*, 110, 644–661.
- MIZON, G. (1995): "Progressive Modeling of Macroeconomic Time Series: The LSE Methodology," in *Macroeconometrics. Developments, Tensions and Prospects*, ed. by K. D. Hoover. Kluwer Academic Publishers.
- NELSON, D. B. (1991): "Conditional Heteroscedasticity in Asset Returns: A New Approach," *Econometrica*, 51, 485–505.
- NEWBY, 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.
- TAUCHEN, G., AND M. PITTS (1983): "The Price Variability-Volume Relationship on Speculative Markets," *Econometrica*, 51, 485–505.
- WHITE, H. (1980): "A Heteroskedasticity-Consistent Covariance Matrix and a Direct Test for Heteroskedasticity," *Econometrica*, 48, 817–838.

Appendix: Data sources and transformations

The data transformations were undertaken in Ox 3.2 and EViews 4.1.

$S_{n(t)}$	$n(t) = 1(t), 2(t), \dots, N(t)$, where $S_{1(t)}$ 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 . Before 1.1.1999 the BID NOK/1EUR rate is obtained by the formula $\text{BID NOK}/100\text{DEM} \times 0.0195583$, where 0.0195583 is the official DEM/1EUR conversion rate $1.95583 \text{ DEM} = 1 \text{ EUR}$ divided by 100. The first untransformed observation is the opening value of BID NOK/100DEM on Wednesday 6.1.1993 and the last is the BID NOK/1EUR closing value on Friday 26.12.2003. The source of the BID NOK/100DEM series is Olsen and the source of the BID NOK/1EUR series is Reuters.
S_t	$S_{N(t)}$, the last closing value of week t
r_t	$\log S_t - \log S_{t-1}$
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 $n = 1(t), 2(t), \dots, N(t)$ and $1(t) - 1 = N(t - 1)$
v_t^r	$\log V_t^r$
$M_{n(t)}$	$n(t) = 1(t), 2(t), \dots, N(t)$, where $M_{1(t)}$ is the first BID USD/EUR opening exchange rate of week t , $M_{2(t)}$ is the first closing rate, $M_{3(t)}$ is the second opening rate, and so on, with $M_{N(t)}$ denoting the last closing rate of week t . Before 1.1.1999 the BID USD/EUR rate is obtained with the formula $1.95583/(\text{BID DEM}/\text{USD})$. The first untransformed observation is the opening value of BID DEM/USD on Wednesday 6.1.1993 and the last is the closing value on Friday 30.12.2003. The source of the BID DEM/USD and BID USD/EUR series is Reuters.
M_t	$M_{N(t)}$, the last closing value of week t
m_t	$\log M_t$

M_t^w	$\{\{\log[M_t + I(M_t = M_{t-1}) \times k_t] - \log(M_{t-1})\} \times 100\}^2$. $I(M_t = M_{t-1})$ is an indicator function equal to 1 if $M_t = M_{t-1}$ and 0 otherwise, and k_t is a positive number that ensures the log-transformation is not performed on a zero-value. $M_t = M_{t-1}$ occurs for $t = 23/2/1996$, $t = 19/12/1997$ and $t = 20/2/1998$, and the value of k_t was set on a case to case basis depending on the number of decimals in the original, untransformed dataserie. Specifically the values of k_t were set to 0.00009, 0.0009 and 0.00009, respectively.
m_t^w	$\log M_t^w$
M_t^r	$\sum_n [\log(M_n/M_{n-1}) \times 100]^2$, where $n = 1(t), 2(t), \dots, N(t)$ and $1(t)-1 = N(t-1)$
m_t^r	$\log M_t^r$
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	$\log Q_t$. Note that 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 was 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.
$O_{i(t)}$	$n(t) = 2(t), 4(t), \dots, N(t)$, where $O_{2(t)}$ is the first closing value of the Brent Blend spot oilprice in USD per barrel in week t , $O_{4(t)}$ is the second closing value of week t , and so on, with $O_{n(t)}$ denoting the last closing value of week t . The untransformed series is Bank of Norway database series D2001712, which is based on Telerate page 8891 at 16.00.
O_t	$O_{N(t)}$, the last closing value in week t
o_t	$\log O_t$
O_t^w	$\{\log[O_t + I(O_t = O_{t-1}) \times 0.009] - \log(O_{t-1})\}^2$. $I(O_t = O_{t-1})$ 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	$\log O_t^w$
O_t^r	$\sum_n [\log(O_n/O_{n-2})]^2$, where $n = 2(t), 4(t), \dots, N(t)$ and $2(t) - 2 = N(t - 1)$

o_t^r	$\log O_t^r$
$X_{n(t)}$	$n(t) = 2(t), 4(t), \dots, N(t)$, where $X_{2(t)}$ is the first closing value of the main index of the Norwegian Stock Exchange (TOTX) in week t , $X_{4(t)}$ is the second closing value, and so on, with $X_{N(t)}$ denoting the last closing value of week t . The source of the daily untransformed series is EcoWin series ew:nor15565.
X_t	$X_{N(t)}$, the last closing value in week t
x_t	$\log X_t$
X_t^w	$[\log(X_t/X_{t-1})]^2$. $X_t = X_{t-1}$ does not occur for this series.
x_t^w	$\log X_t^w$
X_t^r	$\sum_n [\log(X_n/X_{n-2})]^2$, where $n = 2(t), 4(t), \dots, N(t)$ and $2(t) - 2 = N(t - 1)$
x_t^r	$\log X_t^r$
$U_{n(t)}$	$n(t) = 2(t), 4(t), \dots, N(t)$, where $U_{2(t)}$ is the first closing value in USD of the composite index of the New York Stock Exchange (the NYSE index) in week t , $U_{4(t)}$ is the second closing value, and so on, with $U_{N(t)}$ denoting the last closing value of week t . The source of the daily untransformed series is EcoWin series ew:usa15540.
U_t	$U_{N(t)}$, the last closing value in week t
U_t^w	$[\log(U_t/U_{t-1})]^2$. $U_t = U_{t-1}$ does not occur for this series.
u_t^w	$\log U_t^w$
U_t^r	$\sum_n [\log(U_n/U_{n-2})]^2$, where $n = 2(t), 4(t), \dots, N(t)$ and $2(t) - 2 = N(t - 1)$
u_t^r	$\log U_t^r$
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 Bank of Norway'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.

Table 1: Summary of empirical studies that investigate the impact of information intensity on exchange rate volatility.

Publication	Data	Period	Supportive of MDH?
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, Ben Omrane, and Giot (2005)	Intradaily EUR/USD quotes, quoting frequency and Reuters' news-alert screens	15/5 2001- 14/11 2001	Yes
Bjønnes, Rime, and Solheim (2005)	Daily SEK/EUR quotes and transaction volume	1995-2002	Yes

Table 2: Descriptive statistics of selected variables

	S_t	ΔS_t	$ \Delta S_t $	V_t^w	v_t^w	V_t^r	v_t^r	q_t
<i>Mean</i>	8.208	0.000	0.005	0.578	-2.489	0.488	-1.659	7.5115
<i>Median</i>	8.224	0.000	0.003	0.120	-2.121	0.209	-1.567	7.5192
<i>Max.</i>	9.063	0.044	0.044	19.365	2.963	16.033	2.775	9.1363
<i>Min.</i>	7.244	-0.035	0.000	0.000	-10.757	0.004	-5.497	5.6131
<i>St.dev.</i>	0.352	0.008	0.006	1.679	2.447	1.035	1.413	0.5739
<i>Skew.</i>	0.025	0.774	2.886	7.036	-0.845	8.418	-0.102	-0.3287
<i>Kurt.</i>	2.174	9.399	14.991	65.150	3.884	107.193	2.793	3.4512
<i>Obs.</i>	573	572	572	573	573	573	573	573

	Δq_t	$ \Delta q_t $	M_t	Δm_t	$ \Delta m_t $	m_t^w	f_t^a	f_t^b
<i>Mean</i>	0.004	0.226	1.115	0.000	0.011	-0.556	0.007	0.010
<i>Median</i>	-0.003	0.145	1.124	-0.001	0.009	-0.190	0.000	0.000
<i>Max.</i>	2.141	2.141	1.429	0.047	0.053	3.325	0.500	1.000
<i>Min.</i>	-1.278	0.000	0.838	-0.053	0.000	-9.150	0.000	0.000
<i>St.dev.</i>	0.339	0.252	0.150	0.014	0.009	2.018	0.057	0.081
<i>Skew.</i>	0.530	2.722	0.025	0.099	1.228	-1.015	8.271	9.136
<i>Kurt.</i>	9.008	14.119	2.174	3.377	5.023	4.181	70.632	94.602
<i>Obs.</i>	572	572	573	572	572	573	573	573

Note: Some zero-values are due to rounding and the variables are explained in subsection 2.3 and in the appendix.

Table 3: ARMA-GARCH model of Δq_t

$$\Delta q_t = b_0 + b_1 \Delta q_{t-1} + b_2 e_{t-1} + e_t,$$

$$e_t = \sigma_t z_t, \quad \sigma_t^2 = \alpha_0 + \alpha_1 e_{t-1}^2 + \beta_1 \sigma_{t-1}^2$$

<i>Parameter</i>			<i>Diagnostics</i>	
	<i>Est.</i>	<i>Pval.</i>		
b_0	0.004	0.13	R^2	0.19
b_1	0.569	0.00	$Q(10)$	11.03
b_2	-0.910	0.00	$ARCH_{1-10}$	0.29
α_0	0.034	0.02	JB	691.45
α_1	0.299	0.00	<i>Obs.</i>	571
β_1	0.368	0.06		

Notes: Computations are in EViews 5.1 and estimates are ML with heteroscedasticity consistent 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 *F*-form of the Lagrange-multipier test for serial correlation in the squared standardised residuals up to lag 10, *Skew.* is the skewness of the standardised residuals, *Kurt.* is the kurtosis of the standardised residuals, and *JB* is the Jarque and Bera (1980) test for non-normality of the standardised residuals.

Table 4: Regressions of log of weekly NOK/EUR volatility

	(10)		(11)		(12)	
	<i>Est.</i>	<i>Pval.</i>	<i>Est.</i>	<i>Pval.</i>	<i>Est.</i>	<i>Pval.</i>
<i>const.</i>	-2.917	0.00	-3.887	0.01	-0.660	0.67
v_{t-1}^w	0.019	0.64	0.023	0.59	0.007	0.87
v_{t-2}^w	0.077	0.04	0.078	0.04	0.076	0.05
v_{t-3}^w	0.096	0.03	0.105	0.02	0.099	0.02
q_t			0.141	0.46	0.029	0.88
Δq_t			0.995	0.00	0.986	0.00
m_t^w					0.139	0.00
o_t^w					0.015	0.74
x_t^w					0.123	0.01
u_t^w					0.112	0.01
f_t^a					-0.116	0.92
f_t^b					3.545	0.00
<i>id</i>	4.745	0.00	4.400	0.00	3.563	0.00
<i>sd</i> _{<i>t</i>}	1.396	0.00	1.306	0.00	1.037	0.00
R^2	0.14		0.16		0.21	
AR_{1-10}	0.34	0.97	0.81	0.62	0.32	0.98
$ARCH_{1-10}$	0.99	0.45	0.78	0.64	0.56	0.84
<i>Het.</i>	9.42	0.31	13.40	0.34	24.81	0.42
<i>Hetero.</i>	21.89	0.08	45.40	0.01	79.22	0.63
<i>JB</i>	120.94	0.00	117.12	0.00	146.16	0.00
<i>Obs.</i>	570		570		570	

Notes: Computations are in EViews 5.1 and estimates are OLS with heteroscedasticity consistent 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 *F*-form of the Lagrange-multiplier test for serially correlated residuals up to lag 10, $ARCH_{1-10}$ is the *F*-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, *Skew.* is the skewness of the residuals, *Kurt.* is the kurtosis of the residuals, and *JB* is the Jarque and Bera (1980) test for non-normality in the residuals.

Table 5: Regressions of log of realised NOK/EUR volatility

	(13)		(14)		(15)	
	<i>Est.</i>	<i>Pval.</i>	<i>Est.</i>	<i>Pval.</i>	<i>Est.</i>	<i>Pval.</i>
<i>const.</i>	-1.012	0.00	-1.690	0.00	-1.916	0.02
v_{t-1}^r	0.405	0.05	0.643	0.00	0.483	0.00
v_{t-2}^r	0.078	0.29	0.014	0.81	0.047	0.36
v_{t-3}^r	0.104	0.04	0.086	0.07	0.085	0.05
v_{t-4}^r	-0.059	0.25	-0.065	0.16	-0.050	0.28
v_{t-5}^r	0.122	0.00	0.087	0.03	0.069	0.08
q_t			0.139	0.03	0.173	0.02
Δq_t			0.876	0.00	0.725	0.00
m_t^r					0.194	0.00
o_t^r					-0.021	0.62
x_t^r					0.070	0.08
u_t^r					-0.007	0.85
f_t^a					-0.256	0.63
f_t^b					1.403	0.00
<i>id</i>	4.275	0.00	3.777	0.00	3.985	0.00
<i>sd_t</i>	0.659	0.00	0.382	0.00	0.532	0.00
e_{t-1}	-0.130	0.53	-0.380	0.00	-0.238	0.03
R^2	0.53		0.57		0.60	
AR_{1-10}	0.86	0.57	1.17	0.31	0.81	0.62
$ARCH_{1-10}$	0.44	0.93	1.34	0.20	1.15	0.32
<i>Het.</i>	6.16	0.91	10.91	0.82	33.69	0.21
<i>Hetero.</i>	30.94	0.27	50.15	0.24	128.63	0.10
<i>JB</i>	3.15	0.52	0.44	0.80	0.47	0.79
<i>Obs.</i>	568		568		568	

Notes: See table 4 for details.

Table 6: Parsimonious log-linear specifications obtained by GETS analysis

	v_t^w			v_t^r	
	<i>Est.</i>	<i>Pval.</i>		<i>Est.</i>	<i>Pval.</i>
			<i>const.</i>	-2.526	0.00
$v_{t-2}^w + v_{t-3}^w$	0.095	0.00	$3v_{t-1}^r + v_{t-2}^r + v_{t-3}^r + v_{t-5}^r$	0.091	0.00
			q_t	0.256	0.00
Δq_t	0.998	0.00	Δq_t	0.561	0.00
m_t^w	0.141	0.00	m_t^r	0.206	0.00
$x_t^w + u_t^w$	0.143	0.00	x_t^r	0.075	0.07
f_t^b	3.529	0.00	f_t^b	1.635	0.00
id_t	3.445	0.00	id_t	3.878	0.00
sd_t	0.951	0.00	sd_t	0.628	0.00
R^2	0.21			0.60	
AR_{1-10}	0.39	0.95		0.81	0.62
$ARCH_{1-10}$	0.56	0.85		1.07	0.38
<i>Het.</i>	12.93	0.37		18.20	0.20
<i>Hetero.</i>	20.23	0.78		50.19	0.04
<i>JB</i>	143.16	0.00		0.57	0.75
<i>Obs.</i>	570			568	

Notes: See table 4 for details.

Table 7: EGARCH-analysis of NOK/EUR return volatility

	(18)		(19)		(20)	
	<i>Est.</i>	<i>Pval.</i>	<i>Est.</i>	<i>Pval.</i>	<i>Est.</i>	<i>Pval.</i>
<i>const.(mean)</i>	-0.025	0.37	-0.054	0.05	-0.065	0.01
<i>const.(var.)</i>	-0.303	0.11	-0.994	0.13	0.343	0.70
$ e_{t-1}/\sigma_{t-1} $	0.230	0.09	0.247	0.07	0.169	0.11
e_{t-1}/σ_{t-1}	0.005	0.95	0.085	0.21	0.084	0.18
$\log(\sigma_{t-1}^2)$	0.906	0.00	0.789	0.00	0.587	0.00
q_t^*			0.038	0.38	0.037	0.49
Δq_t^*			0.373	0.00	0.356	0.00
m_t^{w*}					0.148	0.03
o_t^{w*}					-0.057	0.29
x_t^{w*}					0.312	0.00
u_t^{w*}					0.098	0.09
f_t^a					-0.352	0.71
f_t^b					1.611	0.00
id_t	3.002	0.00	2.665	0.00	0.441	0.00
sd_t	0.151	0.19	0.329	0.03	1.552	0.01
<i>LogL.</i>	-710.16		-687.85		-656.38	
<i>Q(10)</i>	11.88	0.29	10.95	0.36	11.91	0.29
<i>ARCH</i> ₁₋₁₀	0.88	0.55	11.69	0.31	12.46	0.26
<i>JB</i>	161.00	0.00	119.73	0.00	20.65	0.00
<i>Obs.</i>	572		572		572	

Notes: See table 3 for details.

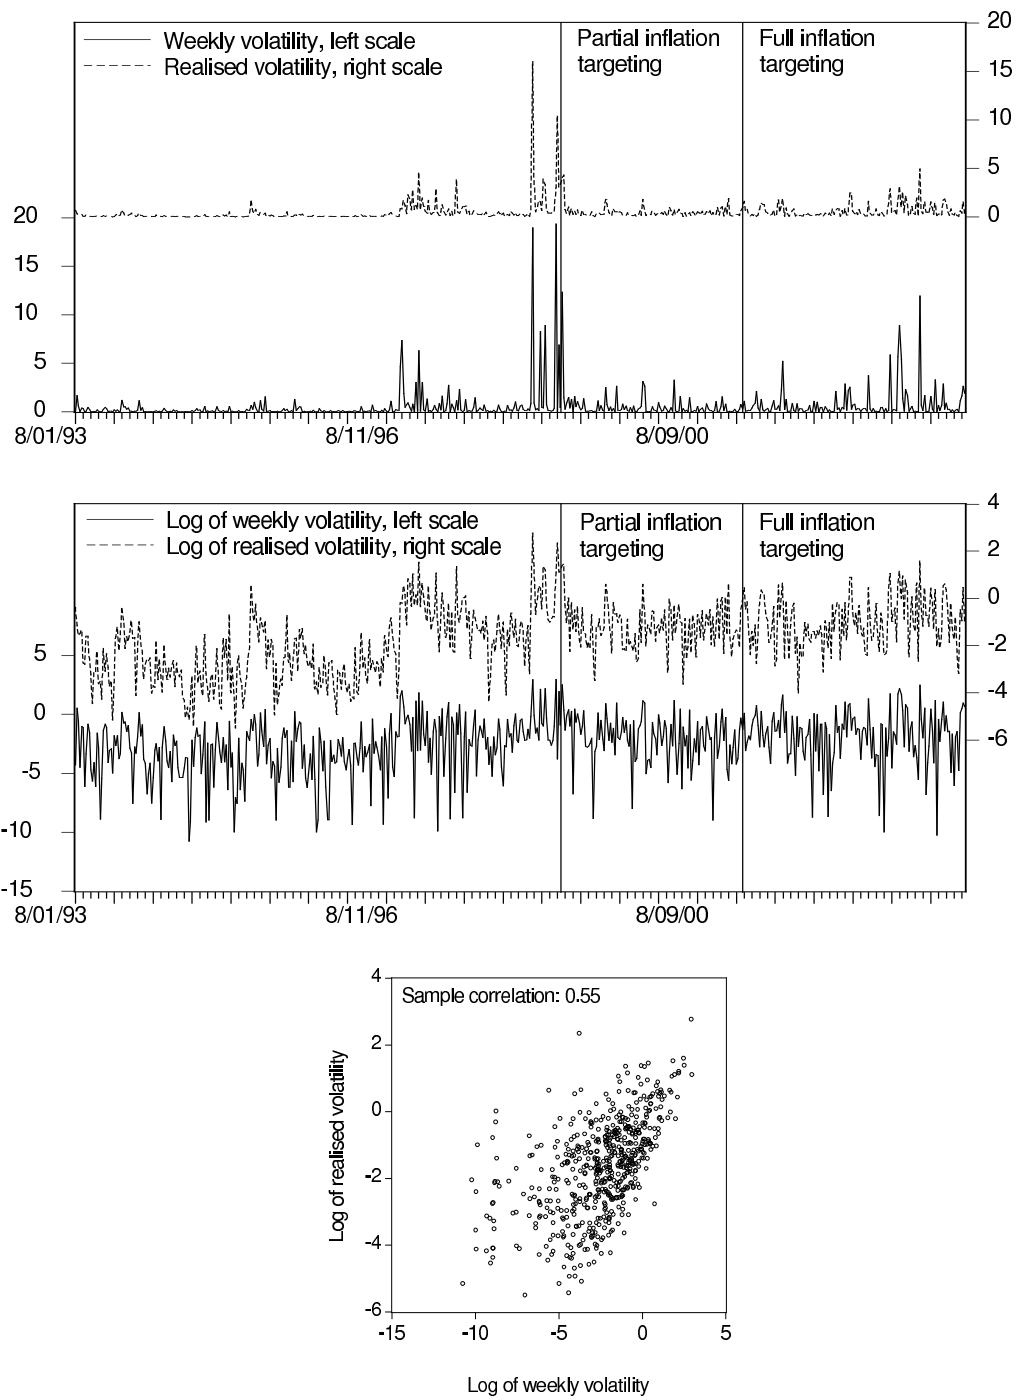


Figure 1: Weekly and realised NOK/EUR volatilities from 8 January 1993 to 26 December 2003 (NOK/DEM before 1 January 1999) in the upper graph, log of weekly and realised volatilities in the middle graph, and a scatter plot of the log-volatilities.

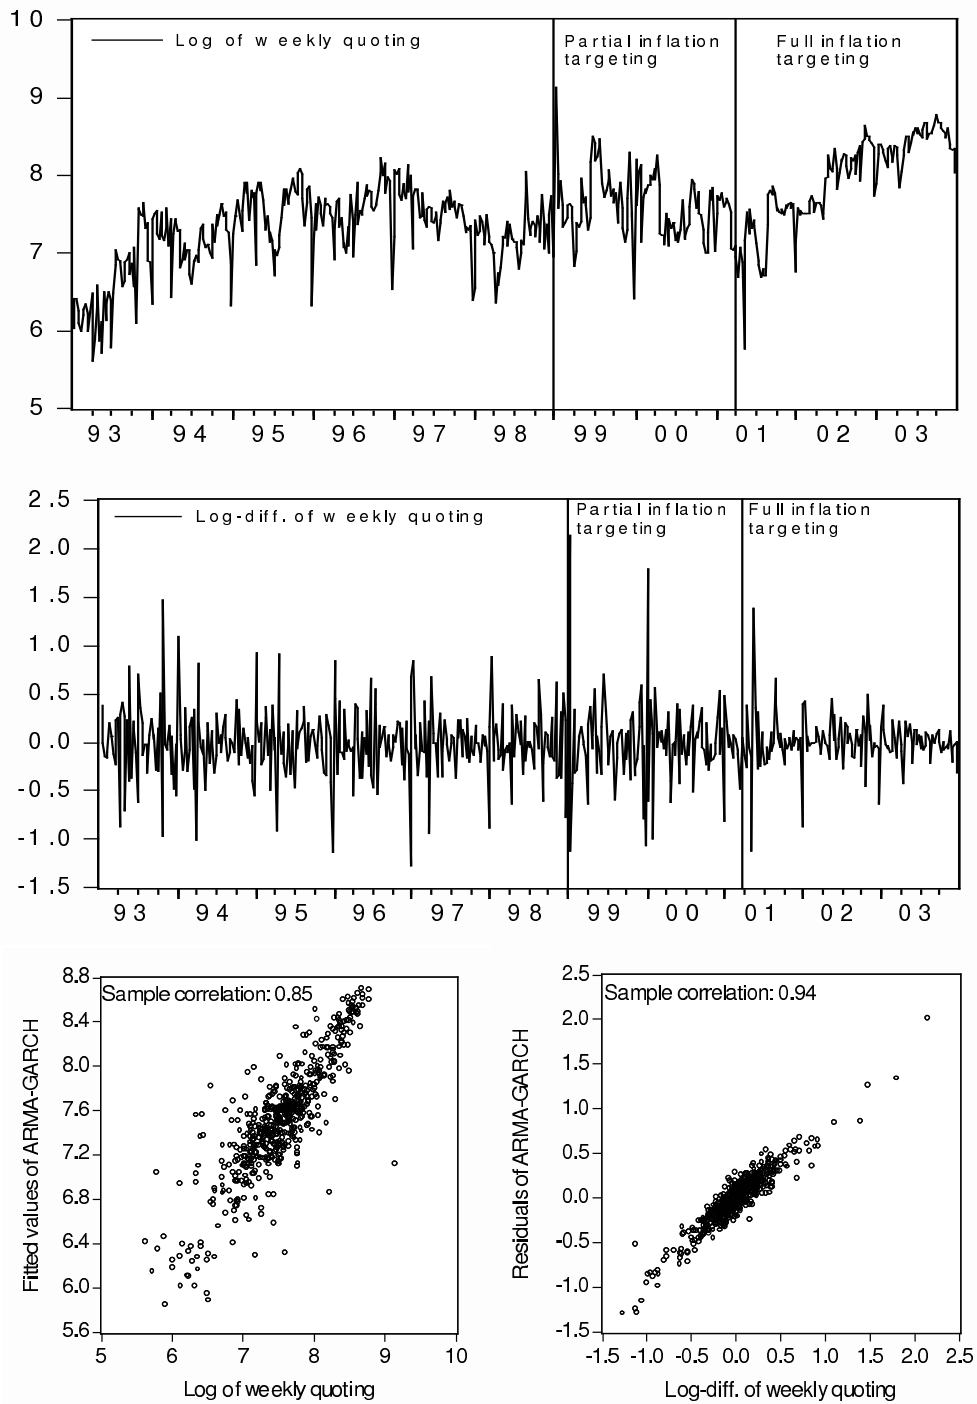


Figure 2: The log of weekly number of BID NOK/EUR quotes (BID NOK/DEM before 1999) in the upper graph, the log-difference of weekly quoting in the middle graph, and scatter plots of q_t vs. \hat{q}_t and Δq_t vs. residual in the bottom graph.




Figure 3: Stability analysis of the impact of q_t in the general unrestricted specification (12) and in the parsimonious specification (17). Computations in PcGive 10.4 with OLS and initialisation at observation number 50.

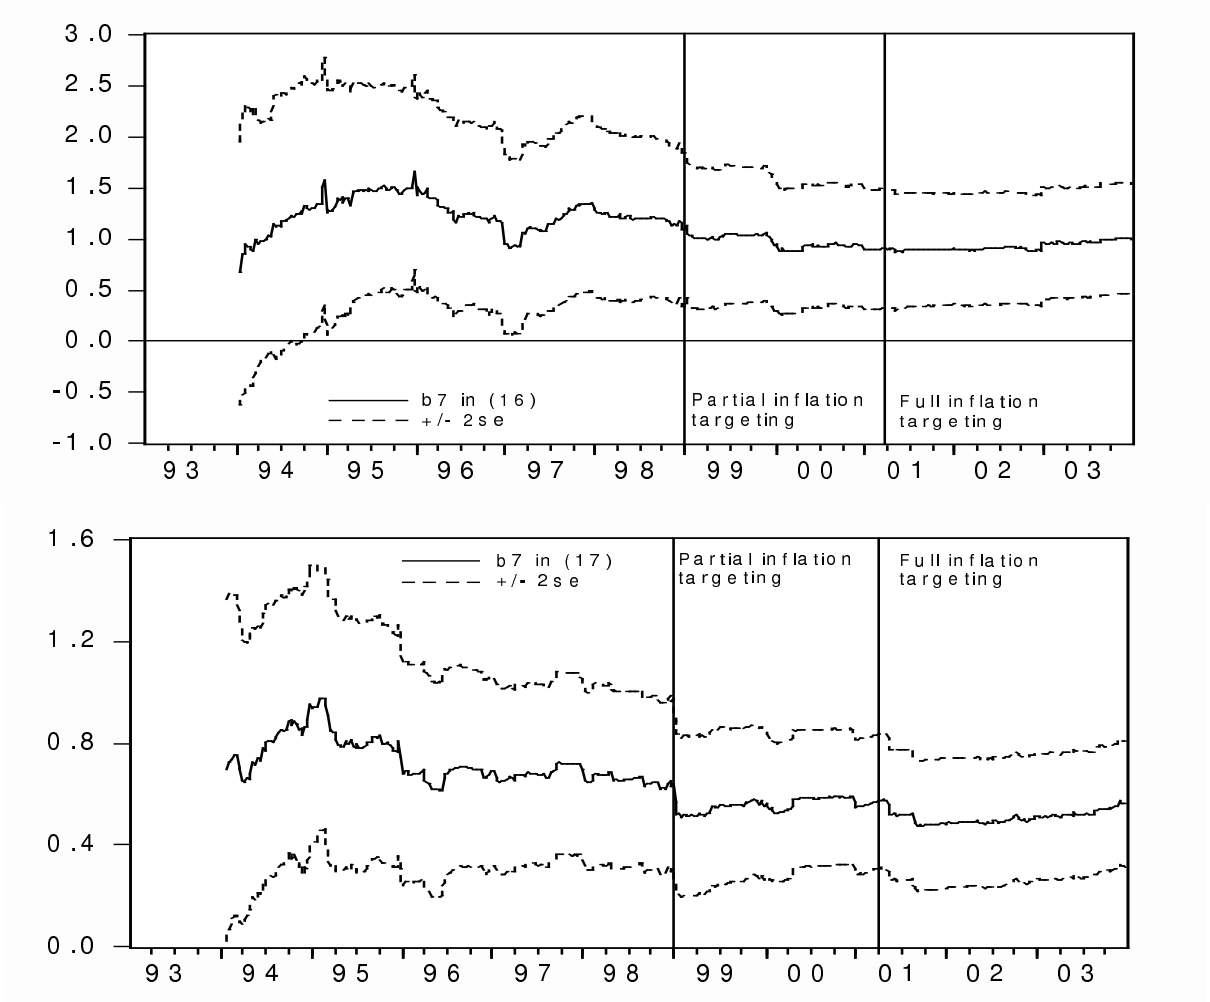


Figure 4: Recursive estimates of b_7 in the parsimonious specifications (16) and (17). Computations in PcGive 10.4 with OLS and initialisation at observation number 50.