A Time Series Model for an Exchange Rate in a Target Zone with Applications

Stefan Lundbergh

Skandia Life Insurance Company Ltd,

SE-103 50 Stockholm, Sweden

Timo Teräsvirta

Department of Economic Statistics,

Stockholm School of Economics,

P.O.Box 6501, SE-113 83 Stockholm, Sweden

Abstract

In this paper we introduce a flexible target zone model that is capable of characterizing the dynamic behaviour of an exchange rate implied by the original target zone model of Krugman (1991) and its modifications. Our framework also enables the modeller to estimate an implicit target zone if it exists. A modelling cycle consisting of specification, estimation, and evaluation stages is constructed. The model is fitted to series of daily observations of the Swedish and the Norwegian currency indices and the estimated models are evaluated.

Keywords: Autoregressive conditional heteroskedasticity, exchange rate dynamics, nonlinear modelling, smooth transition autoregression.

JEL Classification Code: C22, C52, F31.

2

1. Introduction

Over the years, target zones for exchange rates have been a reality in, for example, the Bretton Woods system, the gold standard, the Exchange Rate Mechanism (ERM) of the European Monetary System, Hungary, Scandinavian countries excluding Denmark, and a number of South American countries. In some of the latter, the zone has been a "crawling band" whose centre has been adjusted daily in mini-steps that have been known to the agents in advance. As an example, see Brooks and Revéiz (2002) for a description of such a crawling band for the Colombian peso. For more information, see Darvas (1998).

There is a vast literature about the target zones, both theoretical and empirical. For useful surveys, see Svensson (1992) and Taylor (1995). The modern theoretical literature has its starting point in the target zone model of Krugman (1991) that will be outlined below. The model has later been modified to situations in which its basic assumption, that the interventions of the central bank only occur at the boundaries of the zone, are not satisfied. Examples include Delgado and Dumas (1992) and Torres (2000). The Krugman model has been tested in different ways in many empirical contributions, and a general observation has been that it has failed to adequately characterize the movements of exchange rates in a target zone.

The Krugman model is a continuous time model, and there have been many attempts to fit it to various daily exchange rate series using the simulated method of moments: see for example Smith and Spencer (1992), de Jong (1994), Iannizzotto and Taylor (1999) and Taylor and Iannizzotto (2001). Another, discrete-time, approach (Bekaert and Gray, 1998) has been to model the conditional distribution of the exchange rate within a target zone assuming that the distribution is a truncated normal one, truncation being a result of the existence of a credible zone. In this paper we consider another discrete time model for this problem that we call the smooth transition autoregressive target zone (STARTZ) model. Our aim has been to develop a model that will allow the investigator to both consider the validity to the assumptions of Krugman's model and, at the same time, adequately characterize the dynamic behaviour of an exchange rate fluctuating within a target zone. In order to evaluate the estimated STARTZ model, something that should not be overlooked when carrying out empirical investigations, we construct a number of misspecification tests to tests its adequacy. The plan of the paper is as follows. The classical Krugman model is discussed in Section 2. The STARTZ model is defined in Section 3 and its specification and estimation are considered in Section 4. Sec-

4

tion 5 contains misspecification tests for evaluation and Section 6 an application of the model to two Nordic exchange rates. In Section 7, density forecasts from the estimated model are considered. Finally, Section 8 concludes.

2. The Krugman model and empirical applications

Consider the following continuous-time model for the exchange rate s :

$$s = f + \alpha \mathsf{E}\{\frac{ds}{dt}|\mathcal{F}_t\} \tag{2.1}$$

where $f = m+v$ is the so-called fundamental and $\mathsf{E}\{\frac{ds}{dt}|\mathcal{F}_t\}$ is the expected change of the exchange rate at time t given the information set $\mathcal{F}_t$. The fundamental consists of two components: m represents the policy instruments that the central bank controls, and v contains all the other factors that affect the exchange rate. This component is assumed to follow a Brownian motion. If there is no currency band and the currency floats freely, the central bank does not intervene ($m = 0$) so that f and thus s follow a Brownian motion. This being the case, the expected change of the exchange rate equals zero. Krugman (1991) assumes that there exists a target zone, $s^L \leq s \leq s^U$, and that the authorities intervene through m when the exchange rate reaches either boundary value s^L or s^U. This changes the agents' expectations when the zone is credible. When the exchange rate lies near either boundary, the probability of the exchange rate to move towards the centre is perceived to be higher than the probability that it moves even closer to the boundary. The agents anticipate the intervention, $\mathsf{E}\{\frac{ds}{dt}|\mathcal{F}_t\} \neq 0$, so that the zone creates a nonlinearity called the "honeymoon effect" in the behaviour of the rate as a function of the fundamental. Instead of a straight line, the relationship between the exchange rate and fundamental is characterized as a smooth S-curve. This is a much investigated detail of the model in the empirical literature. A consequence of the S-curve is that the exchange rate will spend more time close to the boundaries that in other segments of the zone, so that the marginal distribution of the exchange rate will be $\cup$-shaped. Finally, the conditional variance should have an $\cap$-shape. We shall investigate the last two implications of the Krugman model in the empirical section of the paper.

As briefly mentioned in the Introduction, a number of authors have fitted the continuous-time model (2.1) to exchange rate data. The formal solution of the symmetric model has the form; see, for example, Krugman (1991), Lindberg and

Söderlind (1994b) or Taylor (1995),

$$s = m + v + 2A \sinh\{\delta(m + v)\}$$

where $\delta = (2/\alpha)^{1/2}\sigma$, with σ being the standard deviation of the innovation in the fundamental and A is a function of the smoothness conditions determining how s approaches the boundaries ("smooth pasting"). The relevant parameters of the model are estimated using the method of simulated moments as described in Lindberg and Söderlind (1994b), Iannizzotto and Taylor (1999) and other articles. Recently, Chung and Tauchen (2001) estimated target zone models with an implicit band using the efficient method of moments. A typical finding is that there is little evidence of the S-shape in the relationship between the exchange rate and the fundamental.

In this paper we construct a discrete time model that is applicable to daily exchange rate series. Monthly series used by many authors who are typically interested in modelling the conditional mean of the process, are too short for our purposes. A central feature of our model is joint modelling of the conditional mean and the conditional variance of the exchange rate in a target zone. Perhaps the closest equivalent to our approach is the one adopted in Bekaert and Gray (1998), henceforth BG, see also Forbes and Kofman (2000), and Klaster and Knot (2002) for a recent application. A feature that our approach shares with theirs is that we also explicitly model the conditional variance of the exchange rate. Furthermore, as in BG, the fundamental is not explicitly observed. BG consider the whole conditional distribution of the first difference of the exchange rate. As they assume normality, this implies modelling the first two moments of the distribution. A simplified form of the model of Δs_t has the following form

$$f(\Delta s_t | \mathcal{F}_t) = \frac{\Phi\left(\frac{\Delta s_t - m_t}{h_t^{1/2}}\right)\frac{1}{h_t^{1/2}}}{\Phi\left(\frac{\Delta_{U_{t-1}} - m_t}{h_t^{1/2}}\right) - \Phi\left(\frac{\Delta_{L_{t-1}} - m_t}{h_t^{1/2}}\right)} \tag{2.2}$$

where Φ is the cumulative distribution function of a standard normal variable, m_t is the conditional mean and h_t the conditional variance. Furthermore, $\Delta_{U_{t-1}} = s^U - x_{t-1}$, the largest possible change of the exchange rate (the zone is perfectly credible), and $\Delta_{L_{t-1}} = s^L - x_{t-1}$ smallest possible change. Thus, (2.2) is a truncated normal density where the boundaries define the truncation points. The conditional mean m_t is a linear function of PB_{t-1}, the position of the exchange rate in the band

at $t-1$, and the conditional variance h_t is described by a GARCH(1,1) process augmented by $|PB_{t-1}|$. Density (2.2) at time t is in fact a density forecast of the change in s_t from $t-1$ to t. Forbes and Kofman (2000) make use of the same general set-up but relax the assumption of a perfectly credible zone by allowing a positive probability for the exchange rate to venture outside the boundaries.

Function (2.2) only forms a part of the BG model because that model also includes a parameterization of jump behaviour of the exchange rate. In this respect it differs from the STARTZ model that does not have such a feature. In this paper we concentrate on modelling the dynamic behaviour of the exchange rate within the band, and in our empirical examples the exchange rate series do not contain jumps. Even if they did, we would have to make the unattractive assumption that all jumps are generated by the same mechanism and that the policy of the central bank remains unchanged after any jump or realignment. The STARTZ model will be considered in detail in the next section.

3. The Model

The basic idea that the STARTZ model share with the model of BG is that the dynamics of both the conditional mean and the conditional variance change when the process approaches the boundary of the target zone. We assume that degree of change depends nonlinearly on the distance between the value of the process and the central parity of the target zone. For example, one might expect the conditional mean to behave as a random walk process close to the central parity, whereas close to the boundary the process will have a tendency to move towards the central parity due to interventions by the central bank.

In general terms, the conditional mean of the model is defined as

$$y_t = m_t(\boldsymbol{\varphi}, \gamma_a, \theta_a, \mu; \mathbf{y}_{t-1}) + \varepsilon_t \tag{3.1}$$

where y_t is the deviation of the exchange rate from the centre of the target zone and $\mathbf{y}_{t-1} = (y_{t-1}, ..., y_{t-p})'$. The function $m_t = m_t(\boldsymbol{\varphi}, \gamma_a, \theta_a, \mu; \mathbf{y}_{t-1})$ is assumed to be bounded and at least twice continuously differentiable for its parameters almost everywhere in the parameter space for any $\mathbf{y}_{t-1}$ belonging to the corresponding sample space. The error process of the model is parameterized as

$$\varepsilon_t = z_t \sqrt{h_t(\boldsymbol{\varphi}, \gamma_a, \theta_a, \mu, \boldsymbol{\eta}, \gamma_b, \theta_b, \delta; E_{t-1})} \tag{3.2}$$

where $\{z_t\} \sim \text{iid}(0,1)$ and $h_t = h_t(\boldsymbol{\varphi}, \gamma_a, \theta_a, \mu, \boldsymbol{\eta}, \gamma_b, \theta_b, \delta)$ is a positive-valued function with $E_{t-1} = (\varepsilon_{t-j} : j \geq 1)$. The structure of (3.2) implies that there is no autocorrelation in the error process $\{\varepsilon_t\}$. Furthermore, $\varepsilon_t = y_t - m_t$ such that $\boldsymbol{\varphi}$ is assumed not to depend on $\boldsymbol{\eta}$. The conditional variance h_t is at least twice continuously differentiable for the parameters almost everywhere in the parameter space. It is also assumed that the moments of y_t necessary for the inference exist and that the parameters are subject to restrictions such that the process defined by (3.1) and (3.2) is stationary and ergodic. This assumption will be satisfied because $\{y_t\}$ is bounded both from below and above due to the target zone.

In order to define m_t and h_t (and to consider the misspecification tests in Lundbergh and Teräsvirta (2002)), let

$$
\begin{aligned}
G^L(s_t; \gamma, \theta, c) &= (1 + \exp(-\gamma(c - s_t)))^{-\theta}, \gamma > 0, \theta > 0 \qquad (3.3) \\
G^U(s_t; \gamma, \theta, c) &= (1 + \exp(-\gamma(s_t - c)))^{-\theta}, \gamma > 0, \theta > 0
\end{aligned}
$$

where s_t is the transition variable, γ a slope parameter, θ an asymmetry parameter and c is a location parameter. The parameter restrictions $\gamma > 0$, and $\theta > 0$ are identifying restrictions. Function (3.3) is a generalized logistic function; see Nelder (1961) and Sollis, Leybourne and Newbold (1999). It contains as a special case ($\theta = 1$) the standard logistic function. The asymmetry parameter θ is essential in this application where the movements of the exchange rate are restricted by the boundaries of the target zone. Note the slight reparameterization (γ instead of γ/θ) compared to Sollis et al. (1999). In growth curve literature, the generalized logistic function is called the Richards growth curve, see Richards (1959).

According to theoretical target zone models, the conditional mean should be a nonlinear function (S-shaped) of the underlying fundamentals with local nonlinearity emerging close to the band ("the honeymoon effect"). This requirement can be met by the following model specification

$$
\begin{aligned}
m_t =\ & \boldsymbol{\varphi}'\mathbf{x}_t \\
& + (\mu s^L - \boldsymbol{\varphi}'\mathbf{x}_t)G^L(y_{t-1}; \gamma_a, \theta_a, \mu s^L) \qquad (3.4) \\
& + (\mu s^U - \boldsymbol{\varphi}'\mathbf{x}_t)G^U(y_{t-1}; \gamma_a, \theta_a, \mu s^U)
\end{aligned}
$$

where $\mathbf{x}_t = (1, y_{t-1}, ..., y_{t-n})'$ is an $(n+1) \times 1$ intercept-lag vector and $\boldsymbol{\varphi} = (\varphi_0, \varphi_1, ..., \varphi_n)'$ the corresponding parameter vector. Vector $\mathbf{x}_t$ implicitly contains all information about the fundamentals at $t = 1$. The linear autoregressive compo-

nent in (3.4) is complemented by two terms that help to characterize the behaviour of the conditional mean close to the lower (s^L) and the upper (s^U) boundary of the target zone. It is assumed in (3.4) that m_t is symmetric in the sense that the local behaviour of the exchange rate is similar in the neighborhood of both boundaries. This assumption can be relaxed, however. Parameter $\mu, 0 < \mu < 1$, adds flexibility to the specification and allows the investigator, among other things, to estimate an implicit band inside the official one, should such a band exist. The slope parameter, $\gamma_a > 0$ and the asymmetry parameter $\theta_a > 0$ jointly tell us how pronounced is the change in the local dynamic behaviour of the exchange rate when one moves from the centre of the target zone to the neighborhood of either boundary.

The conditional mean model (3.4) has the following interpretation. Near the centre of the band the behaviour of the exchange rate is characterized, at least approximately, by a linear combination of its lags, $\varphi' \mathbf{x}_t$ as both $G^L \approx 0$ and $G^U \approx 0$. Close to both the upper and the lower boundary of the target zone the exchange rate depends nonlinearly on $\mathbf{x}_t$. For example, in the case when the exchange rate approaches the upper boundary, $G^U \to 1$, and there is a smooth transition from the autoregressive behaviour represented by $\varphi' \mathbf{x}_t$ towards white-noise like behaviour around μs^U. Obviously, $1 - \mu > 0$ is small. The speed of the transition is determined by γ, θ and c. When the test approaches the lower boundary, $G^L \to 1$ and similar conclusion follows.

According to theoretical target zone models the conditional variance should have a $\cap$-shaped distribution, as the conditional variance of the process must be small close to the boundaries if the band is credible. We parameterize this requirement in a way similar to what was used for the conditional mean. Thus,

$$
\begin{aligned}
h_t \;=\; & \boldsymbol{\eta}' \mathbf{w}_t \\
& + (\delta - \boldsymbol{\eta}' \mathbf{w}_t) G^L(y_{t-1}; \gamma_b, \theta_b, \mu s^L) \\
& + (\delta - \boldsymbol{\eta}' \mathbf{w}_t) G^U(y_{t-1}; \gamma_b, \theta_b, \mu s^U)
\end{aligned}
\tag{3.5}
$$

where constants s^L and s^U again represent the lower and the upper boundary. Parameters γ_b and θ_b are different from γ_a and θ_a respectively, whereas for simplicity, μ is assumed to be the same as in (3.4). This has the technical consequence that when the parameters of (3.4) and (3.5) are estimated by (quasi) maximum likelihood, the information matrix not block diagonal. The conditional mean and variance thus have to be estimated simultaneously. The generalized logistic func-

tions are defined in the same way as they are for the conditional mean. Setting $\boldsymbol{\eta} = (\alpha_0, \alpha_1, ..., \alpha_q, \beta_1, ..., \beta_p)'$ and $\mathbf{w}_t = (1, \varepsilon_{t-1}^2, ..., \varepsilon_{t-q}^2, h_{t-1}, ..., h_{t-p})'$, where $h_t > 0$ almost everywhere, makes $\boldsymbol{\eta}'\mathbf{w}_t$ in (3.5) a standard GARCH(p,q) type specification. Assuming $\delta > 0$ together with the restrictions $\alpha_0 > 0$, $\alpha_j \geq 0$, $j = 1, ...q$; $\beta_j \geq 0$, $j = 1, .., p$; is sufficient for positivity of the conditional variance.

Specification (3.5) implies that the conditional variance is a nonlinear function of the elements of $\mathbf{w}_t$. For example, in the case when the exchange rate approaches the upper boundary there is a smooth transition from a standard GARCH type behaviour represented by $\boldsymbol{\eta}'\mathbf{w}_t$ towards a constant $\delta > 0$ that is expected to be close to zero.

Equations (3.1), (3.2), (3.4) and (3.5) jointly define a Smooth Transition Autoregressive Target Zone (STARTZ) model. Near the boundaries the STARTZ process behaves like an iid process with mean μs^L or μs^U and a (small) variance δ. The process is thus bounded in probability, stationary and ergodic. Furthermore due to the boundness, all moments of $\{y_t\}$ exists.

Although the STARTZ model is aimed at modelling exchange rates restricted by a target zone, it does follow that the exchange rate remain inside the target zone with probability 1. At the boundary the conditional variance of the STARTZ model is small but still positive, as $\delta > 0$. Thus, a shock such that the exchange rate breaks through the boundary of the zone does have a positive probability. This is not unrealistic: even when there is no realignment the exchange rate can momentarily leave the band by a small margin and be quickly brought back again. There is evidence of such events in our data sets.

4. Specification and estimation

The nonlinear STARTZ model defined by (3.1-3.5) is our most general parameterization of the target zone model. In order to carry out the empirical work in an orderly fashion, we propose a modelling strategy that can be described as follows.

1. Select an AR(n) model for the conditional mean according to some suitable criterion such as the AIC (Akaike, 1974) or BIC (Rissanen, 1978, Schwarz, 1978). Select a low-order ARCH or GARCH model for the conditional variance.

2. Estimate an AR(n)-GARCH(p, q) model, reduce the size of the model if necessary and re-estimate for the series to obtain initial values for estimating the STARTZ model.

3. Estimate the parameters of the STARTZ model and test the adequacy of both the conditional mean and the conditional variance specification by appropriate misspecification tests.

4. If the model passes the tests, tentatively accept it. In the opposite case try another specification search or choose another family of models.

All parameter estimates are obtained by maximizing the log-likelihood under the assumption that $\{z_t\}$ is a sequence of independent standard normal errors. In that case the (quasi) log-likelihood function at time t equals

$$l_t = \text{const} - \frac{1}{2} \ln h_t - \frac{1}{2} \frac{\varepsilon_t^2}{h_t} \tag{4.1}$$

where ε_t and h_t are defined in (3.5). We assume that the model under consideration satisfies the necessary regularity conditions needed for the consistency and asymptotic normality of the estimators. In the following section we consider the evaluation of this model. The derivates of the log-likelihood function (4.1) are to be found in Appedix A.

5. Evaluation by misspecification tests

Once we have specified and estimated a model it is important to investigate the validity of the assumptions used in the estimation. We can modify the misspecification tests in Teräsvirta (1994) and Lundbergh and Teräsvirta (2002) for the present situation. The tests in this section only consider misspecification of the $AR(n)$-$GARCH(p,q)$ parameterization within the target zone. The target zone itself is assumed known. In order to describe the tests we first introduce a general structure and thereafter briefly consider each test separately.

5.1. General

Consider the STARTZ model as defined in (3.1) and (3.2). A "quasi-additive" extension of the model may be written as

$$
\begin{aligned}
y_t &= \boldsymbol{\varphi}' \mathbf{x}_t + A(\mathbf{x}_t; \boldsymbol{\pi}_a) \\
&\quad + (\mu s^L - \boldsymbol{\varphi}' \mathbf{x}_t - A(\mathbf{x}_t; \boldsymbol{\pi}_a)) G^L(y_{t-1}; \gamma_a, \theta_a, \mu s^L) \\
&\quad + (\mu s^U - \boldsymbol{\varphi}' \mathbf{x}_t - A(\mathbf{x}_t; \boldsymbol{\pi}_a)) G^U(y_{t-1}; \gamma_a, \theta_a, \mu s^U) + \varepsilon_t \\
\varepsilon_t &= z_t \sqrt{h_t}
\end{aligned}
\tag{5.1}
$$

$$
\begin{aligned}
h_t \;=\;\; & \boldsymbol{\eta}'\mathbf{w}_t + B(\mathbf{w}_t;\boldsymbol{\pi}_b) \\
& +(\delta - \boldsymbol{\eta}'\mathbf{w}_t - B(\mathbf{w}_t;\boldsymbol{\pi}_b))G^L(y_{t-1};\gamma_b,\theta_b,\mu s^L) \\
& +(\delta - \boldsymbol{\eta}'\mathbf{w}_t - B(\mathbf{w}_t;\boldsymbol{\pi}_b))G^U(y_{t-1};\gamma_b,\theta_b,\mu s^U)
\end{aligned}
$$

where functions $A(\mathbf{x}_t;\boldsymbol{\pi}_a)$ and $B(\mathbf{w}_t;\boldsymbol{\pi}_b)$ are assumed twice continuously differentiable for all $\boldsymbol{\pi}_a$ and $\boldsymbol{\pi}_b$ everywhere in the corresponding sample spaces. For notational simplicity and without loss of generality we assume $A(\mathbf{x}_t;\mathbf{0}) = B(\mathbf{w}_t;\mathbf{0}) \equiv 0$. Furthermore, $\boldsymbol{\eta}'\mathbf{w}_t + B(\mathbf{w}_t;\boldsymbol{\pi}_b)$ is assumed to be positive-valued almost everywhere and $\{z_t\}$ is a sequence of independent standard normal variables. Model (5.1) forms a unifying framework for our tests.

The null hypothesis of no additional structure in (3.1) and (3.2) now has the form $\mathrm{H}_0 : \boldsymbol{\pi}_a = 0$ and $\boldsymbol{\pi}_b = 0$ in (5.1). Let $\boldsymbol{\omega} = (\boldsymbol{\varphi}',\gamma_a,\theta_a,\mu,\boldsymbol{\eta}',\gamma_b,\theta_b,\delta)'$, which comprises all the parameters of the model under this null hypothesis. It is assumed that the maximum likelihood estimator of $\boldsymbol{\omega}$ is consistent and asymptotically normal under any null hypothesis to be considered, which requires $\{y_t\}$ to be stationary and ergodic and that the log-likelihood function satisfies the standard regularity conditions. Since $\{y_t\}$ is restricted by the boundaries of the target zone the necessary moments of $\{\varepsilon_t\}$ implied by the Hessian matrix and required for the asymptotic distribution theory to work do exist. The Lagrange multiplier (or score) statistic is defined as

$$
\mathrm{LM} = T\left(\begin{array}{c} \frac{1}{T}\sum \frac{\partial l_t}{\partial \boldsymbol{\pi}_a}|_{\mathrm{H}_0} \\ \mathbf{0} \\ \frac{1}{T}\sum \frac{\partial l_t}{\partial \boldsymbol{\pi}_b}|_{\mathrm{H}_0} \end{array}\right)' \widehat{\mathbf{I}}(\boldsymbol{\pi}_a,\widehat{\boldsymbol{\omega}},\boldsymbol{\pi}_b)^{-1}|_{\mathrm{H}_0}\left(\begin{array}{c} \frac{1}{T}\sum \frac{\partial l_t}{\partial \boldsymbol{\pi}_a}|_{\mathrm{H}_0} \\ \mathbf{0} \\ \frac{1}{T}\sum \frac{\partial l_t}{\partial \boldsymbol{\pi}_b}|_{\mathrm{H}_0} \end{array}\right) \tag{5.2}
$$

where $\widehat{\mathbf{I}}$ is a consistent estimator of the information matrix under the null hypothesis. The partial derivatives defining the score can be found in Appendix A. Under the null hypothesis, statistic (5.2) is asymptotically χ^2- distributed with $\dim(\boldsymbol{\pi}_a) + \dim(\boldsymbol{\pi}_b)$ degrees of freedom.

As (4.1) indicates the likelihood is constructed under the assumption of normality. It is not certain, however that this assumption is satisfied in the present situation. For this reason we construct a robust version of the LM test following Wooldridge (1990). The test is carried out in a TR^2 form as follows:

1. Regress $\frac{1}{\sqrt{\widehat{h}_t}}\frac{\partial \varepsilon_t}{\partial \boldsymbol{\pi}_a}$ on $\frac{1}{\sqrt{\widehat{h}_t}}\frac{\partial \varepsilon_t}{\partial \boldsymbol{\omega}}$ and $\frac{1}{\sqrt{\widehat{h}_t}}\frac{\partial h_t}{\partial \boldsymbol{\pi}_b}$ on $\frac{1}{\sqrt{\widehat{h}_t}}\frac{\partial h_t}{\partial \boldsymbol{\omega}}$. This results in two residual vectors $\ddot{\lambda}_{\boldsymbol{\pi}_a,t}$ and $\ddot{\lambda}_{\boldsymbol{\pi}_b,t}$ of dimensions $\boldsymbol{\pi}_a$ and $\boldsymbol{\pi}_b$ respectively.

2. Regress 1 on $(\widehat{\varepsilon}_t/\sqrt{\widehat{h}_t})\ddot{\lambda}_{\boldsymbol{\pi}_a,t}$ and $((\widehat{\varepsilon}_t^2 - \widehat{h}_t)/\sqrt{2\widehat{h}_t})\ddot{\lambda}_{\boldsymbol{\pi}_a,t}$ and compute the sum

of squared residuals (SSR).

3. Compute the test statistic, $T - SSR$, which is asymptotically χ^2- distributed with $\dim(\boldsymbol{\pi}_a) + \dim(\boldsymbol{\pi}_b)$ degrees of freedom under the null hypothesis.

This is the form in which we compute our LM and LM-type misspecification tests.

5.2. Test against misspecified lag lengths

To test the null hypothesis of misspecified lag lengths in either the conditional mean or the conditional variance or both, the alternative is stated as remaining lag lengths of order n_a in the ordinary error process and of order p_b in the squared (and standardized) errors. In the general case, this gives the extended model (5.1) with $A(\mathbf{x}_t; \boldsymbol{\pi}_a) = \boldsymbol{\pi}_a' \mathbf{v}_t^a$ and $B(\mathbf{w}_t; \boldsymbol{\pi}_b) = \boldsymbol{\pi}_b' \mathbf{v}_t^b$ where $\boldsymbol{\pi}_a = (\pi_{a,1}, ..., \pi_{a,n_a})'$, $\mathbf{v}_t^a = (y_{t-n-1}, ..., y_{t-n-n_a})'$, $\boldsymbol{\pi}_b = (\pi_{b,1}, ..., \pi_{b,p_b})'$, and $\mathbf{v}_t^b = (h_{t-(p+1)}, ..., h_{t-(p+p_b)})'$. That is, the alternative in the conditional mean is an AR process of order $n + n_a$ and the alternative to test against in the conditional variance the alternative is a higher order GARCH, see Bollerslev (1986). The null hypothesis of no remaining higher dependence in either the conditional mean or in the conditional variance is equivalent to $\boldsymbol{\pi}_a = \mathbf{0}$ and $\boldsymbol{\pi}_b = \mathbf{0}$. Under this null hypothesis, the LM-statistic (5.2) is asymptotically χ^2- distributed with $\dim(\boldsymbol{\pi}_a) + \dim(\boldsymbol{\pi}_b)$ degrees of freedom. Note that $\mathbf{v}_t^b$ may be replaced by $\mathbf{v}_t^{b*} = (\varepsilon_{t-(q+1)}^2, ..., \varepsilon_{t-(q+r)}^2)'$ as in Bollerslev (1986). The test can also be carried out separately for the conditional mean and the conditional variance.

5.3. Test against remaining nonlinearity

If an estimated STARTZ model adequately characterizes all nonlinearity in the exchange rate series, there should be no unmodelled nonlinearity left after fitting the model to the data. This can be checked by testing the hypothesis of no remaining additive nonlinearity. The alternative to this null hypothesis is assumed to be an additive smooth transition component. This alternative is obtained as a special case of (5.1) and with

$$A(\mathbf{x}_t; \boldsymbol{\varphi}_a, \rho_a, c_a) = \boldsymbol{\varphi}_a' \mathbf{x}_t \overline{H}_a^*(y_{t-1}; \rho_a, c_a) \qquad (5.3)$$

$$B(\mathbf{w}_t; \boldsymbol{\eta}_b, \rho_b, c_b) = \boldsymbol{\eta}_b' \varepsilon_{t-1} \overline{H}_b^*(\varepsilon_{t-1}; \rho_b, c_b) \qquad (5.4)$$

where $\overline{H}_i^* = H_i^* - \frac{1}{2}$ where $i = a, b$. Subtracting $1/2$ from H_i^* is just a notational convenience in deriving the test and does not affect the generality of the argument. The definitions of functions $A(\mathbf{x}_t; \boldsymbol{\varphi}_a, \rho_a, c_a)$ and $B(\mathbf{w}_t; \boldsymbol{\eta}_b, \rho_b, c_b)$ are the same as in Lundbergh and Teräsvirta (2002). Function H_i^* is defined as function G^U in (3.3) with $\rho_i = \gamma$, $\mathbf{c}_i$, $i = a, b$, and $\theta \equiv 1$. The joint null hypothesis of no additional nonlinearity in the conditional mean and the conditional variance can be expressed as $\mathrm{H}_0 : \rho_a = \rho_b = 0$.

Parameters in $A(\mathbf{x}_t; \boldsymbol{\varphi}_a, \rho_a, c_a)$ and $B(\mathbf{w}_t; \boldsymbol{\eta}_b, \rho_b, c_b)$ are not identified under the null hypothesis. For example when $\rho_a = 0$, parameters $\boldsymbol{\varphi}_a$ and c_a in (5.3) are unidentified nuisance parameters. Following Luukkonen, Saikkonen and Teräsvirta (1988), see also Eitrheim and Teräsvirta (1996) and Lundbergh and Teräsvirta (2002), we circumvent this identification problem under the null hypothesis by expanding functions $A(\mathbf{x}_t; \boldsymbol{\varphi}_a, \rho_a, c_a)$ and $B(\mathbf{w}_t; \boldsymbol{\eta}_b, \rho_b, c_b)$ into a Taylor series around the null hypothesis. Using the first-order expansion this yields, after reparameterization, a transformed model with $A(\mathbf{x}_t; \boldsymbol{\varphi}_a, \rho_a, c_a) = \widetilde{\boldsymbol{\pi}}_a' \mathbf{v}_t^a + R_1(\mathbf{x}_t; \boldsymbol{\varphi}_a, \rho_a, c_a, \boldsymbol{\omega})$ and $B(\mathbf{w}_t; \boldsymbol{\eta}_b, \rho_b, c_b) = \widetilde{\boldsymbol{\pi}}_b' \mathbf{v}_t^b + R_2(\mathbf{w}_t; \boldsymbol{\eta}_b, \rho_b, c_b, \boldsymbol{\omega})$ where $\mathbf{v}_t^a = (\mathbf{x}_t^{nc} y_{t-1})'$, $\widetilde{\boldsymbol{\pi}}_b = (\tau_{b,1}, \tau_{b,2})'$ and $\mathbf{v}_t^b = (\varepsilon_{t-1}, \varepsilon_{t-1}^3)'$ where $\mathbf{x}_t = (1, (\mathbf{x}_t^{nc})')'$. The two remainders of the Taylor expansions, $R_1(\mathbf{x}_t; \boldsymbol{\varphi}_a, \rho_a, c_a, \boldsymbol{\omega})$ and $R_2(\mathbf{w}_t; \boldsymbol{\eta}_b, \rho_b, c_b, \boldsymbol{\omega})$, do not affect the asymptotic distribution theory because both are identically equal to zero under H_0. The new null hypothesis is $\widetilde{\boldsymbol{\pi}}_a = \mathbf{0}$ and $\widetilde{\boldsymbol{\pi}}_b = \mathbf{0}$, under which the LM-type test statistic (5.2) is asymptotically χ^2- distributed with $\dim(\widetilde{\boldsymbol{\pi}}_a) + \dim(\widetilde{\boldsymbol{\pi}}_b)$ degrees of freedom.

In practice it is most often useful to divide this joint test into separate tests for the conditional mean and the conditional variance. This helps to locate the problem, if any, and thus makes it easier to find a remedy to it.

5.4. Test against nonconstant parameters

We assume that the alternative to constant parameters in either the conditional mean or the conditional variance or both is that the parameters change smoothly over time, see Lin and Teräsvirta (1994) and Lundbergh and Teräsvirta (2002). The changing parameters replacing $\boldsymbol{\varphi}$ in (3.4) and $\boldsymbol{\eta}$ in (3.5) are defined as follows:

$$\boldsymbol{\varphi}_t = \boldsymbol{\varphi}^* + \boldsymbol{\lambda}_{\boldsymbol{\varphi}} \overline{H}_{\boldsymbol{\varphi}}^*(t; \rho_{\boldsymbol{\varphi}}, c_{\boldsymbol{\varphi}}) \qquad (5.5)$$

$$\boldsymbol{\eta}_t = \boldsymbol{\eta}^* + \boldsymbol{\lambda}_{\boldsymbol{\eta}} \overline{H}_{\boldsymbol{\eta}}^*(t; \rho_{\boldsymbol{\eta}}, c_{\boldsymbol{\eta}}) \qquad (5.6)$$

Again we define $\overline{H}_i^* = H_i^* - \frac{1}{2}$ where $i = \varphi, \eta$. H_i^* is the logistic function defined as function G^U in (3.3) with $s_t \equiv t$, $\gamma = \rho_i$, $\theta \equiv 1$ and $c = c_i$. The null hypothesis of parameter constancy can be stated as $H_0 : \rho_\varphi = \rho_\eta = 0$ under which $\varphi^* = \varphi$ and $\eta^* = \eta$. Setting $A(\mathbf{x}_t; \boldsymbol{\lambda}_\varphi, \rho_\varphi, c_\varphi) = \boldsymbol{\lambda}_\varphi \overline{H}_\varphi(t; \rho_\varphi, c_\varphi)$ and $B(\mathbf{w}_t; \boldsymbol{\lambda}_\eta, \rho_\eta, c_\eta) = \boldsymbol{\lambda}_\eta \overline{H}_\eta(t; \rho_\eta, c_\eta)$ it is seen that this alternative is a special case of (5.1). The identification problem under the null hypothesis is circumvented as before by expanding $\overline{H}_i(t; \rho_i, c_i)$ into a Taylor series around the null hypothesis, $\rho_i = 0$. Using the first-order expansion we obtain the extended model (5.1) with $A(\mathbf{x}_t; \rho_\varphi, c_\varphi) = \widetilde{\boldsymbol{\pi}}_a' \mathbf{v}_t^a + R_1(\mathbf{x}_t; \rho_\varphi, c_\varphi)$ and $B(\mathbf{w}_t; \rho_\eta, c_\eta) = \widetilde{\boldsymbol{\pi}}_b' \mathbf{v}_t^b + R_2(\mathbf{w}_t; \rho_\eta, c_\eta)$ where $\mathbf{v}_t^a = \mathbf{x}_t t$ and $\mathbf{v}_t^b = \mathbf{w}_t t$. The joint null hypothesis of parameter constancy in both the conditional mean and variance consists of the restrictions $\widetilde{\boldsymbol{\pi}}_a = \mathbf{0}$ and $\widetilde{\boldsymbol{\pi}}_b = \mathbf{0}$. As in the preceding section the two remainder terms of the Taylor expansions $R_1(\mathbf{x}_t; \rho_\varphi, c_\varphi, \boldsymbol{\omega}) \equiv R_2(\mathbf{w}_t; \rho_\eta, c_\eta, \boldsymbol{\omega}) \equiv 0$ under the null hypothesis so that they do not affect the asymptotic distribution theory. The LM-type test statistic (5.2) is thus asymptotically χ^2- distributed with $\dim(\widetilde{\boldsymbol{\pi}}_a) + \dim(\widetilde{\boldsymbol{\pi}}_b)$ degrees of freedom, when the null hypothesis holds. Even here, testing the conditional mean and variance separately is advisable for the same reason as before.

It is also possible to test constancy of μ and δ. The alternative model, in this case, is not a special case of (5.1) but the parameterization is similar to (5.5) and (5.6). Under the alternative,

$$
\begin{aligned}
\mu_t &= \mu^* + \boldsymbol{\lambda}_\mu \overline{H}_\mu^*(t; \rho_\mu, c_\mu) \\
\delta_t &= \delta^* + \boldsymbol{\lambda}_\delta \overline{H}_\delta^*(t; \rho_\delta, c_\delta)
\end{aligned}
$$

where $\overline{H}_\mu^*$ and $\overline{H}_\delta^*$ are defined as before. The null hypothesis of parameter constancy can be stated as $H_0 : \rho_\mu = 0$ and $\rho_\delta = 0$. Once again the identification problem is circumvented by expanding the $\overline{H}_\mu^*$ and $\overline{H}_\delta^*$ into Taylor series around the null hypothesis. This yields $\mu_t = \mu^* + \boldsymbol{\pi}_\mu t + R_3(\mathbf{x}_t; \rho_\mu, c_\mu)$ and $\delta_t = \delta^* + \boldsymbol{\pi}_\delta t + R_4(\mathbf{x}_t; \rho_\delta, c_\delta)$. Even here, the two remainder terms of the Taylor expansions $R_3(\mathbf{x}_t; \rho_\mu, c_\mu) \equiv R_4(\mathbf{x}_t; \rho_\delta, c_\delta) \equiv 0$ under the null hypothesis. The Lagrange multiplier type statistic is obtained as in the previous cases. The relevant elements of the score for observations t can be found in Appendix A. A robust version of the test statistic is used in the applications.

6. Modelling two Nordic currencies

In this section the STARTZ model is applied to the Norwegian and Swedish currencies. In the second half of the 1980s these currencies had unilateral target zones against a trade weighted currency basket. We focus on periods with no realignments and no policy changes. The data for all currencies are daily observations and we model the deviation, in percent, of the exchange rate index from the central parity.

6.1. The Swedish krona: 1985-1991

The daily series of the Swedish exchange rate index contains 1472 observations and covers the period from July 1, 1985 to May 17, 1991. A graph of the index can be found in Figure B.1. The starting point of the series coincides with the introduction of an explicit target zone. In May 1991 the trade weighted currency basket was replaced by the ECU-index. The index under consideration was based on a basket of currencies, so that what is modelled is not a bilateral exchange rate. During the observation period the index was allowed to vary within $\pm 1.5\%$ from its central parity.

Riksbanken (Bank of Sweden) intervened intramarginally during this period; see Lindberg and Söderlind (1994a) for a detailed description of the intervention policy of the Bank. The assumptions of the original Krugman (1991) model are thus not satisfied. When Lindberg and Söderlind (1994b) estimated the parameters of their continuous time model for the period 27 June 1985 to 15 November 1990 they found practically no evidence of the S-shape in the relationship between the exchange rate and the fundamentals.

Figure B.1 shows that the currency index has not approached the upper bound of the zone at any time during the observation period. As the STARTZ model is symmetric, the behaviour of the index at this end of the zone is estimated solely on the basis of the observations close to its lower boundary. The estimates of the parameters of the STARTZ model for the Swedish krona can be found in Table B.1. Note that the value of γ_a, the slope parameter of the transition functions G^L and G^U found in (3.3), has been set to 300. Estimating γ_a and θ_a jointly turned out to be very difficult as the log-likelihood around the maximum was very flat. This is an indication of the fact that a number of parameter combinations yields the transition functions that have the same shape. Conditioning on one of the parameters was therefore necessary. The standard deviation of $\widehat{\theta}_a$ remains large even thereafter, which is another indication of the flatness of the log-likelihood around its maximum

value.

It is seen from the estimates and Figure B.11(a) that the transition from the dynamic behaviour of the index in the centre of the zone to the boundary behaviour is abrupt and occurs only very close to the boundaries. The sum of the estimates of the autoregressive parameters $\widehat{\varphi}_1 + \widehat{\varphi}_2 = 0.996$, so that most of the time (note that $\widehat{\mu} = 0.99$) the behaviour of the index is dominated by a near unit root. This accords with the results of Lindberg and Söderlind (1994b) and the intramarginal interventions. It also suggests that the mean reversion in the exchange rate is in this case extremely weak. Note that previous results suggesting mean reversion such as the ones in Svensson (1993) are obtained by fitting linear autoregressive models without any restrictions to the central parity adjusted series, so that they are not comparable to the ones reported here.

The conditional variance mostly displays rather mild GARCH effects with low persistence: $\widehat{\alpha}_1 + \widehat{\beta}_1 = 0.88$. The estimate of δ is very small ($= 0.0010$) as expected. Even here, the slope parameter γ_b of the transition functions was restricted to 300 to allow the estimation algorithm to converge. The functions are graphed in Figure B.11(a) and show the same abrupt change of behaviour near the boundaries as was discovered in the conditional mean. Figure B.3 shows the 1472 deviations from the central parity graphed against the conditional variances estimated from (3.5). Not unexpectedly, the ∩-shape of the conditional variance is very weak.

Some properties of the standardized residuals can be found in Table B.2. It is seen that the residuals contain some outliers as the kurtosis exceeds five. The results of the misspecification tests of the conditional mean appear in Table B.3, the ones of the conditional variance in Table B.4 and the one of both the conditional mean and variance in Table B.5. The model passes the tests of misspecified lag structure, no additional nonlinearity and constancy of φ_1, φ_2 (conditional mean) and $\delta, \alpha_0, \alpha_1$ and β_1 (conditional variance). Also the model passes the misspecification test of constancy of μ (both in the conditional mean and variance).

In order to illustrate the behaviour of the model in terms of the marginal distribution of the exchange rate, 100000 observations are generated from the estimated STARTZ model and the density of the observations is smoothed using a standard kernel smoother. Figure B.7 shows that the density is hump-shaped. The small earlobes at both tails suggest that there have been interventions close to the lower boundary to offset the pressure on the index to cross the boundary. The shape of the extreme tail is due to the kernel [Epanechnikov] selected for this application and should therefore be interpreted with caution.

6.2. The Norwegian krone: 1986-1988

The Norwegian exchange rate index analyzed in this paper covers the period from October 1, 1986 to October 22, 1990. A graph of the index can be found in Figure B.2. The starting point of the series coincides with a realignment of the zone and, as in the case of Sweden, the observation period ends when the trade weighted currency basket was replaced by the ECU-index. The index as allowed to vary within $\pm 2.25\%$ from its central parity.

Information about the intervention policy of Norges Bank (Central Bank of Norway) can be found in Lysebo and Mundaca (1997) and Mundaca (2000). From October 1986 onwards the Bank first intervened mainly when the index was close to either boundary. In mid-June 1988, a change in the intervention policy was announced, and for the rest of the period the interventions were intramarginal. Late in 1988 Norges Bank started to maintain an implicit target zone that was narrower than the official one. This can be seen from Figure B.2. Obviously, this inofficial "soft zone" was introduced to protect the krone from speculation; for a discussion see Bartolini and Prati (1999) and Ringbom (2003).

The change in the policy regime makes it necessary to split the observation period into three subperiods. The first one, consisting of the observations from October 1, 1986 to June 17, 1988, 431 observations in all, is one during which the crucial assumption in Krugman (1991) of interventions at the boundaries is satisfied. The second period consists of the observations until the end of 1988. According to Lysebo and Mundaca (1997) the wide fluctuations during this period are due to a falling oil price and domestic turbulence, and we leave them unmodelled. The third period with an inofficial zone, contains the observations from January 2, 1989, until the end, in total 449 observations.

We fit a STARTZ model to the first and third period and begin by the first one. The estimates of the parameters of the STARTZ model for the Norwegian krone can be found in Table B.1. Even here, the values of γ_a and γ_b the slope parameters of the transition functions for the conditional mean and variance respectively, have been set to 300 for the same reason as before. Note that the number of observations available for the estimation of parameters is much smaller than in the Swedish case. Functions G_L and G_U, in the conditional mean model, however, now have a shape different from the previous model, and the same is true for the model for the conditional variance. This can be seen, from Figure B.11(b), that the transition from the centre of the zone to the boundaries is smooth. The local random walk

behaviour apparent in Bank of Sweden's currency index is absent here: note that $\widehat{\varphi}_1 = 0.75$. Figure B.4 shows the value of the time-varying AR-coefficient as a function of the observations in the series.

As to the conditional variance specification, the heteroskedasticity consists of a decrease in the conditional variance to a small value when the values of the index change from the central parity to values close to the boundaries. There are no ARCH effects to speak of. As seen from Figure B.5, the conditional variance as a function of the location of the index in the zone has a distinct ∩-shape as the theory prescribes.

Statistics on the standardized residuals in Table B.2 are similar to what has been observed for the Swedish krona. The residuals again contain some outliers as the kurtosis exceeds five. The model passes the misspecification tests: the results appear in Table B.3, Table B.4 and Table B.5. The lowest p-value can be found in the test of parameter constancy but it is still relatively high, 0.07.

The marginal distribution of the exchange rate is investigated as before: 100000 observations are generated from the estimated STARTZ model and the empirical density function of the observations is smoothed using a standard kernel smoother. Figure B.8 shows a very clear ∪-shape with a hump in the centre. This is clear evidence in favour of Krugman's original model. When the interventions of the central bank mostly occurs at the boundaries, the theory predicts an ∪-shape, and this is exactly what we see in the figure. The hump in the middle is not anticipated by the basic Krugman model. It simply indicates that during the observation period the Norwegian index has spent less time in transit to and from the boundaries than either close to the boundaries and near the central parity. Again, the shape of the extreme tail where the density falls to zero should be interpreted with caution.

6.3. The Norwegian krone: 1989-1990

As explained earlier, Central Bank of Norway maintained an inofficial target zone within the official one from the beginning of 1989 till the ECU connection in October 1990. The behaviour of the exchange rate can now be expected to be quite different from what it in 1986-1988. The results from specifying and estimating a STARTZ model for the latest period appear in Table B.1. The currency index shows random walk type behaviour in the sense that $\widehat{\varphi}_1 + \widehat{\varphi}_2 = 0.98$ near the central parity. The transition is about as abrupt as in the Swedish case; see Figure B.11(c). The estimate $\widehat{\mu} = 0.21$, which together with the fact that the transition is abrupt gives us an estimate of the width of the implicit zone that is only about one fifth of the

official zone (about $\pm 0.45\%$). In the case of an implicit zone it is useful to check whether or not the zone is symmetric around the central parity, that is to test the nypothesis $\mu = \mu_{lower} = \mu_{upper}$. The likelihood ratio test against the alternative $\mu_{lower} \neq \mu_{upper}$, results in a p-value equal to 0.95. The conditional mean component of the model passes our misspecification tests.

The equation for the conditional variance at first sight seems like a standard GARCH(1,1) equation augmented with diminishing variance near the inofficial boundaries. However, the estimate of the ARCH parameter α_1 is not significant. If the coefficient were zero, then the GARCH model would not be identified. That again would mean that the standard deviation estimates, including the one for $\widehat{\alpha}_1$, would not be based on standard asymptotic theory. On the other hand, the number of observations is rather small, given the GARCH-type structure of the variance equation, and that could explain the large uncertainty of the estimate while in fact $\alpha_1 \neq 0$. The estimate of α_0 is also insignificant, but the previous arguments apply to it as well. Besides, in Table B.4 there is evidence (if we first assume that the model is identified) of instability of the parameters of the conditional variance model. Finally, it is seen from Figure B.6 that the conditional variance does not display any $\cap$-shape.

The marginal distribution of the index is again considered by generating 100000 observations from the estimated STARTZ model and smoothing the empirical density function with kernel estimation. The graph of the distinctly unimodal density in Figure B.9 deviates from the previous ones in the sense that the density decays smoothly to zero at the tails. Since the policy of Norges Bank was to keep the currency index inside an inofficial zone, there did not exist any need for a stiff ultimate line of defence, and this fact is clear from the figure. In other words, when the currency is defended well inside the official target zone, there is no reason to expect the dynamic properties of the currency index to correspond to the ones predicted by a standard target zone model.

7. Density forecasts

In order to illustrate the conditional distributions of the exchange rates we generated density forecasts from each model in turn from one up to 55 steps ahead. The models were simulated by drawing from the appropriate error distribution and computing the forecasts numerically as in Granger and Teräsvirta (1993). The forecasts have been generated from two starting points. One lies at the central parity, whereas the

other one lies close to the upper boundary of the zone. This means 1.45% for the Swedish krona and 2.15% for the Norwegian krone. For the latter index under the implicit zone regime, the starting point has been 0.5%. Following Wallis (1999), the density forecasts in Figure B.10 are presented as percentiles. The solid curve in the middle is the median, and the remaining ones are the $10, 20, ..., 90\%$ percentiles of the cumulative distribution.

The results show that the strongest mean reversion can be found for the Norwegian krone during the latter period when Norges Bank defended an inofficial zone. In this case the conditional densities, when the starting point is the central parity, are very concentrated even after 55 steps. Density forecasts for the Swedish krona agree with the previous results in that the mean reversion is weak. However, owing to the existence of the boundary, the forecasts densities are strongly skewed when the starting point of the index is the value near it. In that sense, one can speak of mean reversion but then, *median reversion* of the krona has been remarkably weak. Obviously, the model of Bekaert and Gray (1998) that builds on the idea of truncated densities, would yield similar results. The Norwegian krone has had a stronger tendency to fluctuate than the Swedish krona, and the densities are therefore flatter than in the Swedish case. As can be expected, median reversion is stronger for the former than the latter currency.

8. Conclusions

In the target zone literature, the emphasis has been on theoretical models. This paper proposes a rather flexible, empirical, time series model that is capable of characterizing the behaviour implied by theoretical target zone models. The model also enables the investigator to estimate the boundaries of an implicit band, should such a band exist. In order to model empirical data in a systematic way it is important to have a coherent modelling strategy and such a strategy is designed and applied to data here. A statistical advantage of the proposed strategy is that the misspecification tests we use only require standard asymptotic theory and are easy to perform.

The empirical examples indicate that there is structure in data that accords well with theoretical target zone models. For the Swedish krona the behaviour of the currency index within the target zone is in line with what theory suggests for a currency when the central bank intervenes intramarginally. In the case of the Norwegian krone 1986-88, where the Central Bank intervened only at the edges

of the band, the behaviour of the estimated model is in harmony with the results implied by the basic target zone model. For the remaining period of the Norwegian krone, consisting of observations in 1989-1990, the estimated model suggests that the implicit band maintained by the Central Bank during this period was not strictly enforced by the bank.

The STARTZ model may also be used to model other economic variables restricted by explicit or implicit boundaries, such as unemployment or interest rate series. The single-equation STARTZ model may also easily be made multivariate. In the present case this would allow the possibility of incorporating fundamentals into the time series target zone model, a topic which is left for further research.

References

Akaike, H. (1974). A new look at the statistical model identification, *IEEE Transactions of Automatic Control* **19**: 716–723.

Bartolini, L. and Prati, A. (1999). Soft exchange rate bands and speculative attacks: Theory, and evidence from the ERM since August 1993, *Journal of International Economics* **49**: 1–29.

Bekaert, G. and Gray, S. F. (1998). Target zones and exchange rates: An empirical investigation, *Journal of International Economics* **45**: 1–35.

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

Brooks, C. and Revéiz, A. H. (2002). A model for exchange rates with crawling bands - an application to the Colombian peso, *Journal of Economics and Business* **54**: 483–503.

Chung, C.-S. and Tauchen, G. (2001). Testing target-zone models using efficient method of moments, *Journal of Business and Economic Statistics* **19**: 255–269.

Darvas, Z. (1998). *Csúszó árfolyamrendszerek. Elmélet és nemzetközi tapasztalatok*, Magyar Nemzeti Bank, Budapest.

de Jong, F. (1994). A univariate analysis of EMS exchange rates using a target zone model, *Journal of Applied Econometrics* **9**: 31–45.

Delgado, F. and Dumas, B. (1992). Target zones, broad and narrow, *in* P. Krugman and M. Miller (eds), *Exchange Rate Targets and Currency Bands*, Cambridge University Press, Cambridge, pp. 35–56.

Eitrheim, Ø. and Teräsvirta, T. (1996). Testing the adequacy of smooth transition autoregressive models, *Journal of Econometrics* **74**: 59–75.

Fiorentini, G., Calzolari, G. and Panattoni, L. (1996). Analytic derivatives and the computation of GARCH estimates, *Journal of Applied Econometrics* **11**: 399–417.

Forbes, C. S. and Kofman, P. (2000). Bayesian target zones. Working Paper, Monash University.

Granger, C. W. J. and Teräsvirta, T. (1993). *Modelling Nonlinear Economic Relationships*, Oxford University Press, Oxford.

Iannizzotto, M. and Taylor, M. P. (1999). The target zone model, non-linearity and mean reversion: Is the honeymoon really over?, *The Economic Journal* **109**: C96–C110.

Klaster, M. A. and Knot, K. H. W. (2002). Toward an econometric target zone model with endogenous devaluation risk for a small open economy, *Economic Modelling* **19**: 509–529.

Krugman, P. R. (1991). Target zones and exchange rate dynamics, *Quarterly Journal of Economics* **106**: 669–682.

Lin, C.-F. and Teräsvirta, T. (1994). Testing the constancy of regression parameters against continuous structural change, *Journal of Econometrics* **62**: 211–228.

Lindberg, H. and Söderlind, P. (1994a). Intervention policy and mean reversion in exchange rate target zones: The Swedish case, *Scandinavian Journal of Economics* **96**: 499–513.

Lindberg, H. and Söderlind, P. (1994b). Testing the basic target zone model on Swedish data 1982-1990, *European Economic Review* **38**: 1441–1469.

Lundbergh, S. and Teräsvirta, T. (2002). Evaluating GARCH models, *Journal of Econometrics* **110**: 417–435.

Luukkonen, R., Saikkonen, P. and Teräsvirta, T. (1988). Testing linearity against smooth transition autoregressive models, *Biometrika* **75**: 491–499.

Lysebo, T. A. and Mundaca, B. G. (1997). Does the credibility of the exchange rate depend on the width of the target zone. Norges Bank.

Mundaca, B. G. (2000). The effect of interventions on realignment probabilities, *International Financial Markets, Institutions and Money* **10**: 323–347.

Nelder, J. A. (1961). The fitting of a generalisation of the logistic curve, *Biometrics* **17**: 89–110.

Richards, J. F. (1959). A flexible growth curve model for empirical use, *Journal of Experimental Botany* **10**: 290–300.

Ringbom, S. (2003). Narrow target zones within broad zones: A non-speculative exchange rate solution with limited resources, *Open Economies Review* **14**: 319–341.

Rissanen, J. (1978). Modeling by shortest data description, *Automatica* **14**: 465–471.

Schwarz, G. (1978). Estimating the dimension of a model, *Annals of Statistics* **6**: 461–464.

Smith, G. W. and Spencer, M. G. (1992). Estimation and testing in models of exchange rate target zones and process switching, *in* P. Krugman and M. Miller (eds), *Exchange Rate Targets and Currency Bands*, Cambridge: Cambridge University Press, pp. 211–239.

Sollis, R., Leybourne, S. and Newbold, P. (1999). Unit roots and asymmetric smooth transitions, *Journal of Time Series Analysis* **20**: 671–677.

Svensson, L. E. O. (1992). An interpretation of recent research on exchange rate target zones, *Journal of Economic Perspectives* **6**: 119–144.

Svensson, L. E. O. (1993). Assessing target zone credibility, *European Economic Review* **37**: 763–802.

Taylor, M. P. (1995). The economics of exchange rates, *Journal of Economic Literature* pp. 13–47.

Taylor, M. P. and Iannizzotto, M. (2001). On the mean-reverting properties of target zone exchange rates: A cautionary note, *Economic Letters* **71**: 117–129.

Teräsvirta, T. (1994). Specification, estimation, and evaluation of smooth transition autoregressive models, *Journal of the American Statistical Association* **89**: 208–218.

Torres, J. L. (2000). Stochastic intramarginal interventions in target zones, *International Financial Markets, Institutions and Money* **10**: 249–262.

Wallis, K. F. (1999). Asymmetric density forecasts and the Bank of England's fan chart, *National Institute Economic Review* **167**: 106–112.

Wooldridge, J. (1990). A unified approach to robust, regression-based specification tests, *Econometric Theory* **6**: 17–43.

A. Analytical derivatives

In this section we consider the analytical derivatives of the suggested model. These first-order derivatives are used in the estimation and later on in the evaluation of the estimated model. The derivatives of the model are straightforward to compute, see for example Fiorentini, Calzolari and Panattoni (1996). Let $\boldsymbol{\omega} = (\boldsymbol{\varphi}', \gamma_a, \theta_a, \mu, \boldsymbol{\eta}', \delta, \gamma_b, \theta_b)'$; consider the model defined by (3.1-3.5):

$$
\begin{aligned}
y_t &= m_t + \varepsilon_t \\
\varepsilon_t &= z_t \sqrt{h_t}
\end{aligned}
$$

where $m_t = m_t(\boldsymbol{\varphi}', \gamma_a, \theta_a, \mu)$ and $h_t = h_t(\boldsymbol{\varphi}', \gamma_a, \theta_a, \mu, \boldsymbol{\eta}', \delta, \gamma_b, \theta_b)$ are the functions of the conditional mean and the conditional variance. It is worth noting that $\boldsymbol{\varphi}', \gamma_a, \theta_a$ are associated with the conditional mean and that $\boldsymbol{\eta}', \delta, \gamma_b, \theta_b$ with the conditional variance, whereas the parameter μ is associated with both. Only in the special case when we condition the model on μ, do we have block diagonality between the conditional mean and variance. Assuming that $\{z_t\}$ is a sequence of independent standard normal errors, the log-likelihood function for observation t is:

$$
l_t = \text{const} - \frac{1}{2} \ln h_t - \frac{1}{2} \frac{\varepsilon_t^2}{h_t}
$$

where $\varepsilon_t = y_t - m_t$.

A.1. Partial derivative of l_t

The first-order partial derivative (the gradient) of the log likelihood function at time t is

$$
\frac{\partial l_t}{\partial \boldsymbol{\omega}'} = \frac{\varepsilon_t}{h_t} \frac{\partial m_t}{\partial \boldsymbol{\omega}'} + \frac{1}{2h_t} \left(\frac{\varepsilon_t^2}{h_t} - 1 \right) \frac{\partial h_t}{\partial \boldsymbol{\omega}'}
$$

The expectation of the matrix of the second-order partial derivatives of the log-likelihood function (the Hessian) at time t equals

$$
E \left[\frac{\partial^2 l_t}{\partial \boldsymbol{\omega} \partial \boldsymbol{\omega}'} \right] = -E \left[\frac{1}{h_t} \frac{\partial m_t}{\partial \boldsymbol{\omega}} \frac{\partial m_t}{\partial \boldsymbol{\omega}'} + \frac{1}{2h_t^2} \frac{\partial h_t}{\partial \boldsymbol{\omega}} \frac{\partial h_t}{\partial \boldsymbol{\omega}'} \right]
$$

A.2. Partial derivative of the conditional mean m_t

The conditional mean is defined as in (3.1) and (3.4). For notational convenience let $m_t^* = \boldsymbol{\varphi}' \mathbf{x}_t$, $\xi_a^L = \gamma_a(\mu s^L - y_{t-1})/2$ and $\xi_a^U = \gamma_a(y_{t-1} - \mu s^U)/2$. We can then

write the conditional mean as

$$m_t = m_t^* + (\mu s^L - m_t^*)G_a^L + (\mu s^U - m_t^*)G_a^U$$

where $G_a^L = \exp(\xi_a^L \theta_a)/(2\cosh\xi_a^L)^{\theta_a}$ and $G_a^U = \exp(\xi_a^U \theta_a)/(2\cosh\xi_a^U)^{\theta_a}$. The first-order partial derivatives with respect to the conditional mean parameters are

$$\begin{aligned}
\frac{\partial m_t}{\partial \boldsymbol{\varphi}'} &= (1 - H_a^L - H_a^U)\frac{\partial m_t^*}{\partial \boldsymbol{\varphi}'} \\
\frac{\partial m_t}{\partial \gamma_a} &= \frac{\theta_a e^{-\xi_a^L}}{(2\cosh\xi_a^L)}(-y_{t-1} + \mu s^L)(\mu s^L - m_t^*)G_a^L + \\
&\quad \frac{\theta_a e^{-\xi_a^U}}{(2\cosh\xi_a^U)}(y_{t-1} - \mu s^U)(\mu s^U - m_t^*)G_a^U \\
\frac{\partial m_t}{\partial \theta_a} &= (\xi_a^L - \ln(2\cosh\xi_a^L))(\mu s^L - m_t^*)G_a^L + \\
&\quad (\xi_a^U - \ln(2\cosh\xi_a^U))(\mu s^U - m_t^*)G_a^U.
\end{aligned}$$

The corresponding partial derivatives with respect to μ equals

$$\begin{aligned}
\frac{\partial m_t}{\partial \mu} &= (1 + \frac{\gamma_a \theta_a e^{-\xi_a^L}}{(2\cosh\xi_a^L)}(\mu s^L - m_t^*))s^L G_a^L + \\
&\quad (1 - \frac{\gamma_a \theta_a e^{-\xi_a^U}}{(2\cosh\xi_a^U)}(\mu s^U - m_t^*))s^U G_a^U.
\end{aligned}$$

The necessary derivatives evaluated under the null hypothesis used in the misspecification tests are easily obtained by replacing m_t^* with $m_t^{alt} = \boldsymbol{\varphi}'\mathbf{x}_t + A(\mathbf{x}_t; \boldsymbol{\pi}_a)$.

A.3. Partial derivative of the conditional variance h_t.

The conditional variance is parameterized as in (3.2) and (3.5). We rewrite the expression in the same way as the one for the conditional mean. Thus by setting $g_t = \boldsymbol{\eta}'\mathbf{w}_t$, $\xi_b^L = \gamma_b(-y_{t-1} + \mu s^L)/2$ and $\xi_b^U = \gamma_b(y_{t-1} - \mu s^U)/2$ the conditional variance becomes

$$h_t = g_t + (\delta - g_t)G_b^L + (\delta - g_t)G_b^U$$

where $G_b^L = \exp(\theta_b \xi_b^L)/(2\cosh\xi_b^L)^{\theta_b}$ and $G_b^U = \exp(\theta_b \xi_b^U)/(2\cosh\xi_b^U)^{\theta_b}$. To initialize the iterative computation of h_t, the conditional variance is estimated with the sample (unconditional) variance in the pre-sample case. This is done for all $t \leq 0$ by setting $h_t = \varepsilon_t^2 = \frac{1}{T}\sum_{i=1}^{T}\varepsilon_i^2$ where $\varepsilon_i = y_i - m_i$. The first-order derivatives may be computed iteratively by using the following expressions:

The derivatives with respect to conditional mean parameters:

$$\frac{\partial h_t}{\partial \boldsymbol{\varphi}'} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \boldsymbol{\varphi}'}$$

$$\frac{\partial h_t}{\partial \gamma_a} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \gamma_a}$$

$$\frac{\partial h_t}{\partial \theta_a} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \theta_a}$$

The derivatives with respect to μ:

$$\frac{\partial h_t}{\partial \mu} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \mu} +$$

$$\frac{\gamma_b \theta_b e^{-\xi_b^L}}{(2\cosh \xi_b^L)}(\delta - g_t)s^L G_b^L -$$

$$\frac{\gamma_b \theta_b e^{-\xi_b^U}}{(2\cosh \xi_b^U)}(\delta - g_t)s^U G_b^U$$

The derivatives with respect to the conditional variance parameters :

$$\frac{\partial h_t}{\partial \boldsymbol{\eta}'} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \boldsymbol{\eta}'}$$

$$\frac{\partial h_t}{\partial \delta} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \delta} + (G_b^L + G_b^U)$$

$$\frac{\partial h_t}{\partial \gamma_b} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \gamma_b} +$$

$$\frac{\theta_b e^{-\xi_b^L}}{(2\cosh \xi_b^L)}(-y_{t-1} + \mu s^L)(\delta - g_t)G_b^L +$$

$$\frac{\theta_b e^{-\xi_b^U}}{(2\cosh \xi_b^U)}(y_{t-1} - \mu s^U)(\delta - g_t)G_b^U$$

$$\frac{\partial h_t}{\partial \theta_b} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \theta_b} +$$

$$(\xi_b^L - \ln(2\cosh \xi_b^L))(\delta - g_t)G_b^L +$$

$$(\xi_b^U - \ln(2\cosh \xi_b^U))(\delta - g_t)G_b^U$$

where the necessary derivatives under the null hypothesis used in the evaluation tests are easily obtained by replacing g_t with $g_t^{alt} = \boldsymbol{\eta}'\mathbf{w}_t + B(\mathbf{w}_t; \boldsymbol{\pi}_b)$.

A.4. Partial derivatives of the mean and variance model

A.4.1. GARCH(1,1) type

Note that the GARCH type of model is constructed in such a way that the volatility model g_t is driven by the total volatility h_t (including the target zone structure).

$$g_t = \boldsymbol{\eta}' \mathbf{w}_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta h_{t-1}$$

where $\varepsilon_t = y_t - m_t$ where $m_t = m_t^* + (\mu s^L - m_t^*) H_a^L + (\mu s^U - m_t^*) H_a^U$. The derivatives then becomes:

$$
\begin{aligned}
\frac{\partial g_t}{\partial \varphi'} &= -2\alpha_1 \varepsilon_{t-1} \frac{\partial m_{t-1}}{\partial \varphi'} + \beta \frac{\partial h_{t-1}}{\partial \varphi'} \\
\frac{\partial g_t}{\partial \gamma_a} &= -2\alpha_1 \varepsilon_{t-1} \frac{\partial m_{t-1}}{\partial \gamma_a} + \beta \frac{\partial h_{t-1}}{\partial \gamma_a} \\
\frac{\partial g_t}{\partial \theta_a} &= -2\alpha_1 \varepsilon_{t-1} \frac{\partial m_{t-1}}{\partial \theta_a} + \beta \frac{\partial h_{t-1}}{\partial \theta_a} \\
\frac{\partial g_t}{\partial \mu} &= -2\alpha_1 \varepsilon_{t-1} \frac{\partial m_{t-1}}{\partial \mu} + \beta \frac{\partial h_{t-1}}{\partial \mu} \\
\frac{\partial g_t}{\partial \alpha_0} &= 1 + \beta \frac{\partial h_{t-1}}{\partial \alpha_0} \\
\frac{\partial g_t}{\partial \alpha_1} &= \varepsilon_{t-1}^2 + \beta \frac{\partial h_{t-1}}{\partial \alpha_1} \\
\frac{\partial g_t}{\partial \beta} &= h_{t-1}^* + \beta \frac{\partial h_{t-1}}{\partial \beta} \\
\frac{\partial g_t}{\partial \delta} &= \beta \frac{\partial h_{t-1}}{\partial \delta} \\
\frac{\partial g_t}{\partial \gamma_b} &= \beta \frac{\partial h_{t-1}}{\partial \gamma_b} \\
\frac{\partial g_t}{\partial \theta_b} &= \beta \frac{\partial h_{t-1}}{\partial \theta_b}
\end{aligned}
$$

A.5. Parameter constancy of μ and δ

The alternative to test μ against is $\mu_t = \mu^* + \pi_\mu t + R_3(\mathbf{x}_t; \rho_\mu, \mathbf{c}_\mu)$ and under the null hypothesis is $\pi_\mu = 0$. The partial derivatives, under H_0, with respect to π_μ are:

$$
\begin{aligned}
\frac{\partial m_t}{\partial \pi_\mu} &= (1 + \frac{\gamma_a \theta_a e^{-\xi_a^L}}{(2 \cosh \xi_a^L)} (\mu s^L - m_t^*)) ts^L G_a^L + \\
&\quad (1 - \frac{\gamma_a \theta_a e^{-\xi_a^U}}{(2 \cosh \xi_a^U)} (\mu s^U - m_t^*)) ts^U G_a^U \\
\frac{\partial h_t}{\partial \pi_\mu} &= (1 - G_b^L - G_b^U) \frac{\partial g_t}{\partial \pi_\mu} +
\end{aligned}
$$

$$\frac{\gamma_b \theta_b e^{-\xi_b^L}}{(2\cosh \xi_b^L)}(\delta - g_t)ts^L G_b^L -$$

$$\frac{\gamma_b \theta_b e^{-\xi_b^U}}{(2\cosh \xi_b^U)}(\delta - g_t)ts^U G_b^U$$

The alternative to test δ against is $\delta_t = \delta^* + \pi_\delta t + R_4(\mathbf{x}_t; \rho_\delta, \mathbf{c}_\delta)$ and under the null hypothesis is $\pi_\delta = 0$. The partial derivative, under H_0, with respect to π_μ are:

$$\frac{\partial h_t}{\partial \pi_\delta} = (1 - G_b^L - G_b^U)\frac{\partial g_t}{\partial \pi_\delta} + t(G_b^L + G_b^U)$$

B. Tables and Figures

Parameter	SEK (85-91)	NOK (86-88)	NOK (89-90)
φ_1	0.85 (0.030)	0.75 (0.065)	0.69 (0.048)
φ_2	0.15 (0.030)	.	0.29 (0.047)
α_0	0.00092 (0.0041)	0.047 (0.013)	0.00012 (0.0054)
α_1	0.17 (0.034)	0.10 (0.11)	0.0051 (0.071)
β_1	0.71 (0.029)	.	0.92 (0.025)
μ	0.99 (0.0052)	0.96 (0.0049)	0.21 (0.0057)
γ_a	300 (.)	300 (.)	300 (.)
θ_a	0.42 (0.82)	0.0049 (0.00092)	0.63 (1.5)
δ	0.0010 (0.00057)	0.0048 (0.00089)	0.0070 (0.0014)
γ_b	300 (.)	300 (.)	300 (.)
θ_b	0.11 (0.090)	0.032 (0.018)	0.083 (0.039)

Table B.1: Parameter estimates of the STARTZ models (standard deviations in parentheses) for the Swedish krona (SEK) and the Norwegian krone (NOK).

	SEK (85-91)	NOK (86-88)	NOK (89-90)
Characteristics:			
Min	-5.5	-4.5	-3.3
Max	4.8	4.7	2.4
Mean	-0.016	-0.028	0.058
Variance	1.0	1.0	1.0
Skewness	0.16	0.44	-0.25
Kurtosis	5.1	5.3	3.2

Table B.2: Characteristics of the standardized residuals of the STARTZ model.

	SEK (85-91)	NOK (86-88)	NOK (89-90)
Misspecified lag length (p-values)			
$p_a = 1$	0.47	0.92	0.22
$p_a = 2$	0.38	0.86	0.65
Parameter constancy (p-values)			
	0.65	0.96	0.27
Remaining nonlinearity (p-values)			
	0.46	0.84	0.58

Table B.3: p-values of specification tests for the conditional mean for the estimated STARTZ model. LM tests for the conditional mean. The test of misspecified lag lengths is computed against the alternative of additional lags up to the given lag, p_a. The test of parameter constancy is computed against an alternative of time dependence given by a logistic function with time as the transition variable. The test against nonlinearity is of LSTAR type.

	SEK (85-91)	NOK (86-88)	NOK (89-90)
Misspecified lag length (p-values)			
$p_b = 1$	0.15	0.08	0.36
$p_b = 2$	0.34	0.22	0.34
$p_b = 3$	0.27	0.23	0.50
Parameter constancy (p-values)			
ARCH/GARCH	0.55	0.07	0.01
δ	0.89	0.95	0.51
Remaining nonlinearity of STGARCH type (p-values)			
All parameters	0.48	0.11	0.27
Without intercept	0.38	0.13	0.11

Table B.4: p-values of specification tests for the conditional variance of the estimated STARTZ model. LM tests for the conditional variance. The tests of no remaining serial dependence in the squared and standardized residuals are computed against the alternative of remaining dependence up to the given lag, p_b. The test of parameter constancy is computed against an alternative of time dependence given by a logistic function with time as the transition variable. The test against remaining nonlinearity is a test against an STGARCH structure.

	SEK (85-91)	NOK (86-88)	NOK (89-90)
Parameter constancy (p-values)			
μ	0.67	0.37	0.69

Table B.5: p-values of specification tests for both the conditional mean and variance of the estimated STARTZ model. LM tests for parameter constancy is computed against an alternative of time dependence given by a logistic function with time as the transition variable.
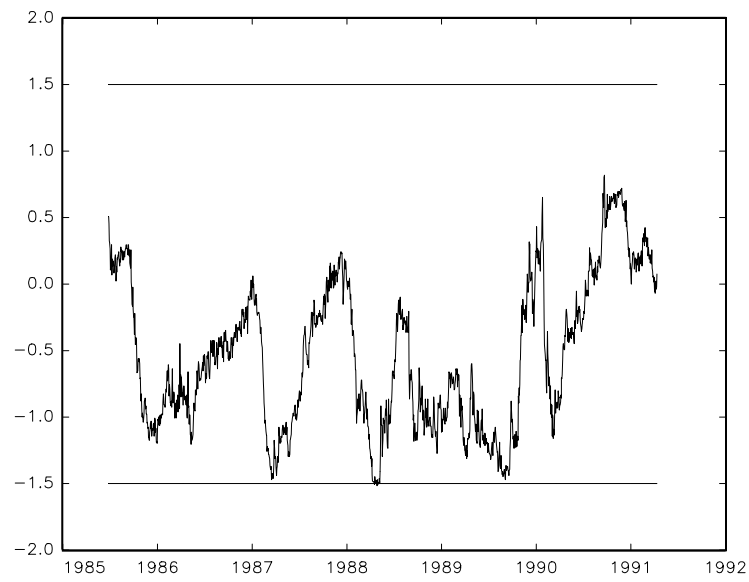
Figure B.1: The deviation (in percent) from central parity for the daily Swedish exchange rate index, July 1, 1985 to May 17, 1991. During this period the central parity of the target zone was 132 and the exchange rate index was allowed to vary within ±1.5 percent from the central parity.
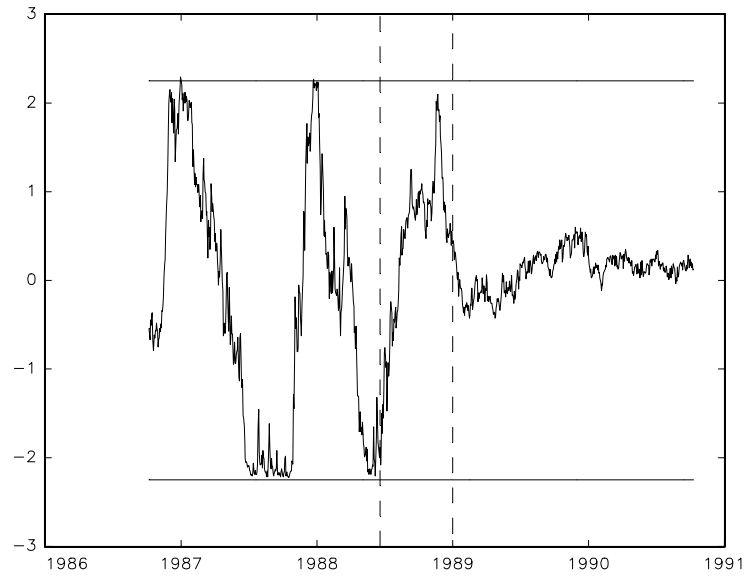
Figure B.2: The deviation (in percent) from the central parity for the daily Norwegian exchange rate index, October 1, 1986 to October 19, 1990. During this period the central parity of target zone was 112 and the exchange rate index was allowed to vary within ±2.25 percent from the central parity. The first dashed line corresponds to June 17, 1988. At that date the authorities changed their intervention policy. The second dashed line corresponds to January 2, 1989.
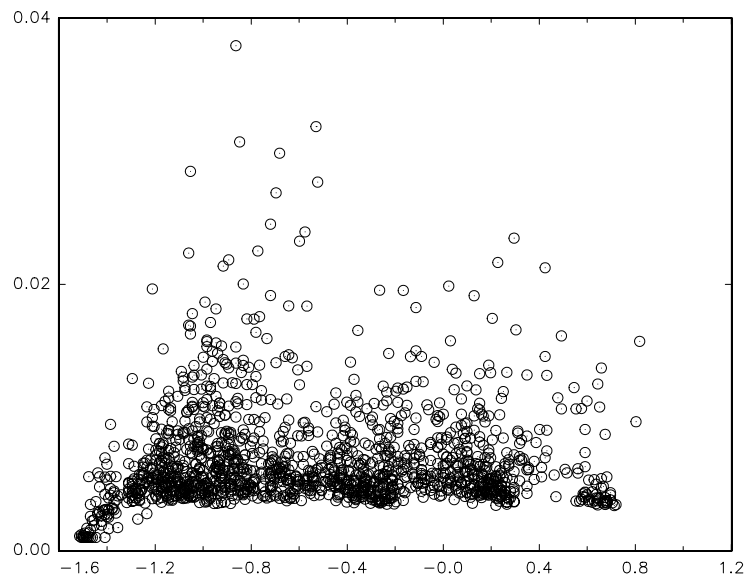
Figure B.3: The daily Swedish exchange rate index, July 1, 1985 to May 17, 1991. The parameterization of conditional variance, h_t, on the y-axis is plotted against the observed deviation from the central parity (in percent) on the x-axis.
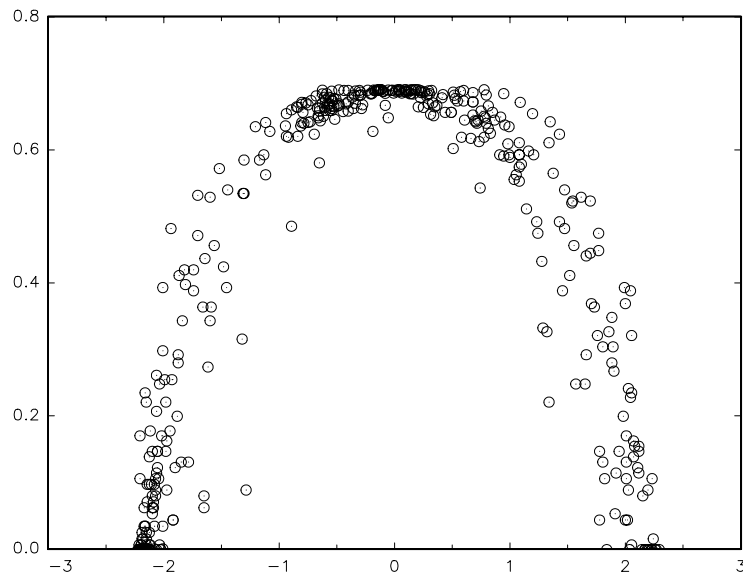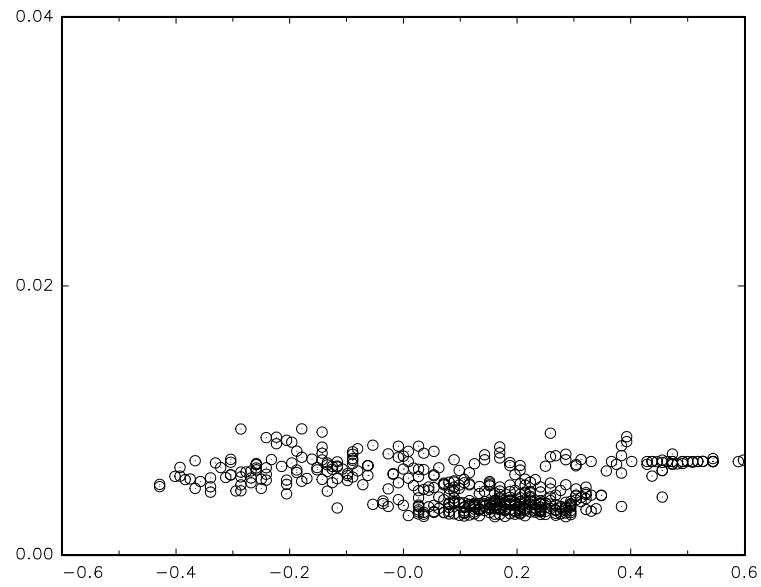
Figure B.4: The daily Norwegian exchange rate index, October 1, 1986 to June 17, 1988. The value of the restricted parameter in the conditional mean $\left[\varphi' - \varphi'G^L - \varphi'G^U\right]$ on the y-axis is plotted against the observed deviation from the central parity (in percent) on the x-axis.
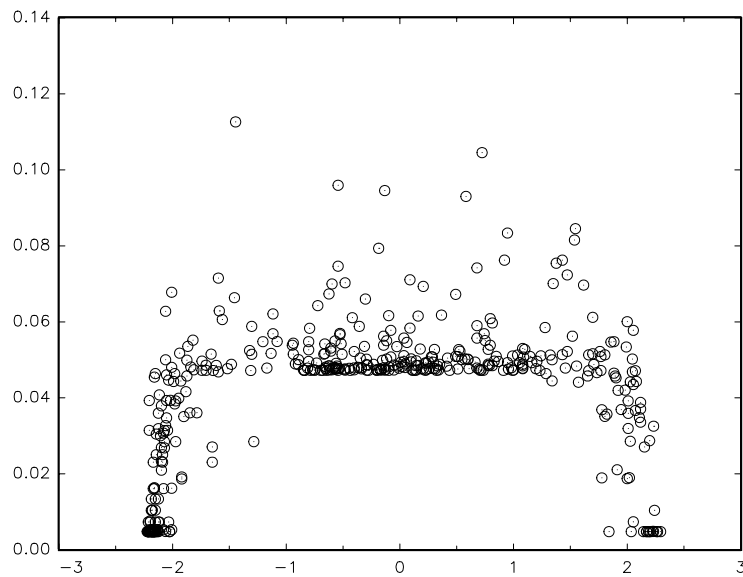
Figure B.5: The daily Norwegian exchange rate index, October 1, 1986 to June 17, 1988. The parameterization of conditional variance, h_t, on the y-axis is plotted against the observed deviation from the central parity (in percent) on the x-axis.

Figure B.6: The daily Norwegian exchange rate index, January 2, 1989 to October 21, 1990. The parameterization of conditional variance, h_t, on the y-axis is plotted against the observed deviation from the central parity (in percent) on the x-axis.
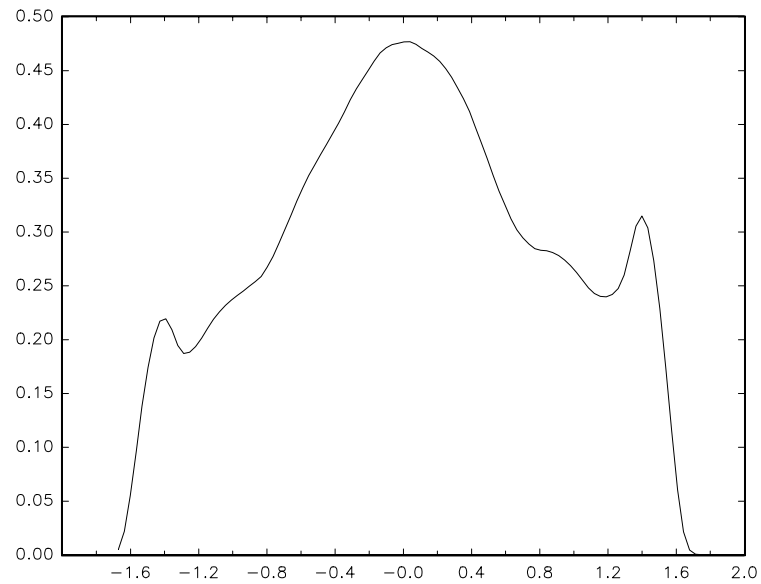
Figure B.7: Simulated marginal density from the STARTZ model for the Swedish krona for the period July 1, 1985 to May 17, 1991. A kernel estimate of the marginal density based on 100000 generated data points is plotted in the figure.
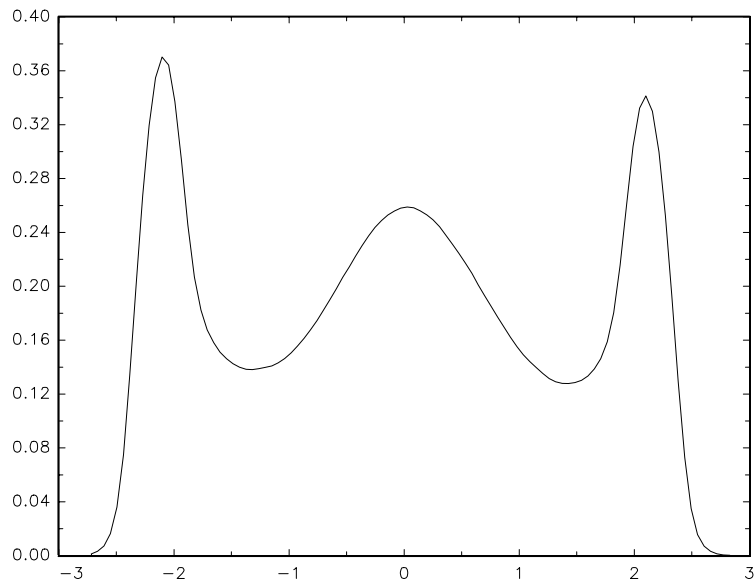
Figure B.8: Simulated marginal density from the STARTZ model for the Norwegian krone in the period October 1, 1986 to June 17, 1988. A kernel estimate of the marginal density based on 100000 generated data points is plotted in the figure.
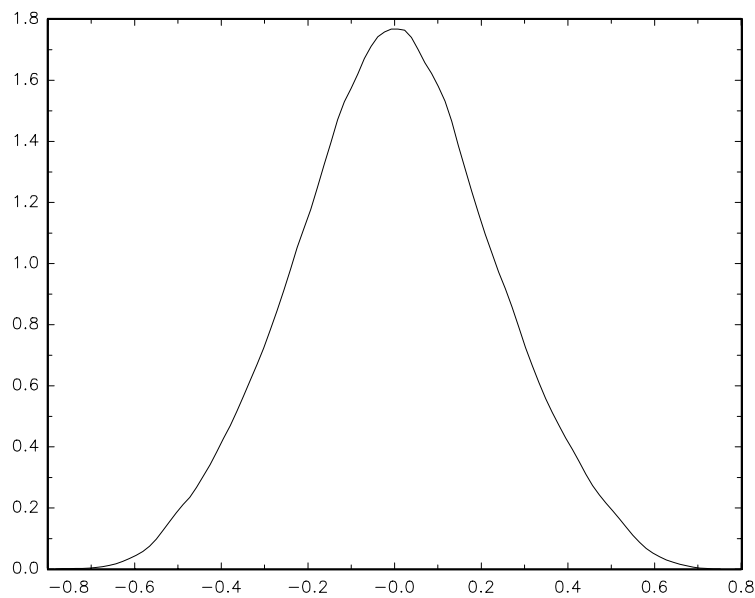

Figure B.9: Simulated marginal density for the STARTZ model for the Norwegian krone for the period January 2, 1989 to October 21, 1990. A kernel estimate of the marginal density based on 100000 generated data points is plotted in the figure.

(a) Fan charts for the Swedish krona:1985-1991

(b) Fan charts for the Norwegian krone: 1986-1988

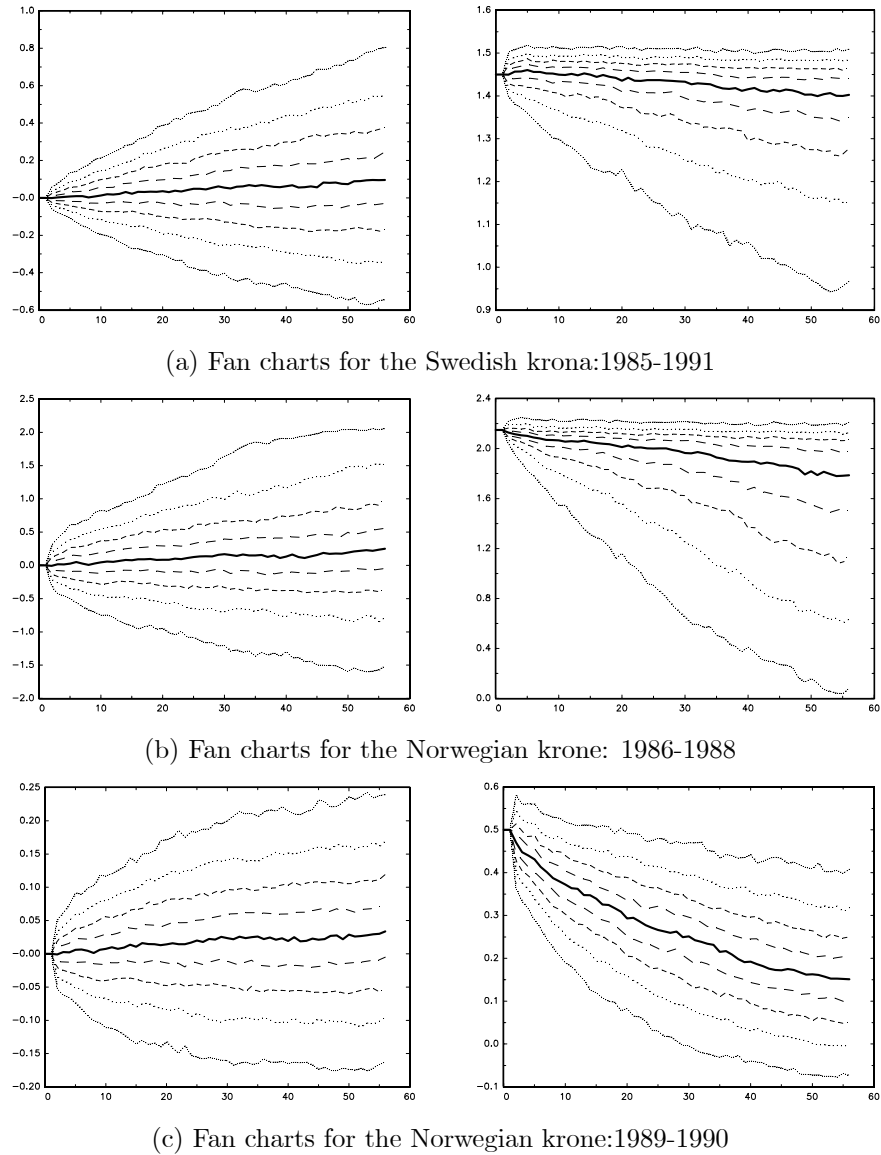(c) Fan charts for the Norwegian krone:1989-1990

Figure B.10: Fan charts of forecast densities from one to 55 steps ahead for the Swedish krona and the Norwegian krone. The fan charts on the left-hand side have the central parity as the starting point. The fan charts on the right-hand side represent forecasts given that the index is close to the official or inoffical (lowest panel) boundary. The deciles on the y-axis are plotted against the number of steps ahead on the x-axis. The solid line in the middle is the median forecast. The fan charts are generated using 1000 independent realisations for 55 steps ahead.
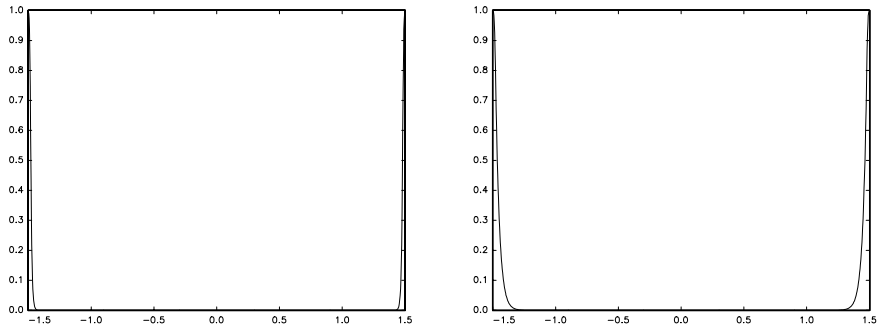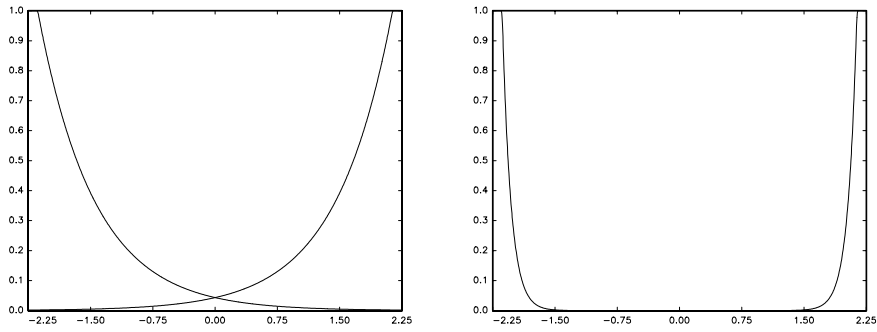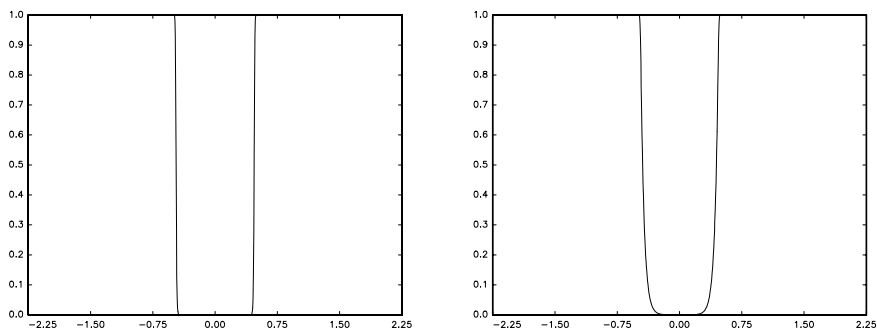
(a) Estimated functions G^L and G^U: Swedish krona 1985-1991

(b) Estimated functions G^L and G^U: Norwegian krone 1986-1988

(c) Estimated functions G^L and G^U: Norwegian krone 1989-1990

Figure B.11: Left-hand panels: Transition functions G^L and G^U for the conditional mean process. Right-hand panels: Transition functions G^L and G^U for the conditional variance process. Value of the transition on the vertical axis and the deviation from the central parity (in percent) on the horizontal axis.