

**Research Unit for Statistical
and Empirical Analysis in Social Sciences (Hi-Stat)**

**A General Framework for Observation Driven
Time-Varying Parameter Models**

Drew Creal
Siem Jan Koopman
André Lucas

March 2009

A General Framework for Observation Driven Time-Varying Parameter Models*

Drew Creal^a, Siem Jan Koopman^{a,c}, André Lucas^{b,c}

^(a) *Department of Econometrics, VU University Amsterdam*

^(b) *Department of Finance, VU University Amsterdam, and Duisenberg school of finance*

^(c) *Tinbergen Institute, Amsterdam*

October 23, 2008

Abstract

We propose a new class of observation driven time series models that we refer to as Generalized Autoregressive Score (GAS) models. The driving mechanism of the GAS model is the scaled likelihood score. This provides a unified and consistent framework for introducing time-varying parameters in a wide class of non-linear models. The GAS model encompasses other well-known models such as the generalized autoregressive conditional heteroskedasticity, autoregressive conditional duration, autoregressive conditional intensity and single source of error models. In addition, the GAS specification gives rise to a wide range of new observation driven models. Examples include non-linear regression models with time-varying parameters, observation driven analogues of unobserved components time series models, multivariate point process models with time-varying parameters and pooling restrictions, new models for time-varying copula functions and models for time-varying higher order moments. We study the properties of GAS models and provide several non-trivial examples of their application.

Keywords: dynamic models, time-varying parameters, non-linearity, exponential family, marked point processes, copulas.

JEL classification codes: C10, C22, C32, C51.

*Corresponding author: André Lucas, VU University Amsterdam, FEWEB, De Boelelaan 1105, 1081 HV Amsterdam, Netherlands, phone: +31 20 598 6039, fax: +31 20 598 6020, email: alucas@feweb.vu.nl.

1 Introduction

Time series models with time-varying parameters have been categorized by Cox (1981) into two classes: observation driven models and parameter driven models. In this paper we develop a new, general framework for building observation driven time series models. In the observation driven approach, time variation of the parameters is introduced by making them dependent on (functions of) their own lagged values, past observations, and exogenous variables. Although the parameters are stochastic, they are perfectly predictable given past information. This greatly simplifies likelihood evaluation and explains why the models have become popular in the applied econometrics and statistics literature. Typical examples of observation driven models are the generalized autoregressive conditional heteroskedasticity (GARCH) model of Engle (1982), Bollerslev (1986) and Engle and Bollerslev (1986), the autoregressive conditional duration and intensity (ACD and ACI, respectively) models of Engle and Russell (1998) and Russell (2001), the dynamic conditional correlation (DCC) model of Engle (2002a) and the Poisson count model discussed by Davis, Dunsmuir, and Streett (2003). Our approach encompasses many of the existing observation driven models as mentioned above. In addition, it gives rise to a wide range of new models.

The alternative to observation driven models are parameter driven models. In parameter driven models, the parameters are stochastic processes which are subject to their own source of error. Given past and concurrent observations, the parameters are not perfectly predictable. Typical examples include the stochastic volatility (SV) model, see Shephard (2005), and the stochastic intensity models of Bauwens and Hautsch (2006) and Koopman, Lucas, and Monteiro (2008). Estimation is usually more involved for these models because the associated likelihood functions are not available in closed-form. Exceptions include linear Gaussian state space models and discrete-state hidden Markov models, see Harvey (1989) and Hamilton (1989). In most other cases, computing the likelihood function requires evaluating a high-dimensional integral usually by simulation methods such as importance sampling and Markov chain Monte Carlo; see, e.g., Shephard and Pitt (1997). However, parameter driven models offer a conceptually straightforward way of introducing time-varying parameters in a wide class of non-linear and non-Gaussian models.

It is widely accepted that in many settings of empirical interest, time variation in a selection of parameters of the model is important for capturing the dynamic behavior of (multivariate) time series processes. Obvious examples include the models for volatility clustering as often found in financial return data, see, e.g., the surveys of Bollerslev, Engle, and Nelson (1994), Ghysels, Harvey, and Renault (1996), Engle (2001), Shephard (2005) and Bauwens, Laurent,

and Rombouts (2006). There are also other examples. In the case of univariate and multivariate (marked) point processes, the intensities can be made time-varying, see, e.g., Engle and Russell (1998), Russell (2001), Bauwens and Hautsch (2006) and Koopman, Lucas, and Monteiro (2008). In the case of copula models, the dependence parameters can be considered as time-varying, see Dias and Embrechts (2004), van den Goorbergh, Genest, and Werker (2005), and, in particular, Patton (2006). The higher order moments of the distribution can also be made to vary over time, see Hansen (1994), Harvey and Siddique (1999), Jondeau and Rockinger (2003) and Brooks, Burke, Heravi, and Persaud (2005). Correlations in multivariate density functions can be allowed to vary over time, see Engle (2002a). Finally, researchers have considered settings where the density itself or its quantiles are time-varying, see Park and Qian (2006) and Engle and Manganelli (2004).

The main contribution of this paper is the proposal of a common framework for developing observation driven models with time-varying parameters. The main difficulty in formulating a unified framework lies in the choice of a function that links the past observations to future parameter values. Such a function should be applicable to a wide class of non-linear and non-Gaussian models. In this paper, we argue that the scaled score of the model density at time t is an effective choice for the driving mechanism of the time-varying parameters. By choosing the scaling appropriately, standard observation driven models such as the GARCH, ACD, and ACI models are recovered. The scaled score is equally applicable to non-standard models and leads to the formulation of new observation driven models that mimic the behavior of parameter driven models for which no observation driven equivalents have been specified.

We will refer to our observation driven model with a scaled score function as the generalized autoregressive score (GAS) model. The GAS model has similar advantages as the GARCH model. Likelihood evaluation is straightforward. Extensions to asymmetric, long memory, and other more complicated dynamics can be considered without introducing further complexities. Other general frameworks for observation driven models have been suggested in the literature including the generalized linear autoregressive (GLAR) models of Shephard (1995), the generalized autoregressive moving average (GARMA) models of Benjamin, Rigby, and Stanispoulos (2003) and the vector multiplicative error models (MEM) of Cipollini, Engle, and Gallo (2006). In contrast to these proposals, GAS models exploit the complete density structure rather than only means and higher moments.

To illustrate the applicability of GAS models, we study a number of interesting, non-trivial settings for observation driven models. We consider linear and nonlinear regression models with time-varying coefficients a typical class of models that we can treat within the GAS framework. An example is the Nelson and Siegel (1987) model for analyzing the term structure of inter-

est rates which emphasizes that GAS can also treat multivariate models, see Subsection 4.1. Multivariate non-Gaussian models for pooled marked point-processes with a GAS specification for latent factors driving the log-intensities is a new model specification that can be used for the modeling of credit rating transitions, see Subsection 4.2. The GAS representations of unobserved components time series models are discussed in Subsections 4.3 and 4.4. A new class of time-varying copulas models based on the GAS framework is formulated in Subsection 4.5. A challenging direction of research is to model higher-order moments of financial returns as time-varying processes. In Subsection 4.6 we show that GAS provides a generic tool to develop models that have time-varying higher-order moments. A particular case is to consider linear and GARCH models with Student t distributions where the degrees of freedom parameter is made time-varying. Observing trade by trade transaction prices on a discrete grid leads to some interesting research directions. For example, price changes can be viewed as realizations of a multinomial distribution which need to be subject to time-varying processes. In Subsection 4.7, we discuss a GAS treatment for the modeling of discrete price changes. Finally, in Subsection 4.8 we discuss a methodology for dynamic mixture distributions based on the GAS framework.

The remainder of the paper is organised as follows. In Section 2 we provide the basic GAS specification together with a set of motivating examples. Section 3 includes a discussion of the statistical properties of GAS models. Section 4 contains a range of non-trivial examples of GAS models, where we develop several new observation driven models. In Section 5 we provide simulation evidence for the statistical properties of the estimators. Finally, Section 6 concludes and provides directions for future research.

2 The model

2.1 Basic model specification

Denote y_t as the dependent variable of interest, f_t as the time-varying parameter vector, x_t as a vector of exogenous variables (covariates), all at time t , and θ as a vector of static parameters. The available information set at time t consists of $Y_1^{t-1} = \{y_1, \dots, y_{t-1}\}$, $F_1^{t-1} = \{f_1, \dots, f_{t-1}\}$ and $X_1^t = \{x_1, \dots, x_t\}$. We assume that observation y_t is generated by the observation density

$$p(y_t | f_{t-1}, Y_1^{t-1}, X_1^t, F_1^{t-2}; \theta). \quad (1)$$

Furthermore, we assume that the mechanism for updating the time-varying parameter f_t is given by the familiar autoregressive updating equation

$$f_t = \omega + \sum_{i=0}^{p-1} A_i s_{t-i} + \sum_{j=1}^q B_j f_{t-j}, \quad (2)$$

where ω is a vector of constants, coefficient matrices A_i and B_j have appropriate dimensions for $i = 0, \dots, p-1$ and $j = 1, \dots, q$ while s_t is the scaled score function and depends on past observations Y_1^{t-1} , the time-varying parameters in F_1^{t-1} , and the static parameter θ . Furthermore, the unknown coefficients in (2) are functions of θ as well, that is $\omega = \omega(\theta)$, $A_i = A_i(\theta)$ and $B_j = B_j(\theta)$. Our main contribution is the particular choice for the driving mechanism s_t that is applicable uniformly over a wide class of densities and, possibly non-linear, models.

Equation (1) differs from a typical parameter driven model specification. In a parameter driven model, the parameter f_t would evolve subject to its own independent source of error, say η_t . In particular, s_t would be replaced by η_t in (2). Estimation of f_t is then based on the conditional (filtered) density $p(f_t|Y_1^t, X_1^t, F_1^{t-1}; \theta)$. For linear Gaussian state space models, this density can be computed in closed form by the Kalman filter. In non-linear and non-Gaussian models, conditional densities are generally evaluated via simulation methods; see, e.g., Durbin and Koopman (2001) and Doucet, de Freitas, and Gordon (2001). The simulations are often most effective when they are based on second order expansions of the log observation density (1). We use the same intuition to update the time-varying parameter from f_{t-1} to f_t via (2) with

$$s_t = S(t, Y_1^{t-1}, F_1^{t-1}) \cdot \nabla_t = S_{t-1} \cdot \nabla_t, \quad (3)$$

where

$$\nabla_t = \partial \ln p(y_t|f_{t-1}, Y_1^{t-1}, X_1^t, F_1^{t-2}; \theta) / \partial f_{t-1}, \quad S_{t-1} = S(t, Y_1^{t-1}, F_1^{t-1}; \theta), \quad (4)$$

with time dependent scaling matrix $S(\cdot)$. Given the reliance of the driving mechanism in (2) on the scaled score vector (3), we let the equations (1) – (3) constitute the generalized autoregressive score model with orders p and q which we abbreviate by GAS(p, q).

There are several intuitive choices for the scaling matrix that we investigate here. Our first choice is to set S_{t-1} equal to the (pseudo)-inverse information matrix based on the density (1), that is

$$S_{t-1} = \mathcal{I}_{t-1}^{-1} = E_{t-1} [\nabla_t \nabla_t']^{-1} = -E_{t-1} \left[\frac{\partial^2 \ln p(y_t|f_{t-1}, Y_1^{t-1}, X_1^t, F_1^{t-2}; \theta)}{\partial f_{t-1} \partial f_{t-1}'} \right]^{-1}. \quad (5)$$

The updating mechanism (2) for f_t reduces closely to a Gauss-Newton updating step for every new observation y_t that becomes available through time. Using this particular choice for scaling the score vector, the GAS model encompasses the well-known observation driven GARCH, ACD, and ACI models as well as most of the Poisson count models considered by Davis et al. (2003). When the scaling matrix is the identity matrix, that is $S_{t-1} = I$ in (3), the recursion captures models such as the autoregressive conditional multinomial (ACM) model of Russell and Engle (2005). In addition, it gives rise to a number of useful observation driven models that

have not been investigated before. We first give some introductory examples of GAS models. In Section 4, we provide a more systematic review of new, non-trivial models within the GAS family that provide interesting directions for future research.

2.2 Some examples

Example 1 (GARCH models): Consider the model $y_t = \sigma_{t-1}\varepsilon_t$ where the Gaussian disturbance ε_t has mean zero and variance one while σ_t is a time-varying standard deviation. It is a basic exercise to show that the GAS(1,1) model with $S_{t-1} = \mathcal{I}_{t-1}^{-1}$ for $f_t = \sigma_t^2$ reduces to

$$f_t = \omega + A_0 (y_t^2 - f_{t-1}) + B_1 f_{t-1}, \quad (6)$$

which is equivalent to the standard GARCH(1,1) specification of Bollerslev (1986). However, if we assume that ε_t follows a Student t distribution scaled to have variance one and with ν degrees of freedom, that is $\varepsilon_t \sim t_\nu$, the GAS(1,1) specification for the conditional variance leads to the updating equation

$$f_t = \omega + A_0 (1 + 3\nu^{-1}) \cdot \left(\frac{(1 + \nu^{-1})}{(1 - 2\nu^{-1})(1 + \nu^{-1}y_t^2 / (1 - 2\nu^{-1}) f_{t-1})} y_t^2 - f_{t-1} \right) + B_1 f_{t-1}. \quad (7)$$

This collapses to (6) in case of the Gaussian distribution when $\nu^{-1} = 0$. The recursion in (7), however, has an important difference with the standard t-GARCH(1,1) model of Bollerslev (1987) which has the Student t density in (1) and the updating equation (6). The denominator of the second term in the right-hand side of (7) causes a more moderate increase in the variance for a large realization of $|y_t|$ as long as ν is finite. The intuition is clear: if the errors are modeled by a fat-tailed distribution, a large realization in y_t does not necessitate a substantial increase in the variance. The GAS updating mechanism for the model with Student t errors therefore is substantially different from its familiar GARCH counterpart. We return to this example in more detail in Section 4.

Example 2 (MEM, ACI, and ACD models): Consider the model $y_t = \mu_{t-1}\varepsilon_t$ where ε_t has a gamma distribution with shape and rate parameters both expressed as a parameter γ . Given standard properties of the gamma distribution, μ_{t-1} is the conditional mean of y_t . Let $f_t = \mu_t$. The GAS(1,1) model in this case is

$$f_t = \omega + A_0/2 (y_t - f_{t-1}) + B_1 f_{t-1}, \quad (8)$$

which is equivalent to the multiplicative error model (MEM) proposed by Engle (2002b) and extended in Engle and Gallo (2006). The exponential distribution is a special case of the gamma distribution making the ACD and ACI models a special case of MEM and consequently GAS.

To see this, suppose y_t is exponentially distributed with time-varying mean $\lambda_t = \exp(f_{t-1})$ and density function $p(y_t|f_{t-1}) = \lambda_t \exp(-\lambda_t y_t)$. The GAS(1,1) model has the updating equation

$$f_t = \omega + A_0 (1 - y_t \exp(f_{t-1})) + B_1 f_{t-1}, \quad (9)$$

which is equivalent to the standard ACI(1, 1) model of Russell (2001). If we parameterize the expected duration rather than the intensity, so that $p(y_t|f_{t-1}) = \lambda_t^{-1} \exp(-\lambda_t^{-1} y_t)$, and define the factor as $\lambda_t = \exp(-f_{t-1})$, we obtain

$$f_t = \omega + A_0 y_t / \exp(f_{t-1}) + B_1 f_{t-1}, \quad (10)$$

which is equivalent to the standard ACD(1, 1) model of Engle and Russell (1998), see also Russell and Engle (2005).

Example 3 (regression models): We next consider the linear regression model $y_t = x_t \beta_{t-1} + \varepsilon_t$ with normally distributed disturbances $\varepsilon_t \sim N(0, \sigma^2)$. Let $f_t = \beta_t$, then the scaled score function with $S_{t-1} = \mathcal{I}_{t-1}^{-1}$ is

$$s_t = x_t (y_t - x_t f_{t-1}) / (x_t' x_t), \quad (11)$$

and the GAS(1, 1) specification becomes

$$f_t = \omega + A_0 x_t \frac{y_t - x_t f_{t-1}}{x_t' x_t} + B_1 f_{t-1}. \quad (12)$$

For the time-varying intercept β_t with $x_t \equiv 1$, the updating equation (12) reduces to the exponentially weighted moving average (EWMA) recursion by setting $\omega = 0$ and $B_1 = 1$, that is

$$f_t = f_{t-1} + A_0 (y_t - f_{t-1}). \quad (13)$$

In this case, we obtain the observation driven analogue of the local level (parameter driven) model,

$$y_t = \mu_{t-1} + \varepsilon_t, \quad \mu_t = \mu_{t-1} + \eta_t,$$

where the unobserved level component μ_t is modeled by a random walk process and the disturbances ε_t and η_t are mutually and serially independent, and normally distributed, see Durbin and Koopman (2001). A direct link between the parameter and observation driven models is established when we set $\eta_t = \alpha(y_t - \mu_{t-1}) = \alpha \varepsilon_t$ while in (13) we set $\alpha \equiv A_0$ and consider f_{t-1} as the (filtered) estimate of μ_{t-1} . The local level model illustration shows that the GAS representation is closely related to the single source of error (SSOE) framework as advocated by Ord, Koehler, and Snyder (1997). However, the GAS framework allows straightforward extensions for this class of models. For example, the EWMA scheme in (13) can be extended

by including σ^2 as a time-varying factor and recomputing the scaled score function in (11) for the new time-varying parameter vector $f_{t-1} = (\beta_{t-1}, \sigma_{t-1}^2)'$.

We note that for the linear regression model, the updating function (12), with $\omega = 0$ and $B_1 = I$, does not reduce to the recursive least squares update which is given by

$$f_t = P_{t-1} x_t \frac{y_t - x_t f_{t-1}}{x_t' P_{t-1} x_t + \sigma^2} + f_{t-1}, \quad \text{with} \quad P_t = \sigma^2 \left(\sum_{j=1}^t x_j' x_j \right)^{-1}.$$

The GAS updating function (12) however reveals that if $x_t' x_t$ is close to zero, the GAS driving mechanism might become unstable. In this case, the recursive least squares update remains stable, at least when σ^2 is sufficiently large. As a remedy for such instabilities, we provide an information smoothed variant of the GAS driving mechanism which we discuss in the next section. Alternatively, we may want to consider using the identity matrix to scale the score with $S_{t-1} = I$ and $s_t = x_t(y_t - x_t f_{t-1})$.

Example 4 (Dynamic exponential family models): The exponential family of distributions has a long history in statistics due to many convenient mathematical properties. Parameter driven models for some members of this family, like the normal (GARCH), the exponential (ACI and ACD), and the multinomial (ACM), have been developed separately in the literature. Shephard (1995) and Benjamin et al. (2003) proposed parameter-driven models for other distributions in this family. The GAS framework introduces time-varying parameters for other members of the exponential family in a natural way. As mentioned earlier, each GAS model also has the advantage that it exploits the full density structure to update the time-varying parameters.

The main obstacle for using GAS models is often computing the information matrix given the parameterization used. To facilitate this task, we provide Table 1 which contains the gradients and elements of the information matrix for a variety of models in the exponential family. In addition to the GARCH and MEM classes of models, the GAS framework also encompasses the time-varying binomial models of Cox (1958) and Rydberg and Shephard (2003), the ACM model of Russell and Engle (2005), and some of the Poisson models in Davis et al. (2003). The latter three models can be obtained by scaling the relevant score vector from Table 1 with either an identity scaling matrix $S_{t-1} = I$ or the matrix square root of $S_{t-1} = \mathcal{I}_{t-1}^{-1}$.

2.3 Other GAS specifications

An important advantage of the GAS(p, q) specification is that its applicability is not restricted to one specific model or choice of model parameterization. In contrast, the recursion scheme

Table 1: Dynamics for some Exponential Family Distributions

Distribution	f_t	∇_t	\mathcal{I}_t
Normal (1) $\frac{\exp(-0.5(y-\mu)^2)}{(2\pi\sigma^2)^{1/2}}$	μ_t σ_t^2	$0.5(y_t - \mu_t)/\sigma_t^2$ $-0.5\sigma_t^{-2} + 0.5\sigma_t^{-4}(y_t - \mu_t)^2$	$\mathcal{I}_{t,11} = 0.5\sigma_t^{-2}$ $\mathcal{I}_{t,22} = 0.5\sigma_t^{-4}$ $\mathcal{I}_{t,12} = 0$
Normal (2) $\frac{\exp(-0.5(y-\mu)^2)}{(2\pi\sigma^2)^{1/2}}$	μ_t $\ln(\sigma_t^2)$	$0.5(y_t - \mu_t)/\sigma_t^2$ $-0.5 + 0.5\sigma_t^{-2}(y_t - \mu_t)^2$	$\mathcal{I}_{t,11} = 0.5\sigma_t^{-2}$ $\mathcal{I}_{t,22} = 0.5$ $\mathcal{I}_{t,12} = 0$
Exponential $\lambda \exp(-\lambda y)$	$\ln(\lambda_t)$	$1 - \lambda_t y_t$	$\mathcal{I}_t = 1$
Gamma $\frac{y^{\alpha-1} \exp(-y/\beta)}{\beta^\alpha \Gamma(\alpha)}$	$\ln(\alpha_t)$ $\ln(\beta_t)$	$\alpha_t (\ln(y_t) - \ln(\beta_t) - \Psi(\alpha_t, 1))$ $y_t/\beta_t - \alpha_t$	$\mathcal{I}_{t,11} = \alpha_t^2 \Psi(\alpha_t, 2)$ $\mathcal{I}_{t,22} = \alpha_t$ $\mathcal{I}_{t,12} = \alpha_t$
Dirichlet	$\ln(\alpha_{it})$	$\alpha_{it} (\Psi(\sum \alpha_{jt}, 1) - \Psi(\alpha_{it}, 1))$ $+ \alpha_{it} \ln(y_{it})$	$\mathcal{I}_{t,ii} = \alpha_{it} (1 + \Psi(\alpha_{it}, 1) + \alpha_{it} \Psi(\alpha_{it}, 2) - \Psi(\sum \alpha_{jt}, 1) - \alpha_{it} \Psi(\sum \alpha_{jt}, 2))$ $\mathcal{I}_{t,ij} = \alpha_{it} \alpha_{jt} \Psi(\sum \alpha_{jt}, 2)$
Poisson $\frac{e^{-\mu} \mu^y}{y!}$	$\ln(\mu_t)$	$y_t - \mu_t$	$\mathcal{I} = \mu_t$
Negative Binomial $(\frac{y+r-1}{k}) p^r (1-p)^y$	$\ln(r_t)$ $\ln(p_t/(1-p_t))$	$r_t (\ln(p_t) + \Psi(y_t + r_t, 1) - \Psi(r_t, 1))$ $r_t (1-p_t) - y_t p_t$	$\mathcal{I}_{t,11} = r_t^2 (\Psi(r_t, 2) - \mathbb{E}[\Psi(r_t + y_t, 2)])$ $\mathcal{I}_{t,22} = r_t (1-p_t)$ $\mathcal{I}_{t,12} = -p_t$
Multinomial $\frac{n! \prod_{j=1}^J p_j^{y_j}}{y_1! \cdots y_J!}$ $y_J = n - \sum_{j < J} y_j$ $p_J = 1 - \sum_{j < J} p_j$	$\ln\left(\frac{p_{it}}{1 - \sum_{j=1}^{J-1} p_{jt}}\right)$ $j = 1, \dots, J-1$	$y_{it} - np_{it}$	$\mathcal{I}_{t,ii} = np_{it}(1-p_{it})$ $\mathcal{I}_{t,ij} = -np_{it}p_{jt}$

We have defined ∇_t in (4) and \mathcal{I}_t in (5). The (i, j) element of \mathcal{I}_t is denoted by $\mathcal{I}_{t,ij}$. We further note that $\Psi(x, k) = \partial^k \ln \Gamma(x) / \partial x^k$.

is applicable to a wide class of models that are characterized by a parametric likelihood specification. This turns out to be particularly relevant for the applications in Section 4, where we generalize some well-known models with time-varying parameters outside their usual area of application. For example, if the time-varying parameter is common across different observations, the specification in (3) gives an automatic and model consistent way to weight the information provided by different observations.

A possible difficulty with using the inverse information matrix as the scaling matrix S_{t-1} is that it may be difficult to compute for specific models. As an alternative one could use a unit scaling matrix $S_{t-1} \equiv I$. In this case, the updating mechanism only uses the unscaled gradient making it close to a steepest-descent optimization step of the likelihood at time t . Our

experience, however, is that this type of updating mechanism is often less stable. Therefore, we prefer to stick to the information scaled score step whenever possible.

A further useful feature of the GAS model in (3) is that under the correct model specification s_t is a martingale difference, $E_{t-1}[s_t] = 0$. This follows directly from the properties of the score vector. Due to scaling by the matrix S_{t-1} , we also obtain $E_{t-1}[s_t s_t'] = S_{t-1} \cdot \mathcal{I} \cdot S_{t-1}'$. This simplifies to \mathcal{I}_{t-1}^1 and \mathcal{I}_{t-1} for $S_{t-1} = \mathcal{I}_{t-1}^{-1}$ and $S_{t-1} = I$, respectively. The first two moments of the driving mechanism s_t thus easily link up with the theoretical properties of the postulated model density (1).

An important ingredient of the GAS model is the scaling of the score. As mentioned in Example 3 above, in some cases scaling by the inverse information matrix may cause numerical instability. The simplest example is given by a time-varying AR(1) model without intercept,

$$y_t = \phi_{t-1} y_{t-1} + \varepsilon_t,$$

with ε_t standard normally distributed. The information scaled score step in this case reduces to $s_t = (y_t - \phi_{t-1} y_{t-1})/y_{t-1}^2$. The GAS updating scheme becomes numerically unstable if the process y_{t-1} is close to zero. In this case, the information matrix $\mathcal{I} = y_{t-1}^2$ is close to zero and s_t can jump to extreme values. To accommodate this numerical problem, we introduce a form of information smoothing over the most recent stretch of observations, that is $S_{t-1} = (\mathcal{I}_{t-1}^c)^{-1}$ where

$$\mathcal{I}_{t-1}^c = \alpha \mathcal{I}_{t-2}^c + (1 - \alpha) \mathcal{I}_{t-1}. \quad (14)$$

for some $0 \leq \alpha \leq 1$. This is an EWMA smoothing scheme. Other weighting schemes for smoothing the information are also possible. The smoothing parameter α determines the number of observations that S_{t-1} takes into account. For $\alpha \rightarrow 0$, we recover the standard GAS model with information scaling. For $\alpha \rightarrow 1$, the model tends to average the information over all past observations. The optimal smoothing parameter could be fixed a priori, or be determined from the data itself by treating α as part of the static parameter vector θ in the likelihood.

The basic dynamics of (2) may be further extended in obvious directions. For example, it may be interesting to include exogenous variables in (2), or to generalize the evolution of f_t to include other non-linearities such as regime-switching. In addition, it may be useful in some applications to consider long-memory versions of (2), for example

$$f_t = \omega + \sum_{i=0}^{\infty} \frac{(i+d-1)!}{i!(d-1)!} s_{t-1-i},$$

for fractional integration parameter $d < 1/2$, such that we obtain a fractionally integrated GAS or FIGAS model specification, similar to the ARFIMA and FIGARCH literature, see the seminal paper of Hosking (1981). We leave such extensions for future research.

3 Statistical properties

3.1 Estimation and inference

An important advantage of observation driven models is that parameter estimation is relatively straightforward using maximum likelihood (ML). This also holds for the GAS model. For an observed time series y_1, \dots, y_n and adopting the standard prediction error decomposition, we can express the maximization problem as

$$\max_{\theta} \sum_{t=1}^n \ell(\theta; y_t, f_t, Y_1^{t-1}, X_1^t, F_1^{t-1}), \quad (15)$$

where $\ell(\theta; y_t, f_t, Y_1^{t-1}, X_1^t, F_1^{t-1}) = p(y_t|f_t, Y_1^{t-1}, X_1^t, F_1^{t-1}; \theta)$ for an observed value y_t . Similar to the GARCH model, the GAS model defines the filter for the time-varying parameters. This makes likelihood evaluation particularly simple. It only requires the implementation of the GAS updating function (2) and the evaluation of $p(y_t|f_t, Y_1^{t-1}, X_1^t, F_1^{t-1}; \theta^*)$ for a particular value θ^* of θ .

It is possible to formulate recursions for computing the gradient of the likelihood with respect to the static parameters θ . Gradient recursions for the GARCH model have been developed by Fiorentini, Calzolari, and Panattoni (1996). For the GAS(1,1) specification, we obtain

$$\frac{d\ell_t}{d\theta'} = \frac{\partial \ln p_t}{\partial \theta'} + \frac{\partial \ln p_t}{\partial f'_{t-1}} \frac{\partial f_{t-1}}{\partial \theta'}, \quad (16)$$

$$\frac{\partial f_t}{\partial \theta'} = \frac{\partial \omega}{\partial \theta'} + A_0 \frac{\partial s_t}{\partial \theta'} + B_1 \frac{\partial f_{t-1}}{\partial \theta'} + (s'_t \otimes I) \frac{\partial \vec{A}_0}{\partial \theta'} + (f'_{t-1} \otimes I) \frac{\partial \vec{B}_1}{\partial \theta'}, \quad (17)$$

$$\frac{\partial s_t}{\partial \theta'} = S_{t-1} \frac{\partial \nabla_t}{\partial \theta'} + (\nabla'_t \otimes I) \frac{\partial \vec{S}_{t-1}}{\partial \theta'}, \quad (18)$$

where $\ell_t = \ell(\theta; y_t, f_t, Y_1^{t-1}, X_1^t, F_1^{t-1})$, $p_t = p(y_t|f_t, Y_1^{t-1}, X_1^t, F_1^{t-1}; \theta)$, $\vec{A} = \text{vec}(A)$ denotes the vector with the stacked columns of the matrix A and \otimes is the Kronecker matrix product. Higher order GAS specifications can be dealt with similarly by formulating the GAS model updating equation in companion form. The log-likelihood derivatives can thus be computed simultaneously with the time-varying parameters f_t . However, computing the analytic derivatives, in particular for (18), may be cumbersome. In practice, we therefore often turn to likelihood maximization based on numerical derivatives.

The easiest way to conduct inference for GAS models is to apply a standard limiting result and use the inverse information matrix at the optimum to compute standard errors and t -values for the estimated parameters. In particular, if θ gathers the static parameters of the model, we conjecture that under standard regularity conditions, the maximum likelihood estimator $\hat{\theta}$ of θ is consistent and satisfies

$$T^{1/2}(\hat{\theta} - \theta) \rightarrow N(0, H^{-1}),$$

with $H = E[\partial^2 \ell / \partial \theta \partial \theta']$.

It is not clear, however, that standard statistical results apply directly. As an example, even though $\{s_t\}$ forms a martingale difference sequence, it is not directly evident that the GAS(1,1) model will be stable as long as $|B| < 1$. This follows because the volatility of s_t also changes over time in a stochastic way. For example, the GAS specification per se does not exclude the variance of s_t from becoming unbounded if s_t becomes large or small. If the model density is such that the inverse information matrix with respect to f_{t-1} is uniformly bounded, standard stability results apply for $|B| < 1$. This holds for a number of examples we discuss in Section 4. It is clear that given the generality of the GAS specification, the conditions for standard asymptotic theory to apply need to be checked on a case by case basis. We leave this to future research and concentrate here on the conceptual issues. To provide some indications of the statistical convergence, we complement our empirical examples of Section 4 by a Monte-Carlo simulation experiment for a selected set of examples in Section 5.

3.2 Parameterization and identification issues

The GAS specification also allows freedom of choice regarding the model's parameterization. In the GARCH example from Section 2, the time-varying parameter was $f_t = \sigma_t^2$. If one would like to account for the positivity of σ_t^2 , an obvious alternative would be to parameterize the model in terms of $f_t = \ln(\sigma_t^2)$. After some manipulations, the GAS(1,1) specification for this 'new' model is

$$f_t = \omega + A_1 \cdot \left(\frac{y_t^2}{\sigma_{t-1}^2} - 1 \right) + B_1 f_{t-1}. \quad (19)$$

The GAS dynamics automatically adapt to the choice of parameterization. In general, assume that one prefers a different parameterization $\tilde{f}_t = h(f_t)$ for some invertable mapping $h(\cdot)$. Let $\dot{h}_t = \partial h(f_t) / \partial f_t'$. Note that \dot{h}_t is deterministic given all information up to and including time t . Let s_t^f denote the GAS information scaled score step for parameterization f_t . For well behaved densities, the information matrix equals both the expected outer product of scores and the expected second derivative of the log density. This allows the information scaled score step to be written as

$$s_t^{\tilde{f}} = \left(E_{t-1}[\dot{h}_{t-1} \nabla_t \nabla_t' \dot{h}_{t-1}'] \right)^{-1} \dot{h}_{t-1} \nabla_t = \dot{h}_{t-1}' s_t^f. \quad (20)$$

Reparameterizing the model thus reduces to rescaling the information scaled score step by the gradient of the mapping $h(\cdot)$ in each period.

Another important issue concerns parameter identification. Consider a model density of the following form,

$$p(y_t; \Phi f_{t-1}), \quad (21)$$

where f_t follows a GAS(1,1) specification and Φ is a matrix of constants. For example, Φf_t can be a vector of volatilities of a vector time series, while f_t is the common factor driving all of them. Then it is not possible to define both Φ and all GAS parameters ω , A , and B , simultaneously. Take the model in (21) and introduce an invertible matrix K . Define $\tilde{f}_t = K f_t$, $\tilde{s}_t = K s_t$, $\tilde{\Phi} = \Phi K^{-1}$, $\tilde{\omega} = K \omega$, $\tilde{A} = K A K^{-1}$, $\tilde{B} = K B K^{-1}$. The likelihoods for $p(y_t; \Phi f_{t-1})$ and $p(y_t; \tilde{\Phi} \tilde{f}_{t-1})$ are obviously identical. Pre-multiplying the GAS(1,1) transition equation for the original parameterization by K , we obtain

$$K f_t = K \omega + K B K^{-1} K s_t + K A K^{-1} K f_{t-1} \Leftrightarrow \quad (22)$$

$$\tilde{f}_t = \tilde{\omega} + \tilde{A} \tilde{f}_{t-1} + \tilde{B} \tilde{s}_{t-1}. \quad (23)$$

Introducing the notation $h_t = \Phi f_t = \tilde{\Phi} \tilde{f}_t$, however, we note that for the parameterization $p(y_t; \tilde{\Phi} \tilde{f}_{t-1})$ the GAS driver equals

$$\begin{aligned} \text{E}[\nabla_{\tilde{f}_{t-1}} \nabla'_{\tilde{f}_{t-1}}]^{-1} \nabla_{\tilde{f}_{t-1}} &= \text{E}[\tilde{\Phi}' \nabla_{h_{t-1}} \nabla'_{h_{t-1}} \tilde{\Phi}]^{-1} \tilde{\Phi}' \nabla_{h_{t-1}} \\ &= K \text{E}[\Phi' \nabla_{h_{t-1}} \nabla'_{h_{t-1}} \Phi]^{-1} \Phi' \nabla_{h_{t-1}} \\ &= \tilde{s}_t. \end{aligned}$$

The GAS equation for the model $p(y_t; \tilde{\Phi} \tilde{f}_{t-1})$ is thus also given by (23). To solve this identification issue, restrictions must be imposed on the matrix Φ . For example, one can equate specific rows of Φ to corresponding rows from the identity matrix. Note that the issue cannot be solved by only imposing restrictions on the matrix A in the GAS equation, for example $A = I$. To see this, put $A = I$, but let K be an arbitrary invertible matrix. The normalization constraint then also holds for the re-parameterized model as $\tilde{A} = K A K^{-1} = K K^{-1} = I$. The other parameter values, however, are clearly different (e.g., Φ versus $\tilde{\Phi} = \Phi K^{-1}$), while the likelihood value remains unchanged. The fact that the identification issue can be solved via restrictions on Φ but not on A in this case is a direct effect of the equivariance of the score and information matrix as a basis for the recursions in the GAS model. Care is therefore needed in normalizing the parameter spaces of models with a factor structure, such as (21).

4 Applications and new models

In this section we apply the GAS model to a number of empirical examples. The examples provide more insight into the potential of the GAS specification for various models. We also provide several new models and non-trivial extensions of existing models.

4.1 Time-varying nonlinear regression models

The term structure of interest rates plays a central role in both macroeconomics and finance as it describes the linkage between monetary policymakers' impact on the short term interest rate and firms' investment decisions at longer horizons. We develop a parameter driven analogue for the popular term structure model of Nelson and Siegel (1987) given by the partial nonlinear regression model

$$y_{t,\tau} = x_\tau(\lambda)\beta + \varepsilon_{t,\tau}$$

where $y_{t,\tau}$ is the interest rate at time t for an investment that matures after τ months. The 1×3 covariate vector $x_\tau(\lambda)$ is defined as

$$x_\tau(\lambda) = [1, (\lambda\tau)^{-1}(1 - \exp(-\lambda\tau)), (\lambda\tau)^{-1}(1 - \exp(-\lambda\tau)) - \exp(-\lambda\tau)],$$

where the coefficients λ and β are unknown and with independent disturbance $\varepsilon_{t,\tau} \sim N(0, \sigma^2)$ for a given time t . For the particular choice of $x_\tau(\lambda)$, the three coefficients in β can be interpreted as the level, slope, and curvature of the term structure, respectively. The slope and curvature factors depend on the parameter λ that is nonlinear in $y_{t,\tau}$. At time t , interest rates for m maturities can be observed such that $\tau = \tau_j$ for $j = 1, \dots, m$. Based on these m observations (for a given t), parameters β , λ and σ^2 can be estimated by nonlinear least squares methods. This estimation procedure can be repeated at each time t , resulting in a time series of parameter estimates, see Diebold and Li (2006) for further details.

Recently, Diebold, Rudebusch, and Aruoba (2006) proposed analyzing the time series dimension simultaneously with the maturity dimension by considering the Nelson-Siegel model as a multivariate state space model. For this purpose, they treat β as a 3×1 vector of unobservables, β_t , that are modeled as a vector autoregressive (VAR) process. Furthermore, the interest rates for all m maturities at time t are put into the observation vector $y_t = (y_{t,\tau_1}, \dots, y_{t,\tau_m})'$ and modeled by $y_t = X_t \beta_t + \varepsilon_t$ where $X_t = [x_{\tau_1}(\lambda)', \dots, x_{\tau_m}(\lambda)']'$, with serially independent $m \times 1$ disturbance vector $\varepsilon_t \sim N(0, D)$ and $m \times m$ positive diagonal matrix D . The unknown parameters in the VAR model for β_t , as well as D and λ are then estimated by ML using the Kalman filter in this parameter driven approach. In an observation driven approach, we set $f_{t-1} = \beta_t$ and consider the Gaussian density for $y_t - X_t f_{t-1} \sim N(0, D)$ with the GAS(1,1) updating equation $f_t = \omega + A_0 s_t + B_1 f_{t-1}$ where

$$s_t = (X_t' D^{-1} X_t)^{-1} X_t' D^{-1} (y_t - X_t f_{t-1}),$$

assuming that $m > 3$. The static parameter vector is given by

$$\theta = \left(\text{diag}(D)', \omega', \vec{A}_0', \vec{B}_1', \lambda \right)'.$$

To illustrate our GAS approach to the dynamic Nelson-Siegel model, we consider the dataset from Diebold et al. (2006) which consists of monthly U.S. Treasury yields with maturities 3, 6, 9, 12, 15, 18, 21, 24, 30, 36, 48, 60, 72, 84, 96, 108, and 120 months over the period from January 1972 to December 2000. For comparison purposes, we have estimated both the parameter driven (DNS) and observation driven (GAS) models by ML. The DNS estimates are close to those reported by Diebold et al. (2006) while the GAS estimates are different than those obtained for the parameter driven DNS model. For example, the estimate of λ is 0.0778 in the DNS model while it is 0.0948 in the GAS model. The estimates of other coefficients are also different which emphasizes that the interpretation of “comparable” coefficients in both models are different. Nevertheless, the estimates of the three factors in β_t are similar as we observe from the graphs in Figure 1. The three estimated factors in both models accurately track the changes of the yield curve over time.




Figure 1: *Three factor dynamic Nelson-Siegel model: level, slope and curvature factors. Each panel contains the estimated factor from the GAS(1,1) observation driven model, the estimated (filtered) factor from the DNS parameter driven model, and the data proxy which for the level factor is the 120 month yield, for the slope factor is the spread from the 3 month yield minus the 120 month yield and for the curvature factor is the 24 month yield minus the 3 and 120 month yield.*

4.2 Pooled marked point-process models

Models with time-varying intensities have been given much attention in the finance and microeconomic literature. The principal areas of application in economics include intraday trade data (market microstructure), defaults of firms, credit rating transitions and (un)employment spells over time. To illustrate the GAS model in this setting, we consider an application from the credit risk literature in which pooled marked point-processes are playing an important role. We develop a new and useful modeling framework that is based on the GAS specification.

Recently, a number of promising models with stochastically evolving intensities have been proposed, see Bauwens and Hautsch (2006), Koopman, Lucas, and Monteiro (2008), Duffie, Eckner, Horel, and Saita (2006), and Koopman, Lucas, and Schwaab (2008). The econometric handling of these parameter driven models is intricate while parameter estimation can be computationally demanding. In particular, likelihood evaluation for these models requires computing a (very) high-dimensional integral. Importance sampling techniques or Markov chain Monte Carlo algorithms have been adopted for likelihood-based analyses of these models. The use of such simulation-based techniques, however, obstructs the widespread application of these models in practice. A computationally less-demanding alternative can be based on developing observation driven analogues of these models.

The first step would then be to consider multivariate generalizations of Russell (2001). However, this is not straightforward. Most of the models of Russell (2001) are developed in the context of high frequency data and in particular for stock trades. The structure of data sets of trades is substantially different from the data sets that are used in credit risk. Whereas in high frequency data one typically observes many spells for a limited number of stocks, in modeling credit data one typically works with many different companies that only have very few spells each. This requires the pooling of data over different companies in the sample. Consequently, different events might carry information on where the dynamic parameter is heading at any point in time. The GAS model provides an easy and consistent methodology to address this issue.

Let $y_k(t)' = (y_{1k}(t), \dots, y_{nk}(t))$ be a vector of marks of n competing risk processes for firm $k = 1, \dots, N$. We have $y_{jk}(t) = 1$ if event type j materializes for firm k at time t , and zero otherwise. By following the application in Koopman, Lucas, and Monteiro (2008), we model the log intensities of these processes by

$$\lambda_{jk}(t) = \eta_j + \alpha_j' f_{t-1}, \quad (24)$$

where η_j is the baseline intensity and α_j is the vector of loadings for f_t which is the vector of dynamic factors and specified by the GAS(1,1) updating equation (2) with $\omega = 0$. Since f_t is an

unobserved process, we may impose sign restrictions for α_j to obtain economic interpretations for the factors. This GAS specification states that the intensities of all firms are driven by the same vector of time-varying systematic factors f_t . Model (24) nests the model of Russell (2001) when we set the dimension of f_t equal to the number of firms N . In a credit risk context, we typically have $\dim(f_t) \ll N$. Further parameter restrictions are needed for model identification, see the discussion in Section 3. In the illustration below we set one of the α_j 's equal to unity.

The log-likelihood specification using (24) is

$$\ell_t = \sum_{j,k} y_{jk}(t) \lambda_{jk}(t) - R_{jk}(t) \cdot (t - t^*) \cdot \exp(\lambda_{jk}(t)), \quad (25)$$

where t^* is the last event time before t , and $R_{jk}(t)$ is a zero-one variable indicating whether company k is potentially subject to risk j at time t . Based on the first and second derivative of ℓ_t , we obtain

$$s_{t+1} = \left[\sum_{j,k} w_{jk}(t) \alpha_j \alpha_j' \right]^{-1} \left(\sum_{j,k} y_{jk}(t) \alpha_j - R_{jk}(t) \cdot (t - t^*) \cdot \exp(\lambda_{jk}(t)) \alpha_j \right), \quad (26)$$

where $w_{jk}(t) = R_{jk}(t) \cdot \exp(\lambda_{jk}(t)) / \sum_{j,k} R_{jk}(t) \cdot \exp(\lambda_{jk}(t)) = \text{P}[y_{jk}(t) = 1]$ is the probability of the next event being of type j for company k . Combining all these elements into a GAS specification, we have obtained a new observation driven model for credit rating transitions.

As an illustration, we adopt the model described above for the CreditPro 7.0 data set which contains the Standard and Poor's (S&P) rating histories of all US corporates over the period 1981–2005. We distinguish two complementary credit rating classes: the investment grade (IG) and the sub-investment grade (SIG). The GAS(1,1) model has a univariate (single) factor f_t and the updating equation has the scaled score function (26). The resulting model is estimated under the restrictions $\alpha_3 < 0$ and $\alpha_4 = 1$. The estimation results are presented in Table 2. The GAS parameter A_0 is estimated close to unity which implies a persistent dynamic process for f_t . Given the estimates of α_j , the downgrades appear to be most sensitive to the common factor f_t . In particular, the baseline downgrade from investment grade to default is small with an estimate of -7.4 while it is strongly sensitive to the common factor f_t with a loading estimate of 1.2. Interestingly, the estimated pattern of the systematic intensity factor f_t is close to the estimated pattern of the parameter driven model of Koopman, Lucas, and Monteiro (2008). However, in a GAS framework we do not require computationally intensive methods such as importance sampling for parameter and factor estimation.

It is straightforward in our GAS framework to generalize the model to a three-factor model. In this case, A_0 and B_1 in the GAS updating equation become 3×3 matrices. To obtain

Table 2: Estimation results for the parameters in the one-factor GAS(1,1) intensity model (24) in a two-grade system, with $\alpha_4 = 1$, with the scaled scoring function (26) and based on the S&P ratings of all US corporates between 1981 and 2005. The estimates are reported with asymptotic standard errors in parantheses below the estimates. Parameter A_0 is subject to a logistic transformation during estimation and α_3 is subject to a identifying restriction $\alpha_3 < 0$.

	IG \rightarrow SIG	IG \rightarrow DEF	SIG \rightarrow IG	SIG \rightarrow DEF
j	1	2	3	4
η	-3.920 (0.118)	-7.360 (0.353)	-3.360 (0.109)	-3.330 (0.217)
α	0.520 (0.076)	1.190 (0.330)	-0.470 (0.086)	1.000 —
	logit(A_0)	A_0	B_1	
	6.415 (0.537)	0.998	0.024 (0.003)	




Figure 2: Marked point-process model: the estimated intensities in a two-grade system of the GAS(1,1) model with three credit risk factors, based on the scaled score function (26) and using the S&P rating histories of US corporates for the period 1981–2005.

interpretable dynamic factors and to acquire the necessary restrictions, we set the loading vector α_j equal to the j th column of a 4×4 diagonal matrix for $j = 1, 2, 3$ while $\alpha_4 = \alpha_3$. This parsimonious specification implies that upgrades and downgrades between IG and SIG have different factors while transitions to default also have a distinct factor. The static parameter vector θ contains the elements of A_0 and B_0 together with the baseline intensities η_j and the three unknown coefficients $\alpha_1, \dots, \alpha_3$. These parameters can be estimated in the usual way by ML. After parameter estimation, we obtain similar estimated patterns for the three factors in f_t as for the more involved parameter driven model of Koopman, Lucas, and Monteiro (2008). In particular, we corroborate the finding that the dynamics of upgrades are substantially different from those of downgrades and defaults as can be clearly viewed in Figure 2 where the estimated intensities $\eta_j + \alpha'_j f_{t-1}$ are displayed.

4.3 Unobserved component models with a single source of error

Unobserved components or structural time series models are a popular class of parameter driven models where the unobserved components (UC) have a direct interpretation, see Harvey (1989). In this section, we describe observation-driven analogues to UC models. For a univariate time series y_1, \dots, y_n , a univariate signal ψ_t can be extracted. The dynamic properties of ψ_t can be broken into a vector of factors f_{t-1} that are specified by the updating equation (2). For example, we can specify the signal as the sum of r factors, that is

$$\psi_t = f_{1,t-1} + \dots + f_{r,t-1} \quad (27)$$

with $f_t = (f_{1,t}, \dots, f_{r,t})'$. In the case $r = 2$, we can specify the first factor as a time-varying trend component (random walk plus drift) and the second factor as a second-order autoregressive process with possibly cyclical dynamics. For this decomposition we obtain the GAS(1,2) model with observation model $y_t = \psi_t + \varepsilon_t = f_{1,t-1} + f_{2,t-1} + \varepsilon_t$, observation density $p(y_t | \psi_t; \theta) = N(f_{1,t-1} + f_{2,t-1}, \sigma^2)$ and updating equation

$$f_t = \begin{pmatrix} \omega \\ 0 \end{pmatrix} + \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} s_t + \begin{bmatrix} 1 & 0 \\ 0 & \phi_1 \end{bmatrix} f_{t-1} + \begin{bmatrix} 0 & 0 \\ 0 & \phi_2 \end{bmatrix} f_{t-2}. \quad (28)$$

The constant ω is the drift of the random walk trend factor $f_{1,t}$ and the autoregressive coefficients ϕ_1 and ϕ_2 impose stationarity for the second factor $f_{2,t}$. The scaled score function is given by

$$s_t = y_t - \psi_t = y_t - f_{1,t-1} - f_{2,t-1} = \varepsilon_t, \quad (29)$$

and can be interpreted as the single source of error. The static parameter vector θ , consisting of coefficients $\omega, a_1, a_2, \phi_1, \phi_2$ and σ , can be estimated straightforwardly by ML. The estimates

Table 3: *Estimation results for the parameters in the trend-cycle GAS(1,2) decomposition model (27) with the updating equation (28) and the scaled scoring function (29) based on quarterly log U.S. real GDP from 1947(1) to 2008(2). The estimates are obtained by ML and reported with asymptotic standard errors in parantheses below the estimates. Furthermore, the ML estimates of parameters in the parameter driven trend-cycle UC model (30)–(31) are reported which are based on the same data set.*

	ω	a_1	a_2	ϕ_1	ϕ_2	σ	log-like
GAS	0.825 (0.043)	0.723 (0.206)	0.563 (0.202)	1.328 (0.130)	-0.424 (0.142)	0.905 (0.041)	-324.51
UC	0.825 (0.040)	0.604 (0.098)	0.621 (0.112)	1.501 (0.102)	-0.573 (0.106)	– –	-324.06

of f_t result in a decomposition of y_t into trend, cycle, and noise. This GAS decomposition can be regarded as the observation driven equivalent of the UC models of Watson (1986) and Clark (1989), who also aim to decompose macroeconomic time series into trend and cycle factors. The UC trend-cycle decomposition model is then given by $y_t = f_{1,t} + f_{2,t}$ with

$$f_{1,t} = \omega + f_{1,t-1} + a_1 \xi_{1,t}, \quad \xi_{1,t} \sim N(0, 1), \quad (30)$$

$$f_{2,t} = \phi_1 f_{2,t-1} + \phi_2 f_{2,t-2} + a_2 \xi_{2,t}, \quad \xi_{2,t} \sim N(0, 1), \quad (31)$$

where the disturbances $\xi_{1,t}$ and $\xi_{2,t}$ are mutually and serially independent.

To illustrate the GAS trend-cycle decomposition model, we consider the time series of quarterly log U.S. real GDP from 1947(1) to 2008(2) obtained from the Federal Reserve Bank of St. Louis. The vector of static coefficients θ is estimated by ML and the results are reported in Table 3. The estimated autoregressive polynomial for factor $f_{2,t}$ has roots in the complex range and therefore factor $f_{2,t}$ has cyclical properties. We may interpret $f_{2,t-1}$ as a real-time business cycle indicator for time t which is displayed in Figure 3. To compare this indicator with the indicator produced by the Watson (1986) model, we also report the ML estimates of the corresponding coefficients in an UC trend-cycle model. These estimates are obtained by using the Kalman filter for likelihood evaluation. Parameter estimates for the UC model are reported in Table 3 and the one-step ahead predicted estimate of $f_{2,t}$ is plotted in Figure 3. We find that the parameter estimates from each model correspond closely. The second factor from each model exhibits cyclical behavior and the growth rate of the trend is estimated to be the same. Estimates of the GAS and UC cycle factors in Figure 3 are almost indistinguishable.

The GAS framework is sufficiently general to provide an observation driven alternative for the decomposition of univariate and multivariate time series based on UC models including models with trend, seasonal, cycle and irregular components. For example, the GAS updating equation can also be designed to incorporate the trend and cycle dynamics as formulated by




Figure 3: *Trend-cycle example: estimated cycles from the GAS and UC trend-cycle models based on quarterly log of U.S. real gdp from 1947(1) through 2008(1). NBER recession dates are indicated by the shaded regions.*

Harvey and Jaeger (1993). Regression and intervention effects can also be incorporated in the GAS specification, see the discussion in Subsection 2.1. Since the resulting GAS models are equivalent to single source of error models, we refer to Ord et al. (1997) for a more detailed discussion on this class of models.

4.4 State space models with time-varying GAS parameters

In this subsection, we argue that the GAS framework can also be applied to the static parameters of linear Gaussian state space models. Similar ideas have been considered for combining state space models with ARCH disturbances; see, e.g. Harvey, Ruiz, and Sentana (1992). To illustrate its relevance, we consider the local level model as specified by

$$y_t = \mu_t + \varepsilon_t, \quad \mu_t = \mu_{t-1} + \xi_t, \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2), \quad \xi_t \sim N(0, \sigma_\xi^2), \quad t = 1, \dots, n. \quad (32)$$

We treat the variances of the irregular and level disturbances, σ_ε^2 and σ_ξ^2 respectively, as GAS factors in order to obtain a time-varying UC model. The resulting model is similar in spirit to the model considered by Stock and Watson (2007) for forecasting quarterly U.S. inflation. The details of our time-varying local level model are given below.

The variances σ_ε^2 and σ_ξ^2 of the local level model (32) can be replaced by two GAS factors to obtain a model with time-varying variances. Since variances must remain positive, we specify the two GAS factors as log-variances. Replacing the constant variances in (32) with the GAS

factors, new disturbances for the local level model are given by

$$\varepsilon_t \sim N\{0, \exp(f_{1,t-1})\}, \quad \xi_t \sim N\{0, \exp(f_{2,t-1})\}, \quad t = 1, \dots, n, \quad (33)$$

with $f_t = (f_{1,t}, f_{2,t})'$. Conditional on f_{t-1} , the unobserved level μ_t in (32) remains linear in the observations y_1, \dots, y_t and therefore the Kalman filter can be adopted to produce the optimal estimate of μ_t which is given by $a_t = E(\mu_t | f_{t-1}, Y_1^{t-1}, F_1^{t-2}; \theta)$ with its mean square error p_t . The log-likelihood function for the local level model (32) is given by

$$\ell(\theta) = \sum_{t=1}^n \ell_t(\theta; y_t, Y_1^{t-1}), \quad \ell_t = \ell_t(\theta; y_t, Y_1^{t-1}) = -\frac{1}{2} \ln 2\pi - \frac{1}{2} \ln d_t - \frac{1}{2} e_t^2 / d_t, \quad (34)$$

where the prediction error e_t and its variance d_t are evaluated by the Kalman filter as given by

$$e_t = y_t - a_t, \quad d_t = p_t + \exp(f_{1,t-1}), \quad k_t = p_t / d_t, \quad (35)$$

$$a_{t+1} = a_t + k_t v_t, \quad p_{t+1} = (1 - k_t) p_t + \exp(f_{2,t-1}), \quad (36)$$

for $t = 1, \dots, n$ with diffuse initializations $a_1 = 0$ and $p_1 = \kappa$ while $\kappa \rightarrow \infty$, see Durbin and Koopman (2001). The Kalman filter update can be carried out simultaneously with the GAS updating equation (2) for f_t and with the scaled score function defined by (3). In this case, we have

$$\nabla_{i,t} = \frac{\partial \ell_t}{\partial \exp(f_{i,t-1})} \times \frac{\partial \exp(f_{i,t-1})}{\partial f_{i,t-1}} = \exp(f_{i,t-1}) \frac{\partial \ell_t}{\partial \exp(f_{i,t-1})} \quad (37)$$

for $i = 1, 2$. Given the Kalman filter equations (35)–(36), the latter term of (37) is evaluated by

$$\frac{\partial \ell_t}{\partial \exp(f_{i,t-1})} = -\frac{1}{2} \left(\dot{d}_{it} + 2e_t \dot{e}_{it} \right) / d_t + \frac{1}{2} \dot{d}_{it} (e_t / d_t)^2, \quad (38)$$

where $\dot{e}_{it} = \partial e_t / \partial \exp(f_{i,t-1})$ and $\dot{d}_{it} = \partial d_t / \partial \exp(f_{i,t-1})$ are evaluated by the additional recursions

$$\dot{e}_{it} = -\dot{a}_{it}, \quad \dot{d}_{it} = \dot{p}_{it} + 1(i = 1), \quad \dot{k}_{it} = (\dot{p}_{it} - k_t \dot{d}_{it}) / d_t, \quad (39)$$

$$\dot{a}_{i,t+1} = (1 - k_t) \dot{a}_{it} + \dot{k}_{it} v_t, \quad \dot{p}_{i,t+1} = -\dot{k}_{it} p_t + (1 - k_t) \dot{p}_{it} + 1(i = 2), \quad (40)$$

where $1(i = j)$ equals one if $i = j$ and zero otherwise, with initializations $\dot{a}_{i1} = 0$ and $\dot{p}_{i1} = 0$, for $t = 1, \dots, n$ and $i = 1, 2$. For the local level model, we therefore obtain two additional recursions to evaluate the score and they can also be carried out simultaneously with the Kalman filter. By following Harvey (1989, p 140–2), we approximate the information matrix by

$$\mathcal{I}_{t-1}(i, j) = -E_{t-1} \left(\frac{\partial^2 \ell_t}{\partial f_{i,t-1} \partial f_{j,t-1}} \right) \approx \exp(f_{i,t-1}) \exp(f_{j,t-1}) \times \left(\frac{1}{2} \dot{d}_{it} \dot{d}_{jt} / d_t^2 + \dot{e}_{it} \dot{e}_{jt} / d_t \right), \quad (41)$$

where $\mathcal{I}_{t-1}(i, j)$ is the (i, j) element of \mathcal{I}_{t-1} for $i, j = 1, 2$. The computation of this approximation is feasible given the additional recursions (39)–(40). However, in practice, the information

matrix may become singular or close to singular. We therefore have adopted the EWMA smoothing scheme for the information matrix to obtain the scaled score function, see the discussion in Subsection 2.3.

The Kalman filter, the additional recursions for the score, and the GAS updating equation for f_t can be carried out simultaneously. The parameter vector θ consists of the GAS updating coefficients in ω , A_0, \dots, A_{p-1} , B_1, \dots, B_q and, possibly, the smoothing coefficient α of the EWMA smoothing recursion for the information matrix. The estimation of θ is done by ML via the maximization of $\ell(\theta)$ in (34) with respect to θ . Given a value of θ , the recursions in real-time provide estimates of μ_t via the Kalman filter and estimates of f_t via the score recursions and the GAS updating equation simultaneously.

To illustrate the new GAS model, we consider a time series of quarterly U.S. Consumer Price Index inflation from 1959(1) to 2007(2) obtained from the FRED database. The local level model with a GAS(1,1) updating equation for the log-variances is adopted and the methodology of estimation as discussed above is implemented. The EWMA smoothing scheme is applied to the information matrix depends on α which is estimated as part of θ . The GAS coefficient matrices A_0 and B_1 are chosen to be diagonal so that we need to estimate a total of seven coefficients. The estimation results are given by

$$\hat{\omega} = \begin{pmatrix} 0.122 \\ 0.003 \end{pmatrix}, \quad \hat{A}_0 = \begin{bmatrix} 0.426 & 0 \\ 0 & 0.081 \end{bmatrix}, \quad \hat{B}_1 = \begin{bmatrix} 0.569 & 0 \\ 0 & 0.916 \end{bmatrix}, \quad \hat{\alpha} = 0.674,$$

with the maximized loglikelihood value given by -372.97 . The loglikelihood value for standard local level model is given by -394.19 indicating a substantial improvement in fit with the GAS model specification.

Panels (i) and (ii) of Figure 4 present the estimated factors $\exp(f_{1,t}/2)$ and $\exp(f_{2,t}/2)$, which are the standard deviations for ε_t and ξ_t , respectively. The standard deviation of the observation disturbance is moderate until the end of 2006 onwards at which time it has become relatively high. The standard deviation of the level disturbance increased in the 1970's during the periods of higher inflation and then decreased steadily over the remaining sample. The signal to noise ratio in our GAS framework is defined by the ratio $q_t = \exp(f_{2,t}) / \exp(f_{1,t})$. When it is low, the estimate of μ_t is based on a long range of past observations. When it is high, μ_t is estimated using only a small set of recent observations. The third graph in Figure 4 displays q_t based on the estimate of f_t . As the properties of the model suggested, the estimate of μ_t is only based on recent observations during the years of the oil-crisis, 1974–1976. From 1980 onwards, the level of inflation is more stable and a longer stretch of past observations are used in the estimator of the trend. Finally, the (filtered) estimate of μ_t is displayed in graph (iv) of Figure 4. The GAS framework captures the overall development of U.S. inflation

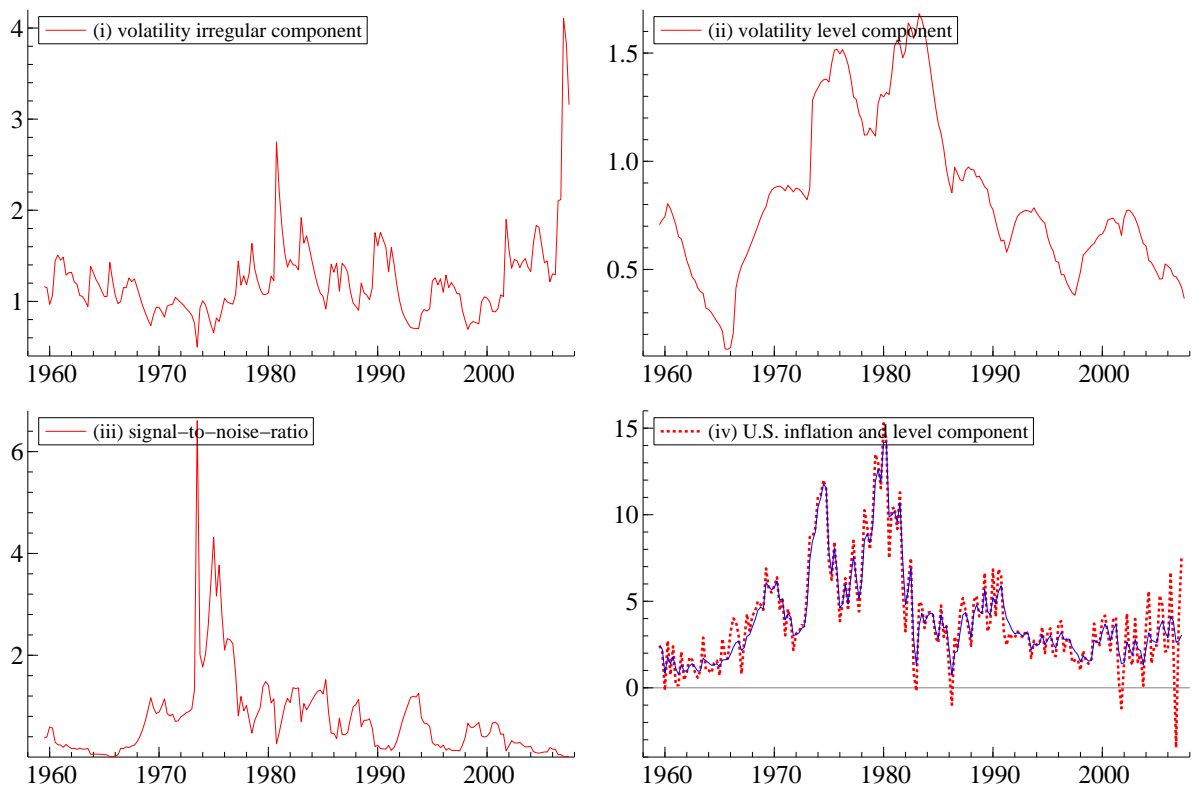


Figure 4: *Estimation results for the time-varying local level model applied to quarterly U.S. inflation from 1959(1) to 2007(2). (i) estimated volatility $\exp(f_{1,t}/2)$ of the irregular component; (ii) estimated volatility $\exp(f_{2,t}/2)$ of the trend component; (iii) estimated signal-to-noise ratio q_t ; (iv) U.S. inflation and estimated trend μ_t .*

effectively. The estimated patterns of the time-varying standard deviations are very similar to those obtained by Stock and Watson (2007) who used Markov chain Monte Carlo to estimate the time-varying variances in their parameter driven model.

4.5 Dynamic copula models

Copulas have become popular over the last decade in the literature on financial risk management. A copula is a multivariate distribution function over a hypercube with uniform marginals. The copula can be used to link marginal distributions into a multivariate distribution using Sklar's theorem. In this subsection, we demonstrate that the GAS framework provides new model specifications for simple copulas such as the bivariate Gaussian copula. We then illustrate some of the numerical extensions of the GAS specification to mixture copulas that allow for asymmetric tail behavior.

We start with simple Gaussian copulas where the GAS model suggests an alternative dynamic structure compared to earlier suggestions in the literature. Patton (2006) introduced the notion of time-varying copulas, see also Dias and Embrechts (2004) and van den Goorbergh, Genest, and Werker (2005). Patton (2006) models¹ the driving mechanism for the dynamic bivariate Gaussian copula as

$$f_t = \omega + A \sum_{i=0}^{m-1} \Phi^{-1}(u_{1,t-i})\Phi^{-1}(u_{2,t-i}) + Bf_{t-1}, \quad (42)$$

where Φ^{-1} is the inverse normal distribution function, u_{1t} and u_{2t} are the probability integral transforms using the univariate marginals, m is a smoothing parameter, and the (Gaussian) correlation parameter ρ_t is a transform of the factor f_t , e.g., $\rho_t = (1 - \exp(-f_t))/(1 + \exp(-f_t))$. Equation (42) is intuitively appealing and builds on our understanding of covariances: if the transformed marginals covary positively, the correlation should increase. The reverse holds if the transformed marginals are of opposite sign.

By using the density of the Gaussian copula, we can derive the GAS specification for the time-varying correlation parameter. Note that the score with respect to the correlation parameter is the same for the Gaussian copula and the bivariate normal. Our results thus also directly apply to the Dynamic Conditional Correlation (DCC) framework of (Engle 2002a).

Define $x_t = \Phi^{-1}(u_{1t})^2 + \Phi^{-1}(u_{2t})^2$ and $y_t = \Phi^{-1}(u_{1t})\Phi^{-1}(u_{2t})$. For $m = 1$, Patton's model (42) then reduces to

$$f_t = \omega + A \cdot y_t + B \cdot f_{t-1}. \quad (43)$$

Deriving the score and information matrix of the bivariate normal for the transformed correlation parameter, one can easily show that the GAS(1,1) specification for f_t equals

$$f_t = \omega + \frac{A}{(1 - \rho_{t-1})(1 + \rho_{t-1})} \cdot \left(y_t - \rho_{t-1} - \rho_{t-1} \frac{x_t - 2}{1 + \rho_{t-1}^2} \right) + Bf_{t-1}. \quad (44)$$

The similarities and differences between (43) and (44) are clear. Both models are driven by y_t as positively clustered transformed marginals should increase the correlation parameter. The additional scaling factor $(1 - \rho_{t-1})(1 + \rho_{t-1})$ in (44) is a consequence of modeling the transformed correlation parameter f_t rather than ρ_t directly. The additional ρ_{t-1} term in parenthesis is just a translation parameter such that $y_t - \rho_{t-1}$ becomes a martingale difference. The most interesting difference between the two models lies in the final term in parenthesis involving x_t . Clearly, $x_t - 2$ is a martingale difference. Note that x_t is large if we observe an extreme observation in u_{1t} , u_{2t} , or particularly in both. The effect of this depends on the current estimate of the

¹We adapt Patton's notation here slightly to correspond with the timing convention used in the current paper, i.e., using f_{t-1} in the copula at time t rather than f_t .

correlation parameter ρ_{t-1} . If the correlation is positive, then x_{t-1} has a negative effect. In this case the x_t term offsets part of the effect of y_t if the latter is positive, i.e., if y_t corresponds with the current positive estimate of ρ_{t-1} . If y_t is negative, however, the x_t term reinforces the magnitude of the GAS step triggered by y_t .

The effects are most clearly visualized by plotting the GAS step for different values of (u_{1t}, u_{2t}) and three different values of the correlation parameter, $\rho_{t-1} = -0.5, 0.2, 0.9$. This is done in Figure 5. If we consider the plot for $\rho = 0.9$, we see two clear differences. First, the GAS step results in a smaller increase in the correlation parameter along the $u_{1t} = u_{2t}$ axis. Particularly if u_{it} are both large or small, the step based on y_t alone results in a more pronounced increase of the transformed correlation parameter. The same holds for the smaller positive correlation parameter of $\rho = 0.2$.




Figure 5: Comparisons between the GAS and Patton drivers for the bivariate Gaussian copula as a function of the uniforms (u_{1t}, u_{2t}) . The top panels contain the graphs for the GAS step in (44) for $\rho_{t-1} = -0.5, 0.2, 0.9$ (left, middle, right). The lower graphs contain the (re-centered) steps $y_t - \rho_{t-1}$ of the Patton model, (43). The vertical axes have the same scale for each column of graphs.

The much more striking feature, however, is the increased sensitivity along the anti-diagonal for positive ρ . If the current estimate of ρ is positive, but one observes a combination of (u_{1t}, u_{2t})

that signals negative rather than positive dependence, the GAS specification is more sensitive to this and is more inclined to rapidly adjust the current estimate of ρ downward. Analogously, for negative values of ρ the left panels show that the effects are reversed. The GAS specification becomes more sensitive to observations along the diagonal than the specification based on y_t alone.

To illustrate this empirically, we extend the example from Patton (2006). We use the daily exchange rates between the German Mark (later Euro) and the US dollar, and between the Japanese Yen and the US dollar. The sample period is January 1986 through August 2008. We model the log returns on the exchange rates by AR-GARCH models and obtain the transformed sample of u_{1t} and u_{2t} as input for the Gaussian copula model. Apart from (43) and (44), we also estimate an ad-hoc implementation of the DCC framework of Engle (2002a). In particular, we model the correlation parameter directly using

$$\rho_t = \omega + A \cdot y_t + B \cdot \rho_{t-1}. \quad (45)$$

To enforce stationarity, we estimate the logit transform of B . The results are presented in Table 4 and Figure 6.




Figure 6: *Comparisons of the correlation parameter for the GAS, Patton, and DCC drivers in (43)–(45). The data are the marginal AR-GARCH transforms of log exchange rates for the German Mark-US dollar and Japanese Yen-US dollar (left panel) and for the German Mark-US dollar and British Pound-US dollar (right panel). The sample period is January 1986–August 2008.*

Table 4 shows that the GAS specification increases the log-likelihood value 25 to 125 points for the same number of parameters. The figures show the empirical estimates of the time-varying correlation. Judging on the value of the parameter B , the GAS specification is the most persistent one, followed by the DCC and the Patton specification. However, the increased sensitivity of the score mechanism to correlation breakdowns, reveals an opposite pattern in the

Table 4: *Parameter estimates for the GAS, Patton, and DCC drivers in (43)–(45). The data are the marginal AR-GARCH transforms of log exchange rates for the German Mark-US dollar and Japanese Yen-US dollar (left panel) and for the German Mark-US dollar and British Pound-US dollar (right panel). The sample period is January 1986–August 2008. Confidence interval in parentheses for B , otherwise standard errors in parentheses.*

	$10^3\omega$	A	$\ln(\frac{B}{1-B})$	B	log-like
German Mark (Euro)–US \$, Japanese Yen–US \$					
GAS	6.11 (2.48)	0.058 (0.009)	5.30 (0.37)	0.995 (0.990,0.998)	1218.16
Patton	-1.60 (0.85)	0.036 (0.003)	4.27 (0.10)	0.986 (0.983,0.989)	1191.51
DCC	1.03 (0.29)	0.008 (0.001)	4.65 (0.09)	0.991 (0.989,0.992)	1184.13
German Mark (Euro)–US \$, British Pound–US \$					
GAS	12.55 (3.55)	0.082 (0.008)	4.97 (0.26)	0.993 (0.988,0.996)	2218.82
Patton	-0.97 (0.84)	0.025 (0.002)	4.71 (0.11)	0.991 (0.989,0.993)	2090.42
DCC	2.64 (0.39)	0.004 (0.000)	4.84 (0.11)	0.992 (0.990,0.994)	2060.43

figures. Due to the sharpe decline at the edges as visualized in Figure 5, the GAS specification reacts much more fiercely to exchange rate returns of opposite sign if the current correlation estimate is positive. This is most clearly seen for the Mark-Pound example, but also the Mark-Yen example shows similar features at the end of 1993 and 2003. The DCC dynamics, and to a lesser extent the Patton dynamics, are much smoother in this sense. The difference between the dynamics for the different specifications may be highly relevant for risk managers, where changes in correlations and in particular correlation breakdowns are a major concern.

Stepping away from the Gaussian copula, it is also interesting to see how the GAS specification generalizes to non-Gaussian copulas. We consider here a mixture of Clayton-type copulas. Patton proposes a generally applicable driving mechanism for copula parameters,

$$f_t = \omega - m^{-1}A \sum_{i=1}^m |u_{1,t-i} - u_{2,t-i}| + Bf_{t-1}, \quad (46)$$

where f_t captures the dependence between the coordinates. The intuition for (46) is clear. If the most recent u_{1t} and u_{2t} are close together, this is a signal of strong dependence and, therefore, f_t is increased. Similarly f_t is decreased if u_{1t} and u_{2t} are far apart.

Though the driving mechanism in (46) is intuitively straightforward, two issues are less clear. First, (46) uses no information contained in the particular choice of the copula. As with the Gaussian copula, such information may be helpful in specifying the dynamics. Second, although (46) provides an easy updating scheme for the bivariate case, the extension to the multivariate case is less obvious. In particular, if one has an Archimedean copula characterized by a single dependence parameter, there are many different ways in which one could use the differences $|u_{it} - u_{jt}|$ for $i \neq j$ to update the dependence parameter. Equation (46) provides little guidance as to how these different and possibly conflicting signals should be weighed.

The Clayton copula for our example is a member of the Archimedean family. Its specification in dimension d is given by

$$C(u_1, \dots, u_d) = \left(1 - d + \sum_{i=1}^d u_i^{-\alpha} \right)^{-1/\alpha}. \quad (47)$$

The Clayton copula is characterized by the dependence parameter α . Low values of α indicate high levels of dependence. This is also captured by the tail dependence coefficient, which measures the probability of joint extreme exceedances. For the Clayton, extreme joint crashes receive positive probability, while joint extreme upward shocks have zero probability.

We specify $\alpha = f_{t-1}$ and define $S(\alpha) = \sum_{i=1}^d u_i^{-\alpha}$. The Clayton copula has pdf

$$c(u_1, \dots, u_d) = (1 - d + S(\alpha))^{-1/\alpha-d} \cdot \prod_{i=0}^{d-1} ((1 - i \cdot \alpha) u_i^{-\alpha-1}). \quad (48)$$

We obtain the score vector

$$\begin{aligned} \nabla_t = & - \sum_{i=0}^{d-1} \left(\frac{i}{1 - i \cdot \alpha} - \ln(u_i) \right) + \frac{1}{\alpha^2} \ln(1 - d + S(\alpha)) + \\ & \left(\frac{1}{\alpha} + d \right) \frac{\sum_{i=1}^d u_i^{-\alpha} \ln(u_i)}{1 - d + S(\alpha)}. \end{aligned} \quad (49)$$

The principal difficulty for some GAS-based dynamic copula models is deriving a closed-form expression for the information matrix. Even for simple copula models, this may quickly become unmanageable analytically. This certainly holds for mixtures of copulas that we consider next. To solve this analytical issue, we compute the information matrix numerically. In our current example, the information matrix can be written as

$$\mathcal{I}_{t-1} = \mathbb{E}_{t-1} [(\nabla_t)^2] \equiv h(f_{t-1}), \quad (50)$$

with the score vector ∇_t as defined in (49). Note that the function $h(\cdot)$ in (50) does not depend on time or on any parameter other than f_{t-1} . We can therefore construct a grid of values $f^{(0)} < \dots < f^{(n)}$ and compute the function value $h(f^{(j)})$ at each of the grid points. Values

at intermediate points can be obtained by cubic spline interpolation or non-parametric kernel smoothing to ensure continuity of first and second derivatives of the likelihood function. The numerical procedure is then as follows. First, choose starting values of the parameter θ and set the starting value f_0 . Using interpolation, compute $h(f_0)$ and use it to scale the score step $s_1 = \nabla_1/h(f_0)$. Compute the new parameter value f_1 through the GAS recursion, and again use interpolation to obtain $h(f_1)$. This process is repeated for the complete sample. Finally, the likelihood can be computed.

The disadvantage of the Clayton copula is that it only has lower tail dependence and no upper tail dependence. Therefore, it is useful to use a symmetrized version of the Clayton copula that allows for non-zero, but different upper and lower tail dependence. The symmetrized Clayton copula is a mixture of the Clayton and the survival Clayton copula. Consider a general mixture of copulas,

$$C(u_1, \dots, u_d) = \sum_{i=1}^m p_i C_i(u_1, \dots, u_d). \quad (51)$$

Define $w_i = p_i C_i / \sum_{j=1}^m p_j C_j$ as the weight of copula i . It is straightforward to derive that for some parameter vector θ ,

$$\frac{\partial \ln c}{\partial \theta} = \sum_{i=1}^m w_i \cdot \frac{\partial \ln c_i}{\partial \theta}, \quad (52)$$

and

$$\frac{\partial^2 \ln c}{\partial \theta \partial \theta'} = \sum_{i=1}^m w_i \cdot \left(\frac{\partial^2 \ln c_i}{\partial \theta \partial \theta'} + \frac{\partial \ln c_i}{\partial \theta} \frac{\partial \ln c_i}{\partial \theta'} \right) - \left(\sum_{i=1}^m w_i \cdot \frac{\partial \ln c_i}{\partial \theta} \right) \left(\sum_{i=1}^m w_i \cdot \frac{\partial \ln c_i}{\partial \theta} \right)', \quad (53)$$

and thus,

$$\mathbb{E}_{t-1} \left[\frac{\partial^2 \ln c}{\partial \theta \partial \theta'} \right] = -\mathbb{E}_{t-1} \left[\left(\sum_{i=1}^m w_i \cdot \frac{\partial \ln c_i}{\partial \theta} \right) \left(\sum_{i=1}^m w_i \cdot \frac{\partial \ln c_i}{\partial \theta} \right)' \right],$$

such that the scores and information matrices of the individual copulas can be used directly to build the driving mechanism of the mixture copula.

The GAS mechanism for the mixture of copulas has an intuitive interpretation. A given observation may have a contribution to the evolution of either α_L or α_U , i.e., to either the upper or lower tail dependence. The contributions are measured in terms of the likelihood of each mixture vis-a-vis the total likelihood. As a result, observations that cluster in the upper tail automatically will contribute to the evolution of α_U , and similarly in the lower tail for α_L . By contrast, Patton's methodology for the symmetrized copula cannot make automatic use of such features, as its driving mechanism is given by averages of $|u_{it} - u_{jt}|$ for both upper and lower tail dependence.

To illustrate the differences between these two, we construct a simulated example. We generate data from the symmetrized Clayton copula. The lower tail dependence coefficient




Figure 7: Comparisons between estimates from the GAS framework and the Patton model for the symmetrized Clayton copula based on a simulated data set.

follows a sinusoid pattern. The pattern of the upper tail dependence is also sinusoid, but with a period that is half as long. This makes it difficult for a model with a uniform observation driving mechanism to capture both upper and lower tail dependence within a single model. We plot the results in Figure 7 for smoothing parameter values $m = 1$ and $m = 10$.

As expected, the driving mechanism based only on averages of $|u_{it} - u_{jt}|$ does capture some of the variation in the dependence coefficients. However, as the same mechanism underlies both types of dependence, it has difficulty in capturing the upper and lower tail dependence dynamics simultaneously. The GAS specification on the other hand is more successful in picking up both dynamics. The estimate is noisier than that of the Patton model, but follows the true dependence pattern much more closely, resulting in a significant increase in the likelihood.

4.6 Time-varying higher order moments

Following the success of the GARCH(1, 1) model, many authors have suggested further generalizations, in particular to the model with Student t errors. Hansen (1994) proposed making the

degrees of freedom time-varying, while Harvey and Siddique (1999), Jondeau and Rockinger (2003) and Brooks et al. (2005) consider models with time-varying skewness and kurtosis. We develop a t-GAS(1,1) model for $y_t = \sigma_{t-1}\varepsilon_t$ where $\varepsilon_t \sim t_{\nu_t}$. The error term is scaled to have unit variance such that σ_{t-1}^2 is the conditional variance while ν_t is a time-varying degrees of freedom. Define the vector of factors as $f_t = (\sigma_t^2, -\ln \left\{ \frac{b-a}{\nu_t-a} - 1 \right\})$ where the latter factor is the inverse of the logit transformation which is used to keep ν_t in the interval $[a, b]$. In our empirical work, we select the interval $[2.01, 30]$ to ensure that the conditional variance exists, i.e. $\nu_t > 2$. We note that it is possible to select the conditional kurtosis as a factor instead of ν_t but for some time series the conditional kurtosis may not exist.

Taking derivatives of the observation density w.r.t. σ_t^2 and ν_t , we obtain the score vector

$$\nabla_t = \begin{bmatrix} -\frac{1}{2\sigma_t^2} + \frac{(\nu_t+1)}{2} \left(1 + \frac{y_t^2}{(\nu_t-2)\sigma_t^2} \right)^{-1} \frac{y_t^2}{(\nu_t-2)\sigma_t^4}, \\ \frac{1}{2} \left\{ \Gamma' \left(\frac{\nu_t+1}{2} \right) - \Gamma' \left(\frac{\nu_t}{2} \right) \right\} - \frac{1}{2\nu_t} - \frac{1}{2} \ln \left(1 + \frac{y_t^2}{(\nu_t-2)\sigma_t^2} \right) + \frac{(\nu_t+1)}{2} \left(1 + \frac{y_t^2}{(\nu_t-2)\sigma_t^2} \right)^{-1} \frac{y_t^2}{(\nu_t-2)^2\sigma_t^4} \end{bmatrix},$$

and with some additional derivations the elements of the information matrix are

$$E_{t-1}[\nabla_t \nabla_t'] = \begin{bmatrix} -\frac{\nu_t}{2\sigma_t^4(\nu_t+3)} & -\frac{3}{2\sigma_t^2(\nu_t+1)(\nu_t+3)(\nu_t-2)} \\ -\frac{3}{2\sigma_t^2(\nu_t+1)(\nu_t+3)(\nu_t-2)} & \frac{1}{4} \left\{ \Gamma'' \left(\frac{\nu_t+1}{2} \right) - \Gamma'' \left(\frac{\nu_t}{2} \right) \right\} + \frac{(\nu_t+4)(\nu_t-3)}{2(\nu_t-2)^2(\nu_t+1)(\nu_t+3)} \end{bmatrix}.$$

The functions Γ' and Γ'' are the digamma and trigamma functions, which are standard in any matrix programming software. Given the results above and the derivatives of the logit transformation, it is straightforward to construct a GAS(1,1) recursion using the reparameterization argument from (20). We label this model the tv-t-GAS(1,1) model.

We consider daily returns on the S&P 500 from February 1989 through April 2008 as an illustration. We compare the tv-t-GAS(1,1) model described above to a t-GAS(1,1) model with constant ν , that is equation (7), and a standard t-GARCH(1,1) model with constant ν as in Bollerslev (1987). Parameter estimates from each of these models are reported in Table 4.6 and estimates of the conditional variance are plotted in panel (i) of Figure 8. Focusing momentarily on the t-GAS(1,1) model versus the t-GARCH(1,1) model, we can see that the log-likelihood values are close to one another with the GARCH(1,1) slightly greater. Both the persistence parameter b_{11} and degrees of freedom are estimated to be larger for the t-GAS(1,1) model than for the t-GARCH(1,1) model. Estimates of the conditional variance in panel (i) are hard to distinguish from one another with the exception of those periods when there are outliers. To see this more clearly, we also plot the differences between the estimates from the two GAS models minus the GARCH model in panel (ii) of Figure 8. In the first half of the sample before 1998, the level of volatility is lower and there are several outliers in the series. The estimated conditional variance from the t-GARCH(1,1) model is larger than from both GAS models. These are the large negative values in panel (ii). The difference in estimated degrees

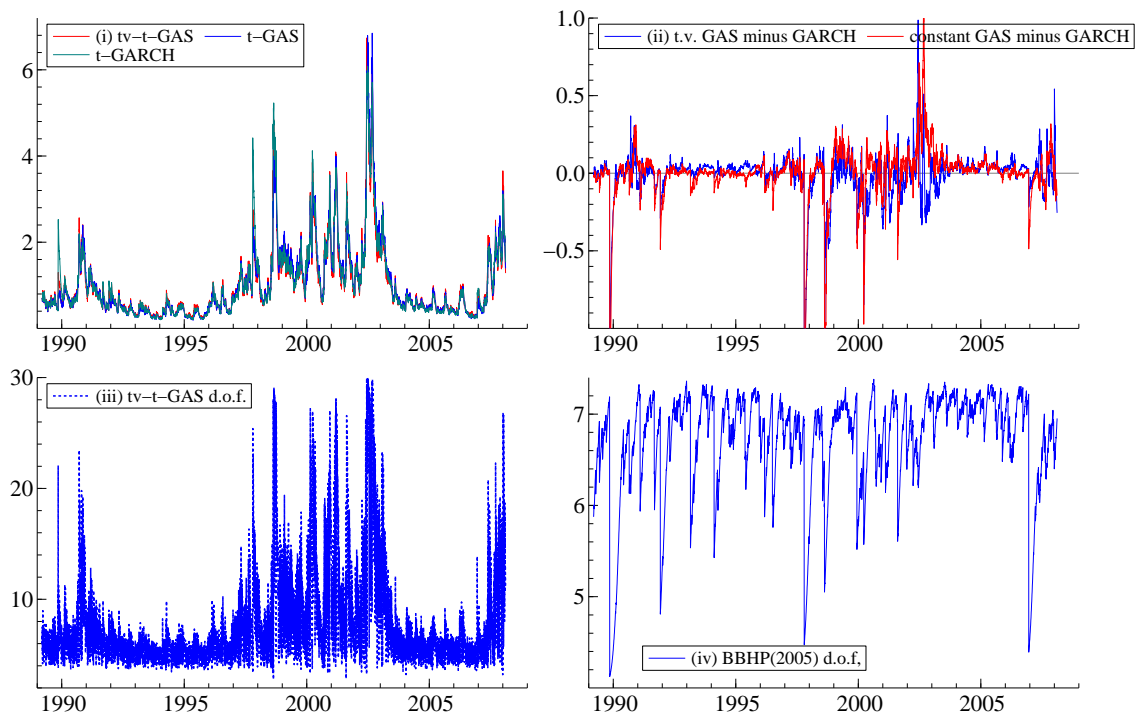


Figure 8: *Time-varying degrees of freedom example: (i) estimated conditional variances from the t -GAS(1,1), t -GARCH(1,1), and tv - t -GAS(1,1) models; (ii) differences between the two GAS(1,1) models and the t -GARCH(1,1) model; (iii) estimated time-varying degrees of freedom from the tv - t -GAS(1,1) model; (iv) estimated time-varying degrees of freedom from the GARCH model of Brooks et al. (2005).*

of freedom is due to the fact that the t -GAS model does not treat outliers like a standard t -GARCH model. From 1998-2003, volatility increases and, relative to this level, large returns are not outliers. Estimates of the conditional variance from the GAS and GARCH models are still significantly different and economically meaningful during this period.

Turning our attention to the tv - t -GAS(1,1) model, the estimated time-varying degrees of freedom from this model is plotted in panel (iii) of Figure 8 and these estimates demonstrate significant variability. The log-likelihood for our new time-varying GAS model increases appreciably relative to the t -GAS(1,1) model. Estimates of the conditional variance in panels (i) and (ii) are reasonably similar to the t -GAS(1,1) model with some differences in 1998-2004 when the time-varying degrees of freedom increases. We compare this model with the time-varying higher-order GARCH model of Brooks et al. (2005), which we label as the tv - t -GARCH(1,1) model. In their model, the conditional kurtosis evolves independently from σ_t^2 according to its own GARCH(1,1) recursion. The implied estimates of ν_t can be calculated straightforwardly.

Table 5: *Estimates from the t-GARCH(1,1), t-GAS(1,1), and tv-t-GAS(1,1) models applied to daily returns of the S&P500 from Feb. 1989 - April 2008. The tv-t-GARCH(1,1) model is from Brooks et. al. (2005). The full sample results are on the left. Split sample results for the t-GAS(1,1) model are on the right.*

	tv-t-GAS	t-GARCH	t-GAS	tv-t-GARCH	t-GAS	t-GAS
					pre-1998	post-1998
ω_1	0.006 (0.005)	0.003 (0.001)	0.004 (0.001)	0.003 (0.001)	0.002 (0.001)	0.007 (0.003)
ω_2	-2.373 (0.310)	-	-	-	-	-
a_{11}	0.057 (0.007)	0.047 (0.007)	0.044 (0.006)	0.049 (0.007)	0.026 (0.006)	0.061 (0.009)
a_{12}	-0.128 (0.043)	-	-	-	-	-
a_{21}	-0.219 (0.033)	-	-	-	-	-
a_{22}	-1.498 (0.002)	-	-	0.005 (0.006)	-	-
b_{11}	0.994 (0.003)	0.951 (0.007)	0.997 (0.002)	0.949 (0.007)	0.997 (0.003)	0.995 (0.004)
b_{12}	0.000 (0.000)	-	-	-	-	-
b_{21}	0.982 (0.154)	-	-	-	-	-
b_{22}	0.026 (0.121)	-	-	0.965 (0.024)	-	-
ν	-	6.699 (0.622)	7.032 (0.677)	-	5.367 (0.610)	10.96 (2.074)
log-like	-6138.18	-6153.02	-6156.46	-6153.44	-2359.55	-3778.63

It is notable result that the estimates of ν_t from our model shown in panel (iii) are significantly different than the implied estimates of ν_t from the tv-t-GARCH(1,1) model of Brooks et al. (2005). In the literature on time-varying higher-order moments, researchers typically force the factors to evolve independently by imposing zero restrictions on b_{12} and b_{21} . The

estimated autoregressive coefficients b_{21} and b_{22} reported in Table 4.6 for the GAS model imply that both σ_t^2 and ν_t are driven by the same factor because b_{22} is close to zero. Accordingly, the estimates of ν_t in panel (iii) exhibit a similar pattern with the conditional variance in panel (i). Estimates of ν_t from the tv-t-GARCH(1,1) model, which imposes these restrictions, result in a different behavior for the time-varying degrees of freedom. The parameter b_{22} is estimated to be significant and persistent in this model.

To investigate this result further, we split the sample in half before and after 1998 and estimated ν using the t-GAS(1,1) model with constant degrees of freedom on the two sub-samples. Estimates from this model on the two sub-samples are reported in the right-hand columns of Table 4.6. The degrees of freedom parameter and its standard error clearly increase in the second half of the sample. Estimates of ν on the two sub-samples from the t-GARCH(1,1) model (not reported) are similar. Although this result may seem counterintuitive initially, the reason is that large returns during this period are no longer extreme outliers because the conditional volatility σ_t^2 is higher. This provides some support for estimates of ν_t from our model and some evidence that modelling higher-order moments independently of the conditional variance may be inappropriate. The models described in this section might be improved further by linking the time-varying behavior of the degrees of freedom with a time-varying level parameter ω_t in the variance. We leave this extension to future research.

4.7 Time-varying multinomial

Trade by trade financial transaction prices lie on a discrete grid with most price changes taking only a small number of values. Russell and Engle (2005) proposed modelling this behavior using a conditional multinomial distribution with time-varying probabilities in conjunction with their ACD model. Here, we construct a GAS version of their model. Consider the case where the observed series y_t , for $t = 1, \dots, n$, has a J -dimensional multinomial distribution with vector of probabilities π_t and let $\pi_{j,t}$ be the j -th element of this vector. The vector of factors f_t will have dimension $J-1$ with elements $f_{jt} = \ln \left(\frac{\pi_{jt}}{1 - \sum_{j=1}^{J-1} \pi_{jt}} \right)$ where the final probability $\pi_{J,t}$ is determined by the constraint that they sum to one. Next, denote \tilde{y}_t and $\tilde{\pi}_t$ as the corresponding $J-1$ dimensional vectors with the J -th element omitted. The score with respect to $f_{j,t-1}$ is

$$\nabla_{jt} = \tilde{y}_{jt} - \tilde{\pi}_{j,t-1}, \quad (54)$$

while the diagonal and off-diagonal elements of the information matrix are

$$\mathcal{I}_{ii,t-1} = \tilde{\pi}_{i,t-1}(1 - \tilde{\pi}_{i,t-1}), \quad (55)$$

$$\mathcal{I}_{ij,t-1} = -\tilde{\pi}_{i,t-1}\tilde{\pi}_{j,t-1}. \quad (56)$$

Combining these results, a GAS(p, q) model for the multinomial distribution reduces to

$$f_t = \omega + \sum_{i=0}^{q-1} A_i \mathcal{S}_{t-i-1} (\tilde{y}_{t-i} - \tilde{\pi}_{t-i-1}) + \sum_{j=1}^p B_j f_{t-j}, \quad (57)$$

where the scale matrix $\mathcal{S}_{t-1} = \mathcal{I}_{t-1}^{-1}$ can be constructed from (55) and (56). The ACM model of Russell and Engle (2005) can be obtained as a special case of the GAS model (57) by selecting the scale matrix \mathcal{S}_{t-1} to be the identity matrix. These authors also add the expected durations from an ACD model as explanatory variables in (57).

As an empirical illustration, we use transaction data from the NYSE TAQ database on Royal Dutch Shell A (RDSA) for the month of November 2007. After retaining trades between 9:30 and 4:00, there are 61,690 trades remaining. Panels (i)-(ii) of Figure 9 contain the observed price changes and observed durations for the first 23,500 trades, while panel (iii) is a histogram of all the trades. The observed durations give evidence of diurnal patterns that are typical of transactions data. In addition, the observed price changes indicate that the probabilities should contain a similar diurnal pattern, as trades with large tick sizes are less likely during opening and closing of the market when volume is higher.

In our sample, 98% of the price changes fall within a ± 5 tick range of zero (see panel (iii)), where a tick is now 1 cent after decimilization of the market in 2001. Decimilization unfortunately causes an increase in the required dimension of the factor f_t and a corresponding increase in the number of parameters to estimate. For this example, f_t will have a minimum of 10 dimensions meaning that the A_0 matrix in an ACM(1,1) model will have 100 parameters. Our solution to this problem is to define new factors \tilde{f}_t as $f_t = d + Z\tilde{f}_t$ where \tilde{f}_t has $\dim(\tilde{f}_t) \ll \dim(f_t)$. The GAS(1,1) model reduces to

$$\tilde{f}_t = A_0 Z' \mathcal{I}_{t-1}^{-1} Z Z' (\tilde{y}_t - \tilde{\pi}_{t-1}) + B_1 \tilde{f}_{t-1}, \quad (58)$$

where the matrix Z must be restricted to identify the model. For illustration purposes, we selected $\dim(\tilde{f}_t) = 3$ and set the upper 3×3 elements of Z equal to the identity matrix for identification. Following Russell and Engle (2005), we include expected durations in (58) and jointly estimate the ACD model. We also restrict the matrices B_j to be diagonal. Specifying a multinomial-GAS(1,2)-ACD(1,2) model for this series, some of the estimated time-varying probabilities for the first third of the data set are shown in panels (v) and (vi) of Figure 9. Panel (v) is a plot of the probability of a price increase of 5 ticks or more while panel (vi) plots the probability of no price movement. The model picks up the diurnal dynamics of the price changes reasonably well with the probability of an increase of 5 ticks or more changing considerably throughout the day. An alternative observation driven model for trade-by-trade data has been proposed by Rydberg and Shephard (2003) using the GLAR methodology of




Figure 9: *Time-varying multinomial GAS(1,2)-ACD(1,2) model. (i) observed price changes; (ii) observed durations; (iii) histogram of price changes; (iv) estimated expected duration from the ACD model; (v) estimated probability of an increase of 5 ticks or more; (vi) estimated probability of a trade with no change in price.*

Shephard (1995). We note that a GAS version of their model will be slightly different but close to their specification.

4.8 Dynamic mixtures of models

The GAS specification can provide a mixture framework for probabilities of several competing, possibly, time-varying models. Assume we have a mixture model with J components where each component or sub-model has a likelihood \mathcal{L}_{jt} . Define the vector of GAS factors as the time-varying mixture probabilities π_{jt} , which defines a new mixture model

$$\mathcal{L}_t = \sum_{j=1}^J \pi_{jt} \mathcal{L}_{jt}. \quad (59)$$

We parameterize the π_{jt} 's using the logit transformation to ensure that the probabilities remain in the zero-one interval. The GAS factors are

$$\pi_{jt} = \frac{e^{f_{jt}}}{1 + \sum_{k=1}^{J-1} e^{f_{kt}}} \Leftrightarrow f_{jt} = \ln(\pi_{jt}) - \ln \left(1 - \sum_{k=1}^{J-1} \pi_{kt} \right), \quad (60)$$

for $j = 1, \dots, J-1$ with the probability of the last component determined by the constraint $\pi_{Jt} = 1 - \sum_{k=1}^{J-1} \pi_{kt}$. Taking the derivative of the log-likelihood with respect to $f_{j,t-1}$, we obtain the elements of the score vector

$$\frac{\partial \mathcal{L}_t}{\partial f_{j,t-1}} = \frac{\pi_{j,t-1} \mathcal{L}_{jt}}{\sum_{k=1}^J \pi_{k,t-1} \mathcal{L}_{kt}} - \pi_{j,t-1}, \quad (61)$$

for $j = 1, \dots, J-1$. The interpretation of (61) is intuitive. The probability of model j is increased if the relative likelihood of model j is above its expectation $\pi_{j,t-1}$. Otherwise, it is decreased. The information matrix for this GAS model is not easy to compute analytically. In our empirical example below, we use a mixture of two normal densities $\phi_j(y)$ for $j = 1, 2$ implying an information matrix of the form

$$E_{t-1}[\nabla_t \nabla_t'] = \pi_{1,t}(1 - \pi_{1,t}) E_{t-1} \left[\left(\frac{\phi_1(y) - \phi_2(y)}{\pi_{1,t}\phi_1(y) + (1 - \pi_{1,t})\phi_2(y)} \right)^2 \right],$$

where the expectation is taken with respect to the mixture distribution. We use numerical integration to compute the information matrix, which is feasible when the mixture model (59) contains say $J = 5$ components or fewer.

To illustrate the methodology, we consider a time series of quarterly log U.S. real GDP growth rates from 1947(2) to 2008(2) obtained from the Federal Reserve Bank of St. Louis. The GAS model is a mixture of two normals with different means μ_i for $i = 1, 2$ and a common variance σ^2 . The GAS factor is the probability that the data comes from the normal distribution

with low mean indicating the probability of a recession. The initial conditions of the factors $f_{k,0}$ are estimated as part of the static parameter vector θ , while the GAS(1,1) updating equation is adopted with an information smoothed scaling matrix S_t as in (14) with $\alpha = 0.05$. This GAS model provides an observation driven alternative to a hidden Markov model (HMM). We compare it to a simplified version of Hamilton's (89) model with no autoregressive dynamics, that is

$$\begin{aligned} y_t &= \mu_t + \varepsilon_t, & \varepsilon_t &\sim \mathcal{N}(0, \sigma^2), \\ \mu_t &= \begin{cases} \mu_1 & \text{if } S_t = 0 \\ \mu_2 & \text{if } S_t = 1 \end{cases} \\ p_{ij} &= P(S_t = j | S_{t-1} = i), & i = 0, 1 & \quad j = 0, 1 \end{aligned}$$

In this model, the latent variable S_t is a regime-switching variable indicating whether the economy is in a recession or expansion. We base our comparison on the one-step ahead predicted estimates produced by the hidden Markov model because the GAS factor is effectively a one-step ahead predictor.

Table 6: *Estimates from the GAS(1,1) mixture and hidden Markov models applied to U.S. log real gdp growth rates from 1947(2) to 2008(2). Standard errors are in parenthesis.*

	μ_1	μ_2	σ	ω	A	B	log-like
GAS	0.203	1.129	0.868	0.364	2.415	0.673	-329.45
	(0.008)	(0.005)	(0.003)	(0.018)	(0.121)	(0.006)	
	μ_1	μ_2	σ	p_{11}	p_{22}	-	
HMM	-0.090	1.106	0.830	0.741	0.918		-333.17
	(0.019)	(0.007)	(0.003)	(0.007)	(0.003)		

Estimates of the parameters of both models are reported in Table 4.8. The estimated values for each mean are reasonably close with the recession parameter μ_1 for the HMM model being slightly smaller and negative. Panel (i) of Figure 10 presents the growth rate of log U.S. real GDP along with the estimated conditional mean $\pi_t \mu_1 + (1 - \pi_t) \mu_2$ from the GAS and HMM models. The GAS and HMM estimates nicely follow the changes in the mean of the series. The estimated probabilities of a recession from each model are plotted in panel (ii) of Figure 10. The estimated probabilities from the GAS model reflect the possibility of the model to rapidly adapt to new signals concerning the current behavior of the time series. As a result, we obtain a clear division of regimes (switches) over time as depicted in the graph. In contrast, the one-step ahead predicted probabilities produced by the hidden Markov model do not change

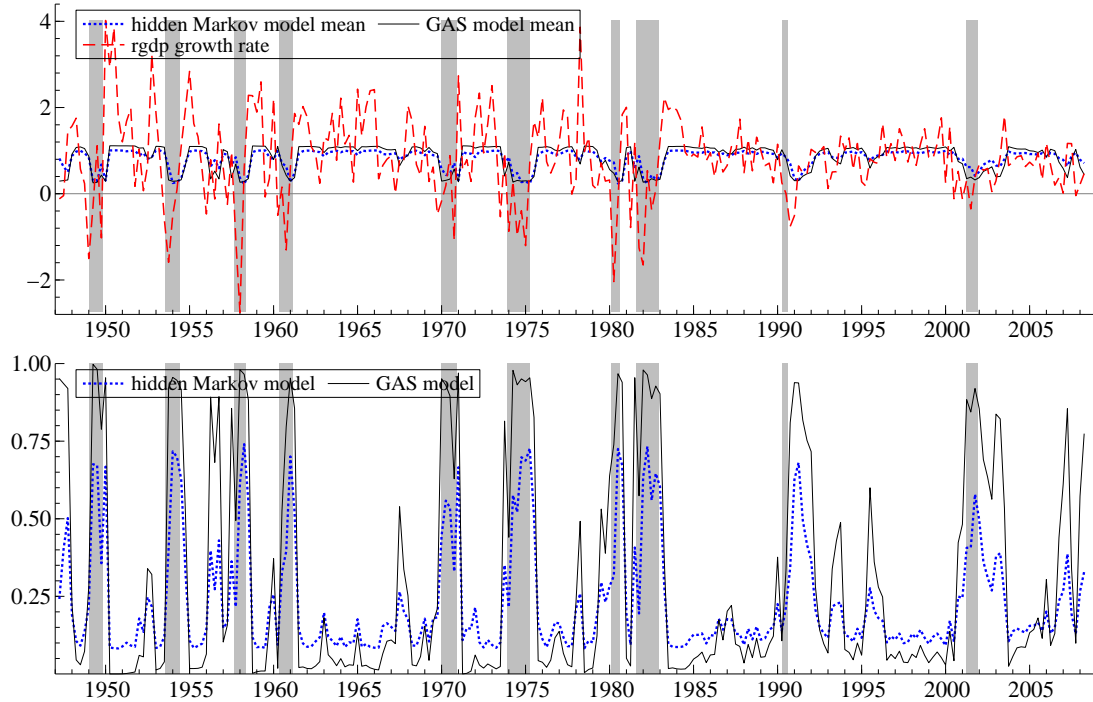


Figure 10: *Mixture model example: (i) growth rate of log U.S. real GDP from 1947(2)-2008(2) and the estimated conditional mean from the GAS(1,1) model and the hidden Markov model; (ii) one-step ahead predicted probability of a recession from each model. NBER recession dates are represented by the shaded regions.*

as rapidly and are not as clear. The GAS model offers a convenient method for forecasting economic downturns. A multivariate model incorporating leading economic variables would be an interesting extension of the GAS model presented here.

5 Simulation experiments

In this section, we provide simulation evidence on the statistical properties of the GAS ML estimators for a selected set of three examples. We concentrate on the marked point process model, the linear state space model with time varying variances, and the Gaussian copula with time-varying correlation.

5.1 The marked point process model

To investigate the statistical properties of the GAS model for the marked point processes of Section 4.2, we consider a simplified version of this model. We consider a cross-section of firms

with two possible ratings, A and B, and possible transitions between them. Neither of the states is absorbing, such that there is no attrition of the panel of firms over time. We consider a panel of $N = 250$ and $N = 2,500$ firms. As the results are comparable, we only present the graphs for $N = 2,500$.

The relevant equations are

$$\begin{aligned}\lambda_{At} &= \eta_1 + f_t, \\ \lambda_{Bt} &= \eta_2 + \alpha f_t, \\ f_t &= A s_t + B f_{t-1},\end{aligned}$$

with $\eta_1 = -3.5$, $\eta_2 = -4.0$, $\alpha = -1$, $A = 0.025$, $B = 0.95$, and where λ_{At} and λ_{Bt} denote the intensities of an A firm becoming a B firm, and of a B firm becoming an A firm, respectively. The parameter values are roughly in line with the empirical estimates concerning the levels of intensities and the magnitude of the systematic factor.

We consider three sample sizes for the time series dimension, namely $T = 20, 50, 100$. We use 1,000 simulations of the true dgp. For each simulation, we compute the ML estimate as well as its t -value based on the numerical second derivative of the likelihood at the optimum. As in the empirical application, we enforce stationarity by parameterizing and estimating the logit transform of B in the GAS equation.

The results are presented in Figure 11 and are very encouraging. The densities of the parameter estimates reveal that for increasing sample sizes T , the estimates peak more at their true values. There is some skewness in the densities of α and A , particularly for smaller sample sizes. If we consider the t -values, however, it appears that the approximation by the normal distribution for purposes of inference is reasonable, even for sample sizes as small as $T = 20$.

5.2 State space models with time-varying variances

The finite sample properties are investigated for the UC model with time-varying variances, see Section 4.4 where the model is considered as an illustration. In particular, the local level model (32) is adopted with the log-variances of the irregular and level disturbances, $\ln \sigma_\varepsilon^2$ and $\ln \sigma_\xi^2$ respectively, treated as GAS factors. The model for the two log-variance factors is given by $f_t = \omega + A s_{t-1} + B f_{t-1}$ where f_t is a 2×1 vector and the 2×2 matrices A and B are both diagonal. The true parameter values in the Monte Carlo study below are chosen to be close to those obtained in the illustration of Section 4.4. To enforce a stable behaviour of the factors over time, the parameters are subject to transformations. The true values for the resulting coefficients are given by

$$\ln \omega_1 = -4.5, \quad \ln \omega_2 = 5.0, \quad \text{logit } A_{11} = 0.1, \quad \text{logit } A_{22} = 1.0,$$




Figure 11: *Simulation densities over 1,000 simulations of a marked point process model. The top panel contains the densities of the parameter estimates, the bottom panels contain the densities of t-values computed using the inverted second derivative of the Hessian at the optimum.*

$$\text{logit } B_{11} = 3.5, \quad \text{logit } B_{22} = 3.2, \quad \text{logit } \alpha = -1.0,$$

where α is the parameter of the EWMA smoothing scheme for the information matrix, see the discussion in Section 2.1. The subscript refers to the element of the associated vector or matrix. Given an observed time series, all coefficients can be estimated using the methods described earlier, see Section 4.4 for further the details.

The Monte Carlo design is similar to the one presented in the previous section. The number of simulations equals 1,000 while the time series dimension is set to $T = 100, 250$ and 1000. Smaller values for T are less interesting for this model. The results of the Monte Carlo study are given in Figure 12.




Figure 12: *Simulation densities over 1,000 simulations of a local level model with time-varying log variances. The top panel contains the densities of the parameter estimates, the bottom panels contain the densities of t -values computed using the inverted second derivative of the Hessian at the optimum.*

5.3 Time-varying Gaussian copula model

In this section, we inspect the finite sample properties of the time-varying Gaussian copula model described in Section 4.5. The model is specified as

$$f_t = \omega + \frac{A}{(1 - \rho_{t-1})(1 + \rho_{t-1})} \cdot \left(y_t - \rho_{t-1} - \rho_{t-1} \frac{x_t - 2}{1 + \rho_{t-1}^2} \right) + B f_{t-1}.$$

where $x_t = \Phi^{-1}(u_{1t})^2 + \Phi^{-1}(u_{2t})^2$ and $y_t = \Phi^{-1}(u_{1t})\Phi^{-1}(u_{2t})$. The parameter settings for the true dgp were $\omega = 0.02$, $A = 0.15$, and $B = 0.96$, while the simulation sample sizes were $T = 200, 400, 600$ observations. To ensure stationarity of the factor f_t and for numerical stability, we restricted A and B during estimation using a logit transformation.

The results from the experiment over 1,000 simulations are shown in Figure 13. The density of the parameter estimates are converging toward their true values as T increases. The rate of convergence appears to be slower for this model than the marked point process model. The densities of the t -values appear slightly biased for the ω and B parameters with the bias diminishing as the sample size increases.




Figure 13: *Simulation densities over 1,000 simulations of a time-varying Gaussian copula model. The top panel contains the densities of the parameter estimates, the bottom panels contain the densities of t -values computed using the inverted second derivative of the Hessian at the optimum.*

6 Conclusions

In this paper, we have introduced Generalized Autoregressive Score (GAS) models. A GAS model is a uniformly applicable observation driven model specification to capture time variation in parameters. We have shown how GAS models encompass other well-known models, such as generalized autoregressive conditional heteroskedasticity models and autoregressive conditional duration and intensity models as well as multiplicative error models and single source of error models. The advantage of the GAS model is that it exploits the full likelihood information. By making a scaled (local density) score step, the time-varying parameter automatically tries to reduce its prediction error at the current observation with respect to current parameter values. Though based on a completely different paradigm, the GAS model provides a powerful and highly competitive alternative to other observation driven models as well as parameter driven models. We have illustrated this extensively by describing a number of non-trivial empirical and simulated examples. Some of these examples are interesting in their own right and provide interesting extensions or alternative specifications for parameter driven models with

time-varying parameters, in particular for state space models with stochastically time-varying parameters, for multivariate marked point processes, and for time-varying copula models.

There are many interesting future research directions. The issues of identification, consistency, stationarity, and asymptotic distribution theory require more work than presented here. Due to its generality and applicability for a wide class of models, however, it appears difficult to come up with an uniform set of conditions for stationarity and consistency that is applicable to all situations of interest. A more promising route may be to formulate conditions for particular sub-sets of models with a GAS specification.

To investigate the finite-sample properties of GAS models in more detail is a second direction for further research. Although we have provided a number of interesting empirical and simulated examples, a more systematic study into the statistical properties of parameter estimates for GAS models may be appropriate.

A third direction for future research concerns the development of misspecification tests for GAS models. On the one hand, we require goodness-of-fit tests and model selection criteria for GAS models. Many of such tests and diagnostics are already developed for the class of GARCH models. On the other hand, the GAS model itself might provide a powerful basis for dynamic misspecification tests. A similar approach to test for the presence of possible GARCH effects is already widely applied in empirical studies. Lagrange multiplier based tests for the possible presence of GAS effects are straightforward extensions of these. Such tests might provide a useful empirical tool for testing for possible time variation in parameters in the context of a large class of nonlinear and non-Gaussian models.

A fourth direction of research is the application of the GAS specification to new models. In this paper, we have tried to review a number of interesting directions of new models with time-varying parameters. However, the GAS framework is not restricted to these, and other new and empirically relevant models with time-varying parameters would provide additional support for the usefulness of GAS as an empirical modeling tool.

Long-memory versions of the GAS model would be another interesting research project. However, the long-memory specification for GAS models is not trivial and therefore more theoretical and empirical research is needed. A related issue is that GAS models may be interpreted as discrete time approximation of their parameter driven counterparts. An interesting research project may be to bridge the gap between GAS models and parameter driven models in a similar continuous time limiting sense as obtained by Nelson (1996) who has bridged the gap between GARCH models and stochastic volatility model specifications.

A sixth direction of future research is to provide a systematic comparison of the advantages and disadvantages of parameter driven versus observation driven models in a wider setting

than GARCH and ACI. Given numerical advances for non-linear and non-Gaussian state space models, and given the general applicability of the current GAS specification, such comparisons have become feasible.

Finally, there are various computational details that need to be studied in further detail. Three issues of particular interest are: finding starting values, finding the required degree of information smoothing for the GAS updating step in particular models, and finding better numerical approximations to the scaling matrix if it cannot be computed analytically. With respect to the first issue, our findings so far are mixed. In relatively straightforward models, the problem of finding appropriate starting values does not exist. In particular, if the information matrix is clearly non-singular for all sample observations, the maximum likelihood maximization algorithm converges quickly and robustly. Introducing information smoothing as well as finding reasonable starting values become more relevant when an observation contains limited or no information on the parameter of interest. This is particularly relevant if there are regions with a degenerate information matrix. In our experience, some degree of information smoothing is indispensable in such cases. In addition, automatic smoothing by estimating the smoothing parameter directly from the data has increased the likelihood value in several cases. In our current implementations, however, the information smoothing is rather rigid. One could consider more involved specifications, where the degree of smoothing also depends on the current position in the sample and the parameter space. The third issue concerns further progress that is needed for models where the information matrix cannot be computed analytically. In the illustration of time-varying copulas in Subsection 4.5, we provided some suggestions based on numerical interpolation techniques using kernel smoothing in low-dimensional parameter spaces. Further extensions are needed to develop computationally feasible estimation methods for GAS models with large parameter spaces and (perhaps even) more complicated specifications.

References

- Bauwens, L. and N. Hautsch (2006). Stochastic conditional intensity processes. *Journal of Financial Econometrics* 4(3), 450.
- Bauwens, L., S. Laurent, and J. Rombouts (2006). Multivariate GARCH models: a survey. *Journal of Applied Econometrics* 21(1), 79–109.
- Benjamin, M. A., R. A. Rigby, and M. Stanispoulos (2003). Generalized autoregressive moving average models. *Journal of the American Statistical Association* 98(461), 214–223.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31, 307–327.
- Bollerslev, T. (1987). A conditionally heteroskedastic time series model for speculative prices and rates of return. *The Review of Economics and Statistics* 69(3), 542–547.

- Bollerslev, T., R. F. Engle, and D. B. Nelson (1994). ARCH models. *Handbook of Econometrics 4*, 2959–3038.
- Brooks, C., S. P. Burke, S. Heravi, and G. Persaud (2005). Autoregressive conditional kurtosis. *Journal of Financial Econometrics 3*(3), 399–421.
- Cipollini, F., R. F. Engle, and G. M. Gallo (2006). Vector multiplicative error models: representation and inference. *NBER Working Paper*.
- Clark, P. K. (1989). Trend reversion in real output and unemployment. *Journal of Econometrics 40*, 15–32.
- Cox, D. R. (1958). The regression analysis of binary sequences (with discussion). *Journal of the Royal Statistical Society, Series B 20*(2), 215–242.
- Cox, D. R. (1981). Statistical analysis of time series: some recent developments. *Scandinavian Journal of Statistics 8*, 93–115.
- Davis, R. A., W. T. M. Dunsmuir, and S. Streett (2003). Observation driven models for Poisson counts. *Biometrika 90*, 777–790.
- Dias, A. and P. Embrechts (2004). Dynamic copula models for multivariate high-frequency data in finance. *Manuscript, ETH Zurich*.
- Diebold, F. X. and C. Li (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics 130*, 337–364.
- Diebold, F. X., G. Rudebusch, and S. Aruoba (2006). The macroeconomy and the yield curve. *Journal of Econometrics 131*, 309–338.
- Doucet, A., N. de Freitas, and N. Gordon (2001). *Sequential Monte Carlo Methods in Practice*. New York, NY: Springer-Verlag.
- Duffie, D., A. Eckner, G. Horel, and L. Saita (2006). Frailty correlated default. *manuscript Stanford University*.
- Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Oxford: Oxford University Press.
- Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica 50*(4), 987–1007.
- Engle, R. F. (2001). GARCH 101: The use of ARCH/GARCH models in applied econometrics. *Journal of Economic Perspectives 15*(4), 157–168.
- Engle, R. F. (2002a). Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics 20*(3), 339–350.
- Engle, R. F. (2002b). New frontiers for ARCH models. *Journal of Applied Econometrics 17*, 425–446.
- Engle, R. F. and T. Bollerslev (1986). Modelling the persistence of conditional variances. *Econometric Reviews 5*(1), 1–50.
- Engle, R. F. and G. M. Gallo (2006). A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics 131*, 3–27.
- Engle, R. F. and S. Manganelli (2004). CAViaR: conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics 22*(4), 367–382.
- Engle, R. F. and J. R. Russell (1998). Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica 66*(5), 1127–1162.

- Fiorentini, G., G. Calzolari, and L. Panattoni (1996). Analytic derivatives and the computation of GARCH estimates. *Journal of Applied Econometrics* 11, 399–417.
- Ghysels, E., A. C. Harvey, and E. Renault (1996). Stochastic volatility. *Handbook of Statistics* 14, 119–191.
- Hamilton, J. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica* 57(2), 357–384.
- Hansen, B. E. (1994). Autoregressive conditional density estimation. *International Economic Review* 35(3), 705–730.
- Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge, UK: Cambridge University Press.
- Harvey, A. C. and A. Jaeger (1993). Detrending, stylised facts and the business cycle. *Journal of Applied Econometrics* 8, 231–247.
- Harvey, A. C., E. Ruiz, and E. Sentana (1992). Unobserved component time series models with ARCH disturbances. *Journal of Econometrics* 52(1-2), 129–157.
- Harvey, C. R. and A. Siddique (1999). Autoregressive conditional skewness. *Journal of Financial and Quantitative Analysis* 34, 465–487.
- Hosking, J. R. M. (1981). Fractional differencing. *Biometrika* 68(1), 165–176.
- Jondeau, E. and M. Rockinger (2003). Conditional volatility, skewness, and kurtosis: existence, persistence, and comovements. *Journal of Economic Dynamics and Control* 27(10), 1699–1737.
- Koopman, S. J., A. Lucas, and A. Monteiro (2008). The multi-state latent factor intensity model for credit rating transitions. *Journal of Econometrics* 142, 399–424.
- Koopman, S. J., A. Lucas, and B. Schwaab (2008). Forecasting cross-sections of frailty-correlated default. *08-029/4*.
- Nelson, C. R. and A. Siegel (1987). Parsimonious modelling of yield curves. *Journal of Business* 60(4), 473–489.
- Nelson, D. B. (1996). ARCH Models as diffusion approximations. *Modelling Stock Market Volatility: Bridging the Gap to Continuous Time*.
- Ord, J. K., A. B. Koehler, and R. D. Snyder (1997). Estimation and prediction for a class of dynamic nonlinear statistical models. *Journal of the American Statistical Association* 92(440).
- Park, J. Y. and J. Qian (2006). Autoregressive modeling of time-varying densities in functional space. *Working paper Texas A&M University*.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review* 47(2), 527–556.
- Russell, J. R. (2001). Econometric modeling of multivariate irregularly-spaced high-frequency data. *Unpublished manuscript, University of Chicago, Graduate School of Business*.
- Russell, J. R. and R. F. Engle (2005). A discrete-state continuous-time model of financial transactions prices and times: the autoregressive conditional multinomial-autoregressive conditional duration model. *Journal of Business & Economic Statistics* 23(2), 166–180.
- Rydberg, T. H. and N. Shephard (2003). Dynamics of trade-by-trade price movements: decomposition and models. *Journal of Financial Econometrics* 1, 2–25.

- Shephard, N. (1995). Generalized linear autoregressions. *Unpublished manuscript, Nuffield College, University of Oxford*.
- Shephard, N. (2005). *Stochastic Volatility: Selected Readings*. Oxford: Oxford University Press.
- Shephard, N. and M. K. Pitt (1997). Likelihood analysis of non-Gaussian measurement time series. *Biometrika* 84, 653–67.
- Stock, J. H. and M. W. Watson (2007). Why Has U.S. Inflation Become Harder to Forecast? *Journal of Money, Credit, and Banking* 39, 3–33.
- van den Goorbergh, R. W. J., C. Genest, and B. J. M. Werker (2005). Bivariate option pricing using dynamic copula models. *Insurance Mathematics and Economics* 37(1), 101–114.
- Watson, M. W. (1986). Univariate detrending methods and stochastic trends. *Journal of Monetary Economics* 18, 49–75.