



# Dynamic copula models and high frequency data<sup>☆</sup>



Irving De Lira Salvatierra, Andrew J. Patton<sup>\*</sup>

Duke University, United States

## ARTICLE INFO

### Article history:

Received 26 June 2013

Received in revised form 28 August 2014

Accepted 18 November 2014

Available online 22 November 2014

### JEL classification:

C32

C51

C58

### Keywords:

Realized correlation

Realized volatility

Dependence

Forecasting

Tail risk

## ABSTRACT

This paper proposes a new class of dynamic copula models for daily asset returns that exploits information from high frequency (intra-daily) data. We augment the generalized autoregressive score (GAS) model of Creal et al. (2013) with high frequency measures such as realized correlation to obtain a “GRAS” model. We find that the inclusion of realized measures significantly improves the in-sample fit of dynamic copula models across a range of U.S. equity returns. Moreover, we find that out-of-sample density forecasts from our GRAS models are superior to those from simpler models. Finally, we consider a simple portfolio choice problem to illustrate the economic gains from exploiting high frequency data for modeling dynamic dependence.

© 2014 Elsevier B.V. All rights reserved.

## 1. Introduction

This paper proposes a class of models to incorporate high frequency (intra-daily) information into models of the dynamic dependence between lower frequency (e.g., daily) asset returns, modeled using a copula-based approach. Our approach is based on a combination of recent work on general time-varying distributions (Creal et al., 2013), on the incorporation of high frequency data into models for lower frequency conditional second moments (Shephard and Sheppard, 2010; Nouredin et al., 2012; Hansen et al., 2011, 2014), and on dynamic copula models for economic time series (see Patton (2013) for a review).

Unlike variances and covariances, the copula of low frequency returns is not generally a known function of the copula of corresponding high frequency returns. Thus the elegant link between high frequency volatility measures (e.g., realized variance and covariance) and their lower frequency counterparts cannot generally be exploited when considering dependence via the copula function.<sup>1</sup> However, it is still likely that high frequency measures such as realized correlation contain information that is useful for modeling dynamic dependence through a copula model, to the extent that the copula model has parameters that are related to correlation. As Andersen et al. (2003) note, “[t]he essence of forecasting is quantification of the mapping from the past and present into the future. Hence, quite generally, superior estimates of present conditions translate into superior forecasts of the future.” It is this intuition that we seek to exploit in this paper.

<sup>☆</sup> We thank Tim Bollerslev, Jia Li, George Tauchen, and seminar participants at Duke for helpful comments. An appendix containing additional results for this paper is available at <http://www.econ.duke.edu/~ap172/research.html>.

<sup>\*</sup> Corresponding author at: Department of Economics, Duke University, 213 Social Sciences Building, Box 90097, Durham, NC 27708-0097, United States.

E-mail address: [andrew.patton@duke.edu](mailto:andrew.patton@duke.edu) (A.J. Patton).

<sup>1</sup> The literature on realized measures is large and still growing, see Andersen et al. (2006) and Barndorff-Nielsen and Shephard (2007) for surveys.

Similar to the class of “GARCH-X” models, we propose augmenting a generalized autoregressive score model (GAS) of Creal et al. (2013) with realized measures such as realized correlation. We call such models “GRAS” models.<sup>2</sup> In addition to a baseline GRAS model using realized correlation, we also consider models that incorporate realized volatilities, measures of co-jumps, and market volatility. We apply these models to stock returns over the period 2000 to 2010, a total of 2773 trading days, and we consider three pairs of stocks, exhibiting low, medium and high levels of correlation.

We compare our proposed GRAS copula models with models that do not exploit high frequency data, both in-sample and out-of-sample. In sample, we find that including a measure of realized correlation significantly improves the performance of the model. This finding is robust across all three pairs of assets, across the particular choice of realized measure, and across the “shape” of the copula that we assume. (We consider the Normal, rotated Gumbel and the Student’s  $t$  copulas.) We also find evidence that other high frequency measures, such as continuous and jump realized correlations and realized volatilities, contain useful information.

We compare the out-of-sample performance of the competing models in two ways. Firstly, we use the density forecast accuracy test for copulas proposed by Diks et al. (2012). We compare the dynamic copula models both in terms of their fit across the entire support, and in particular on the joint tails of the distribution. When we look at the joint tails of the support, the GRAS model uniformly out-performs the constant parameter copula model and the GAS model, and in a majority of comparisons the difference is statistically significant. Second, we consider a simple portfolio decision problem, where a risk averse investor uses one of the competing copula-based density forecasts to compute her optimal portfolio weights, as in Patton (2004) and Jondeau and Rockinger (2012). We evaluate the out-of-sample utility of these portfolios, and find that investors are generally willing to pay a positive management fee to switch from the constant copula model and GAS model to the GRAS model.

This paper is related to work over the past decade on specifications for time-varying conditional copulas. We build on Patton (2006a), Jondeau and Rockinger (2006) and Creal et al. (2013), who consider models of time-varying copulas where a parametric functional form is assumed, and the parameter is allowed to vary through time as a function of lagged information, similar to the ARCH model for volatility, see Engle (1982).<sup>3</sup> We attempt to bridge the gap between the existing time-varying copula models, which have almost exclusively used lower frequency data, and models from the volatility and correlation forecasting literature, which have successfully used high frequency data, see Shephard and Sheppard (2010), Noureldin et al. (2012), and Hansen et al. (2011, 2014) for example. In a recent related paper, Fengler and Okhrin (2012) use a method-of-moments approach to match the covariance structure implied by a copula-based multivariate model with that estimated using high frequency data. This approach suffers from the fact that it cannot be used with copulas that have more than one free parameter (e.g., the Student’s  $t$  copula), and the authors’ simulation study suggests that it can lead to biased estimators when the level of dependence is high. Our proposed approach overcomes these limitations by using the GAS structure, described in the next section, to link high frequency information to low frequency dependence measures.

The remainder paper is organized as follows. Section 2 provides the general formulation of our GRAS models and outlines the estimation procedure. Section 3 presents our in-sample results, and Section 4 presents the out-of-sample forecasting results, for density forecasting and a portfolio choice problem. Section 5 concludes. An online supplemental appendix contains additional details and results.

## 2. Dynamic copula models and high frequency data

The models used in this paper are based on Sklar’s (1959) theorem, extended to apply to conditional distributions in Patton (2006a). This theorem allows the researcher to decompose a (conditional) joint distribution into marginal distributions and a copula. If  $\mathbf{Y}_t = [Y_{1t}, Y_{2t}]$  has a conditional joint distribution  $\mathbf{F}_t$  and conditional marginal distributions  $F_{1t}$  and  $F_{2t}$ , then we can write:

$$\mathbf{Y}_t | \mathcal{F}_{t-1} \sim \mathbf{F}_t = \mathbf{C}_t(F_{1t}, F_{2t}) \tag{1}$$

where  $\mathbf{C}_t$  is the conditional copula of  $\mathbf{Y}_t$  and  $\mathcal{F}_{t-1}$  is some information set, usually taken as  $\sigma(\mathbf{Y}_{t-1}, \mathbf{Y}_{t-2}, \dots)$ . The copula contains all information about the dependence between  $Y_{1t}$  and  $Y_{2t}$ , and is sometimes called the “dependence function.” From a modeling perspective, the usefulness of Sklar’s theorem arises from the fact that we can construct a joint distribution  $\mathbf{F}_t$  by linking together any two marginal distributions,  $F_{1t}$  and  $F_{2t}$ , with any copula; there is no need for these functions to belong to the same family. Thus it provides a great deal of flexibility in modeling joint distributions. In Section 2.3 we describe the models that we use for the conditional marginal distributions in our empirical work, and below we describe various models for the conditional copula. The discussion below exploits the fact that the conditional copula of  $\mathbf{Y}_t$  can be interpreted as the conditional joint distribution of the probability integral transforms of these variables:

$$\begin{aligned} \text{Let } & U_{it} \equiv F_{it}(Y_{it}), \quad i = 1, 2 \\ \text{then } & \mathbf{U}_t | \mathcal{F}_{t-1} \sim \mathbf{C}_t \end{aligned} \tag{2}$$

<sup>2</sup> The name “GAS-X” is already taken for a line of medications. (We thank Drew Creal and Kevin Sheppard for pointing this out to us.)

<sup>3</sup> Alternative dynamic specifications include regime switching models, see Rodriguez (2007) and Okimoto (2008), and “stochastic copula” models, see Hafner and Manner (2012).

## 2.1. The GAS model

As noted above, we build on ARCH-type models for the dynamic copula, and in particular the class of GAS models proposed by Creal et al. (2013). The GAS specification addresses the problem of the choice of “forcing variable” to use in the equation governing the dynamics of the time-varying parameter. For models of the conditional variance, an immediate choice for this variable is the lagged squared residual, as in the ARCH model, but for models with parameters that lack an obvious interpretation the choice is less clear. Creal et al. (2013) propose using the lagged score of the density model (copula model, in our application) as the forcing variable.<sup>4</sup> Specifically, for a copula with time-varying parameter  $\delta_t$  we have:

$$\begin{aligned} \text{Let } \mathbf{U}_t | \mathcal{F}_{t-1} &\sim \mathbf{C}(\delta_t) \\ \text{then } \delta_t &= \omega + \beta\delta_{t-1} + \alpha s_{t-1} \\ \text{where } s_{t-1} &= S_{t-1} \nabla_{t-1} \\ \nabla_{t-1} &= \frac{\partial \log \mathbf{c}(\mathbf{u}_{t-1}; \delta_{t-1})}{\partial \delta_{t-1}} \end{aligned} \quad (3)$$

and  $S_t$  is a scaling matrix (e.g., the inverse Hessian or its square root). While this specification for the evolution of a time-varying parameter is somewhat arbitrary, Creal et al. (2013) motivate it by showing that it nests a variety of popular and successful existing models: GARCH (Bollerslev, 1986) for conditional variance; ACD (Engle and Russell, 1998) for models of trade durations (the time between consecutive high frequency observations); Davis et al.’s (2003) model for Poisson counts. Harvey (2013) further motivates this specification as an approximation to a filter for a model driven by a stochastic latent parameter, or an “unobserved components” model.

## 2.2. The “GRAS” model

Building on the GAS specification in Eq. (3), we propose augmenting the evolution equation with information based on high frequency data, making this a “realized GAS” model, which, continuing the recent tradition of liberal acronym formation, we call the GRAS model. Consider a GRAS model based on a generic realized measure, denoted  $RM$ :

$$\delta_t = \omega + \beta\delta_{t-1} + \alpha s_{t-1} + \gamma RM_{t-1} \quad (4)$$

To the extent that the realized measure provides useful information about the current dependence parameter,  $\delta_t$ , beyond that contained in the score of the copula likelihood, this model will out-perform the underlying GAS model. One can imagine that the usefulness of realized measures for copula parameters will depend on the type of dependence the copula parameter governs, and the type of dependence captured by the realized measure. For example, a copula parameter that primarily captures rank correlation might be well explained by realized (linear) correlation, while a copula parameter that primarily captures asymmetry between joint upper-tail and joint lower tail dependence might have little to do with realized correlation.

It is worth noting here that by augmenting the GAS model with a realized measure, we are formally outside the “generalized autoregressive score” framework, and the process in Eq. (4) no longer has the interpretation as an approximate filter for an unobservable components model. Instead, it is interpreted as a parsimonious way of capturing dynamics in the conditional copula, with a forcing variable guided by the shape of the copula. Like the many extensions of the original ARCH model of Engle (1982), see Bollerslev (2009) for a glossary, the model proposed here is somewhat arbitrary, and must be judged on its empirical performance. Sections 3 and 4 are devoted to answering that question in detail.

Most parametric copulas have parameters that are constrained to lie in a particular range of values (e.g., a correlation parameter forced to take values only inside  $(-1, 1)$ ). To ensure that this is satisfied, Creal et al. (2013) suggest applying a strictly increasing transformation,  $h$ , (e.g., log, logistic, arc tan, etc.) to the parameter,

$$\varphi_t = h(\delta_t) \Leftrightarrow \delta_t = h^{-1}(\varphi_t) \quad (5)$$

and to model the unconstrained transformed parameter, denoted  $\varphi_t$ . This approach extends directly to GRAS models and we employ this in our analysis below:

$$\varphi_t = \omega + \beta\varphi_{t-1} + \alpha s_{t-1} + \gamma RM_{t-1} \quad (6)$$

where  $s_{t-1} = S_{t-1} \nabla_{t-1}$  and  $\nabla_{t-1} = \partial \log \mathbf{c}(\mathbf{u}_{t-1}; \delta_{t-1}) / \partial \delta_{t-1}$  as above. In all of our GAS and GRAS specifications, we use the (Cholesky) square root of the inverse Hessian matrix as our scale matrix, that is,  $S_t = I_t^{-1/2}$ .

<sup>4</sup> Harvey (2013) and Harvey and Sucarrat (2014) propose a similar method for modeling time-varying parameters, which they call a “dynamic conditional score,” or “DCS,” model.

### 2.3. Estimation and inference

We consider marginal distribution models of the following form:

$$Y_{it} = \mu_i(\mathbf{Z}_{t-1}; \theta_i) + \sigma_i(\mathbf{Z}_{t-1}; \theta_i) \varepsilon_{it}, \quad i = 1, 2, \quad \mathbf{Z}_{t-1} \in \mathcal{F}_{t-1} \quad (7)$$

where  $\varepsilon_{it} | \mathcal{F}_{t-1} \sim F_i(0, 1; \theta_i) \quad \forall t$

That is, we allow for general time-varying conditional means and variances of the individual asset returns, and we assume that the standardized residuals,  $\varepsilon_{it}$ , are *iid* from some parametric distribution with zero mean and unit variance.<sup>5,6</sup> In our empirical work below, we use ARMA models for the conditional mean and GJR-GARCH models (see [Glosten et al. \(1993\)](#)) for the conditional variance, with the number of lags chosen using the BIC.<sup>7</sup> We use the skewed *t* distribution of [Hansen \(1994\)](#) for the distribution of the standardized residuals, and we verify its goodness-of-fit using standard tests described in [Section 3.1](#) below.

When combined with parametric models for the conditional marginal distributions, the GRAS model for the conditional copula defines a dynamic parametric model for the joint distribution. The joint likelihood is then

$$L(\theta) \equiv \sum_{t=1}^T \log \mathbf{f}_t(\mathbf{Y}_t; \theta) = \sum_{t=1}^T \log f_{1t}(Y_{1t}; \theta_1) + \sum_{t=1}^T \log f_{2t}(Y_{2t}; \theta_2) + \sum_{t=1}^T \log \mathbf{c}_t(F_{1t}(Y_{1t}; \theta_1), F_{2t}(Y_{2t}; \theta_2); \theta_c) \quad (8)$$

By the structure of the model above, this model can be estimated in stages, first estimating the marginal distributions and then estimating the copula model conditioning on the estimated marginal distribution parameters. This entails, in general, some loss of efficiency relative to estimating the entire joint distribution model in one step, however it greatly simplifies the computational burden, and [Joe \(2005\)](#) and [Patton \(2006b\)](#) find the loss of efficiency to generally be low.

While estimation of the entire parameter vector  $\theta$  is simplified by doing it in stages, inference on the resulting copula parameter estimates is more difficult than usual, as the estimation error from the marginal distribution stages must be taken into account. [White \(1994\)](#) and [Patton \(2006b\)](#) provide methods for doing so based on a modified “Hessian” matrix. An alternative method is to use a block bootstrap, applied to the returns, see [Gonçalves and White \(2004\)](#) and [Gonçalves et al. \(2013\)](#) for technical details. Unlike some applications of the bootstrap, its use here does not lead to any asymptotic refinements, rather it merely enables one to avoid having to compute large and complicated Hessian matrices. We employ the stationary bootstrap of [Politis and Romano \(1994\)](#), using 100 replications and an average block length of 60, for inference in this paper.

## 3. Empirical analysis of U.S. equity returns

### 3.1. Data description and marginal distribution estimation

We use high frequency transaction data over the period January 2000 to December 2010, a total of 2773 trading days, taken from the NYSE's TAQ database. We follow the cleaning rules outlined in [Barndorff-Nielsen et al. \(2009\)](#), see [Li \(2013\)](#) for details. All of our “realized measures” are constructed using five-minute returns from within the trade day, with the overnight returns omitted. Our daily returns are computed as the log-difference of the close prices from the high frequency data base, and are adjusted for stock splits using information from CRSP.<sup>8</sup>

<sup>5</sup> Inference methods for models with dynamic conditional copulas and nonparametric or semiparametric marginal distributions are not yet available in the literature, see [Patton \(2013\)](#), and so we are constrained to consider fully parametric marginal distributions. Also, we note here that the literature on GAS models is growing and some results on the asymptotic properties of MLE for *univariate* GAS models are available (see [Blasques, et al. \(2012\)](#)), but no formal results are yet available for GAS models applied to copulas, and similarly nothing that could be directly used for GRAS models. [Han \(2015\)](#) provides some results on univariate GARCH-X processes, which is related to the extension from GAS to GRAS models, but is not directly applicable. We simply assume that ML estimation of GAS and GRAS models for copulas have the usual asymptotic properties.

<sup>6</sup> It is possible to allow the conditional distribution of the standardized residuals to vary through time (e.g., to have time-varying higher-order moments), but for simplicity we do not consider this here.

<sup>7</sup> As mentioned in the [Introduction](#), recent work in volatility modeling has found that including high frequency variance measures in a GARCH-type equation to improve goodness-of-fit, see [Shephard and Sheppard \(2010\)](#) and [Hansen, et al. \(2011\)](#) and for example. In the supplemental appendix we also present results when the “realized GARCH” model of [Hansen, et al. \(2011\)](#) is used in place of the GJR-GARCH model.

<sup>8</sup> The dates and the corresponding ratio of the splits for each stock are: (i) 2-for-1 for Microsoft on May 23, 1994; December 9, 1996; February 23, 1998; March 29, 1999; February 18, 2003. (ii) 2-for-1 for Apple on June 21, 2000; February 28, 2005. (iii) 2-for-1 for Exxon on April 14, 1997; July 19, 2001. (iv) 2-for-1 for Chevron on June 13, 1994; September 13, 2004. (v) 2-for-1 for Celgene on October 25, 2004; February 27, 2006; 3-for-1 on April 17, 2000. (vi) 2-for-1 for J&J on June 12, 1996; June 13, 2001.

**Table 1**

Summary statistics and marginal distribution estimates.

	Microsoft	Apple	Exxon	Chevron	Celgene	J&J
<i>Panel A: Summary statistics</i>						
Mean	−0.027	0.119	0.022	0.028	0.085	0.010
Std dev	2.209	2.958	1.745	1.770	3.618	1.374
Skewness	−0.083	−0.138	0.178	0.158	−0.140	−0.788
Kurtosis	11.221	6.390	14.374	14.230	8.714	18.733
Correl (linear/rank)	0.437/0.478		0.851/0.803		0.163/0.218	
<i>Panel B: Conditional mean</i>						
Constant	−0.027	0.119	0.026	0.034	0.085	0.010
AR(1)			−0.138	−0.118		
AR(2)				−0.079		
<i>Panel C: Conditional variance</i>						
Constant	0.079	0.106	0.057	0.068	0.059	0.020
ARCH	0.041	0.026	0.029	0.025	0.014	0.031
Asym ARCH	0.074	0.050	0.076	0.088	0.032	0.110
GARCH	0.906	0.937	0.907	0.900	0.964	0.905
<i>Panel D: Skew t density</i>						
DoF	4.498	5.977	9.464	18.768	5.027	5.881
Skew	−0.008	0.031	−0.125	−0.109	0.017	−0.020
<i>Panel E: GoF tests</i>						
KS p-value	0.708	0.088	0.066	0.380	0.543	0.785
CvM p-value	0.507	0.059	0.109	0.315	0.588	0.842

Notes: The top panel of this table presents summary statistics on the daily returns for six U.S. firms, listed in the column headings. The second panel presents the parameter estimates for AR(p) models of the conditional means of these returns, and the third panel presents parameter estimates for GJR-GARCH(1,1) models of the conditional variance. The fourth panel presents parameter estimates for Hansen's (1994) skew *t* density for the standardized residuals. The bottom panel presents simulation-based *p*-values from Kolmogorov–Smirnov and Cramer–von Mises tests of the goodness-of-fit of the density specification.

We consider three pairs of stocks for this analysis: Microsoft and Apple, Exxon and Chevron, Celgene and Johnson & Johnson.<sup>9</sup> These pairs cover a range of (average) levels of dependence, with sample linear correlations of 0.44, 0.85 and 0.16 respectively, to allow us to see whether the gains from using high frequency data are greater or lesser depending on the level of dependence. The top panel of Table 1 presents some summary statistics for the daily returns used in our analysis, and Fig. 1 presents a scatter plot of the daily returns on these assets. (See Figs. 2 and 3.)

As noted in Section 2.3, we use ARMA–GJR–GARCH models to capture the marginal distribution dynamics, with lag lengths chosen using the BIC.<sup>10</sup> The second and third panels of Table 1 present the parameter estimates for the conditional mean and variance models. The estimated parameters are similar to other studies of daily equity returns: generally small AR coefficients, strongly persistent volatility, with the asymmetric ARCH coefficient being two to four times larger than the standard ARCH coefficient, indicating asymmetric volatility dynamics.

We use Hansen's (1994) skewed *t* for the distribution of the standardized residuals, and the estimated parameters are presented in the fourth panel of Table 1. The estimated degrees of freedom parameters, which controls the degree of kurtosis, range from 4.5 to 18.8, indicating excess kurtosis. The skewness parameter, which lies in the interval  $[-1, 1]$  for this distribution, is generally small, indicating evidence of only mild skewness. The bottom panel of Table 1 presents the Kolmogorov–Smirnov and Cramer–von Mises tests of the specification of the marginal distribution. We use a simulation-based approach to obtain critical values that are correct in the presence of estimated parameters, see Patton (2013) for details. All of the series pass both of these tests at the 0.05 level, although there is some mild evidence of misspecification for Apple, and potentially for Exxon. We proceed with this specification and move on to the estimation of the dynamic copula models.

### 3.2. High frequency data and dynamic copula models

In Table 2 we present our first empirical contribution. This table reports the parameter estimates from three different GRAS copula models. The first is based on a Normal copula, and we model the arc tan of this parameter evolving according to Eq. (6). This copula rules out tail dependence and asymmetric dependence, and is presented as a benchmark model, similar to the DCC model of Engle (2002). The second is based on a rotated (or survival) Gumbel copula, which allows for lower tail dependence and asymmetry. The Gumbel copula parameter must lie in  $[1, \infty)$  and we impose this by modeling  $\varphi = \log(\delta - 1)$ . The third model uses a Student's *t* copula,

<sup>9</sup> Moving from bivariate to higher dimension analyses significantly reduces the number of existing copula models that can be applied, and significantly increases the complexity of the problem (see Christoffersen, et al. (2012) and Oh and Patton (2013) for recent work in this area). Addressing these issues in a satisfactory way would take us beyond the scope of the current paper, and we leave this extension for future work.

<sup>10</sup> We consider ARMA (*p, q*) models up to order (5,5). For four of the six stocks the BIC-optimal model is an AR(0), i.e., just a constant, while for Exxon the optimal model is an AR(1) and for Chevron an AR(2) is selected. We consider volatility models in the set ARCH(1), ARCH(2), GARCH(1,1), GARCH(2,2), GJR-GARCH(1,1) and GJR-GARCH(2,2). For all six stocks the BIC-optimal model is a GJR-GARCH(1,1) specification.




Fig. 1. This figure shows a scatter plot of daily returns on three pairs of stocks. We use data over the period January 2000 to December 2010.

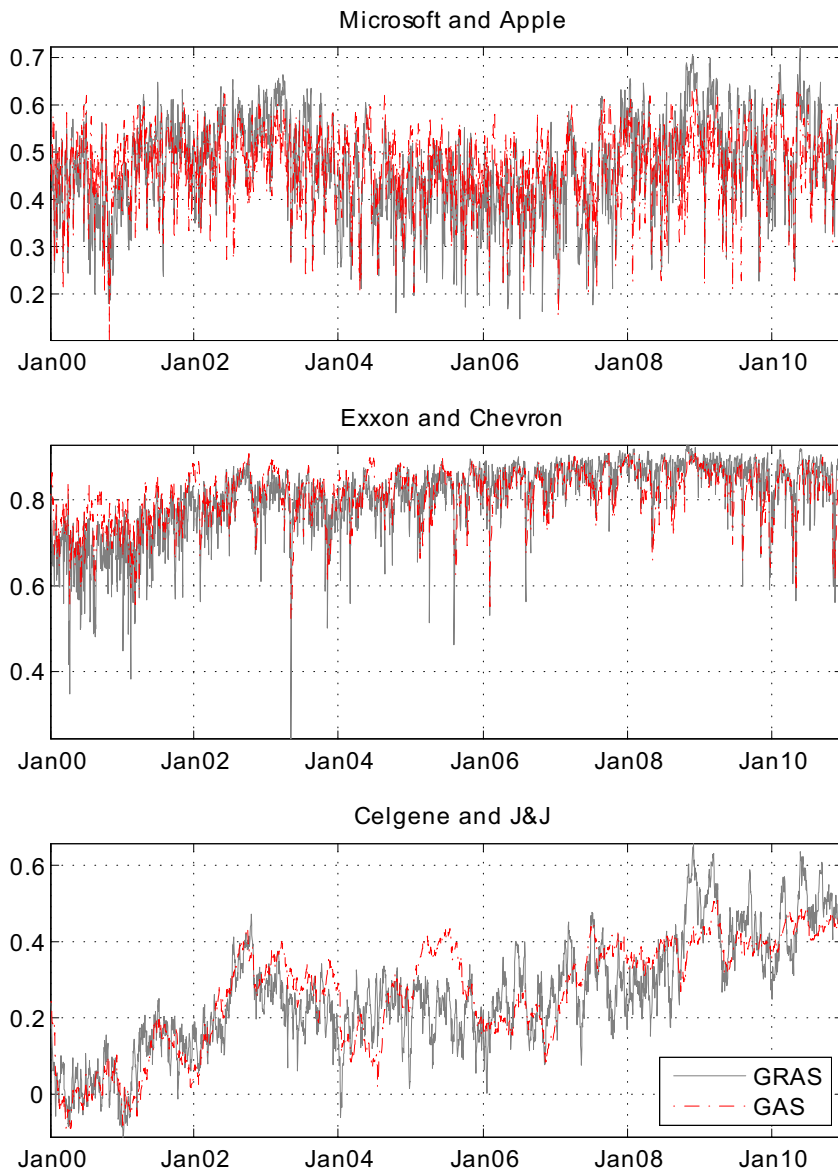
which allows for dependence in both tails, but imposes symmetry. The Student's  $t$  copula has two parameters,  $\rho$  and  $\nu^{-1}$ , and for simplicity we only allow the correlation parameter  $\rho$  to vary through time. Like the Normal copula, we model the arc tan of this parameter. In all three models, the GRAS dynamics are obtained using Eq. (6), with the realized measure being realized correlation, computed in the usual way:

$$RCorr_t = \frac{RV_t^{(1,2)}}{\sqrt{RV_t^{(1,1)}RV_t^{(2,2)}}} \quad (9)$$

where  $RV_t^{(i,j)} = \sum_{k=1}^m r_{t,k}^{(i)} r_{t,k}^{(j)}$

and  $m = 78$  is the number of five-minute returns in a trade day. Our daily returns are whole day (“close to close”), but we use only open-to-close realized measures in the GRAS models, as this is the only period of the day in which high frequency data is available.<sup>11</sup>

<sup>11</sup> In the supplemental appendix we present tables corresponding to Table 2 but using 1-minute, 10-minute and 15-minute sampling to obtain realized correlations. The results using 10- and 15-minute samples are similar to those using 5-minute sampling, but worse when using 1-minute sampling, perhaps reflecting greater market microstructure effects at that frequency.



**Fig. 2.** This figure shows the estimated correlation parameter of the Student's  $t$  GRAS model (solid line) against the estimated correlation parameter of the Student's  $t$  GAS model (dotted line) for three different pairs of stocks.

The coefficient on lagged realized correlation in the GRAS model is denoted  $\gamma_{RC}$ , and is reported in the last row of the parameter estimates. Bootstrap standard errors, which take into account estimation error from the marginal distribution parameters, are presented in parentheses below the estimated parameters. We see from this table that the coefficient on realized correlation is significantly different from zero for all three models across all three pairs of assets. In seven out of the nine models it is significant at the 1% level, and for the other two it is significant at the 5% level. This is strong statistical evidence of the usefulness of high frequency data for modeling dynamics in the conditional copula of lower frequency returns, confirming the general intuition of Andersen et al. (2003) for this application.<sup>12,13</sup>

Fig. 2 presents the estimated time series of the linear correlation<sup>14</sup> from the Student's  $t$  copula model using the GAS and GRAS specifications. (The Student's  $t$  GRAS model beats the Normal and rotated Gumbel specifications for all three pairs of assets, discussed

<sup>12</sup> In the online supplemental appendix we present a corresponding table of parameter estimates when  $\gamma_{RC}$  is imposed to be zero, leading to the GAS specification of Creal et al. (2013). Comparing the values of the log-likelihoods from that model with those in Table 2 confirms the gains from including realized correlation in this specification.

<sup>13</sup> In the supplemental appendix we consider all 15 possible pairs out of these six assets, and we present a table of the  $p$ -values from tests comparing the proposed GRAS model to the GAS model. We find that 10 out of these 15 pairs have significant  $p$ -values, indicating gains from using the GRAS model. Interestingly, all of the non-significant pairs involve Apple or Microsoft, perhaps indicating that these technology stocks have different dynamics than the other stocks under consideration.

<sup>14</sup> The linear correlation implied by a copula-based multivariate model is generally not a closed-form function of the parameters of the model, even when the copula has a "correlation" parameter, see Patton (2013). We use simple numerical integration to obtain the model-implied linear correlation.

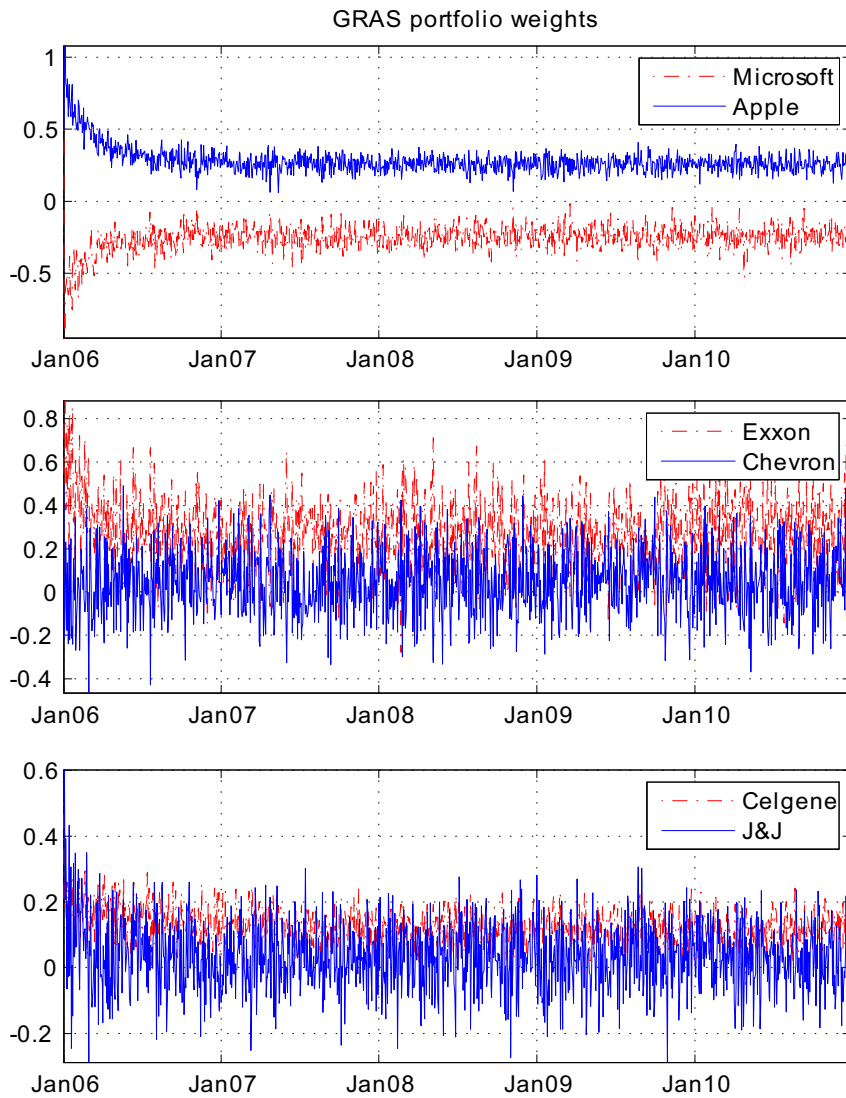


Fig. 3. This figure shows the estimated optimal portfolio weights based on the Student's *t* GRAS model, for an investor with risk aversion of 7, for three pairs of assets.

Table 2  
GRAS with realized correlation.

	Microsoft and Apple			Exxon and Chevron			Celgene and J&J		
	Normal	R. Gumbel	Student's t	Normal	R. Gumbel	Student's t	Normal	R. Gumbel	Student's t
$\omega$	0.099 (0.097)	-0.491 (0.171)	0.094 (0.069)	0.517 (0.238)	-0.188 (0.071)	0.439 (0.205)	0.010 (0.018)	-0.864 (0.271)	0.008 (0.018)
$\alpha$	0.054 (0.031)	0.083 (0.050)	0.084 (0.033)	0.141 (0.034)	0.153 (0.045)	0.166 (0.040)	0.025 (0.020)	0.030 (0.073)	0.028 (0.023)
$\beta$	0.655 (0.158)	0.737 (0.124)	0.688 (0.116)	0.587 (0.145)	0.604 (0.150)	0.610 (0.133)	0.853 (0.097)	0.663 (0.121)	0.879 (0.097)
$\nu^{-1}$			0.111 (0.023)			0.129 (0.021)			0.035 (0.022)
$\gamma_{RC}$	0.614*** (0.179)	0.659*** (0.188)	0.585*** (0.145)	0.717*** (0.201)	0.562*** (0.175)	0.779*** (0.212)	0.392** (0.160)	1.457*** (0.421)	0.329** (0.159)
$p - val_{GAS}$	326.570	331.570	353.320	1506.80	1488.70	1550.80	120.600	113.393	122.400
	0.001	0.000	0.000	0.000	0.001	0.000	0.014	0.001	0.039

Notes: This table presents the estimated parameters from the GRAS model presented in Eq. (6), for the copula of three pairs of assets. The coefficient on realized correlation is denoted  $\gamma_{RC}$ . Estimates of  $\gamma_{RC}$  that are significant at the 1%, 5%, and 10% levels are denoted with 3, 2 and 1 asterisks respectively. The bottom row of this table presents the *p*-value from a test that the realized measure is equal to zero, and thus that the GRAS model simplifies to a GAS model.



**Table 3**

GRAS with realized correlation and realized variance.

	Microsoft and Apple			Exxon and Chevron			Celgene and J&J		
	Normal	R. Gumbel	Student's t	Normal	R. Gumbel	Student's t	Normal	R. Gumbel	Student's t
$\omega$	0.085 (0.161)	-0.812 (0.477)	0.121 (0.164)	0.483 (0.364)	-0.364 (0.496)	0.395 (0.374)	-0.540 (0.405)	-3.254 (0.960)	-0.507 (0.426)
$\alpha$	0.053 (0.036)	0.080 (0.066)	0.078 (0.037)	0.148 (0.032)	0.156 (0.041)	0.171 (0.041)	-0.018 (0.033)	-0.046 (0.097)	-0.010 (0.035)
$\beta$	0.657 (0.200)	0.671 (0.130)	0.688 (0.114)	0.549 (0.167)	0.570 (0.157)	0.589 (0.150)	0.177 (0.185)	0.257 (0.137)	0.191 (0.175)
$\nu^{-1}$			0.109 (0.020)			0.130 (0.020)			0.025 (0.023)
$\gamma_{RC}$	0.528** (0.221)	0.746*** (0.256)	0.500*** (0.180)	0.851*** (0.278)	0.641*** (0.216)	0.869*** (0.267)	0.945*** (0.181)	1.810*** (0.347)	0.924*** (0.188)
$\gamma_{RV1}$	0.055* (0.031)	0.055* (0.033)	0.053** (0.025)	0.122** (0.055)	0.038 (0.058)	0.070* (0.039)	-0.247*** (0.066)	-0.467*** (0.101)	-0.245*** (0.057)
$\gamma_{RV2}$	-0.066* (0.039)	-0.091** (0.042)	-0.059** (0.027)	-0.127* (0.067)	-0.054 (0.071)	-0.075 (0.061)	0.107** (0.050)	0.199* (0.090)	0.110* (0.047)
log L	330.565	336.124	355.494	1509.90	1489.30	1551.70	131.080	126.280	132.020
$p - val_{GAS}$	0.010	0.010	0.000	0.000	0.000	0.000	0.000	0.000	0.000
$p - val_{GRAS}$	0.050	0.020	0.000	0.050	0.500	0.220	0.000	0.000	0.000

Notes: This table presents the estimated parameters from the GRAS model presented in Eq. (6), for the copula of three pairs of assets. The coefficient on realized correlation is denoted  $\gamma_{RC}$ . Estimates of  $\gamma_{RC}$ ,  $\gamma_{RV1}$ , or  $\gamma_{RV2}$  that are significant at the 1%, 5%, and 10% levels are denoted with 3, 2 and 1 asterisks respectively. The penultimate row of this table presents the  $p$ -value from a test that all coefficients on realized measures are equal to zero. The bottom row presents the  $p$ -value from a test that all coefficients on realized measures except for that on the lagged realized correlation are equal to zero.

further in Section 3.4 below.) The bottom panel of this figure, showing the results for Celgene–Johnson & Johnson is perhaps the easiest to interpret. We see from this figure that the GRAS model is able to capture lower and higher values of dependence than the GAS model, which is based only on daily information. As expected, a model based on higher frequency information is able to adjust to changes in dependence faster than a model based only on daily information. The same inference can be made from the other two plots in this figure, although the parameter dynamics in those models are such that it is harder to observe.

### 3.3. Extensions of the baseline GRAS model

We next consider extensions of the GRAS model presented in Eq. (6), where we augment that specification with additional realized measures.

Firstly, we consider an extension motivated by Engle (2002) and Andersen et al. (2006), who find that conditional correlation is affected by the past level of volatility. We augment our GRAS specification from above to include not only realized correlation, but also the log realized variance of each asset. Thus the evolution equation becomes:

$$\varphi_{t+1} = \omega + \beta\varphi_t + \alpha s_t + \gamma_{RC}RCorr_t + \gamma_{RV1} \log(RV_t^{(1,1)}) + \gamma_{RV2} \log(RV_t^{(2,2)}) \quad (10)$$

The results for this model are presented in Table 3. We again see that the coefficient on lagged realized correlation is highly significant across all models. Further, the coefficients on lagged realized volatility are significant for two out of three pairs of assets. The last two rows of this table present joint tests for the significance of the coefficients on realized measures. The penultimate row tests the restriction that all realized measures have coefficients equal to zero, which simplifies this model to a GAS model, and this is strongly rejected. The bottom row tests that only realized correlation has a non-zero coefficient, and this is strongly rejected for the Microsoft–Apple and Celgene–J&J pair, but is not rejected for the Exxon–Chevron pair. Thus this table presents evidence that other measures computed from high frequency data may be helpful in capturing the dynamics of dependence between asset returns.

Next we consider a “heterogeneous autoregressive” (HAR) version of our baseline model, to see whether longer-run dependence in realized measures is useful for modeling the dynamics of daily copulas, see Corsi (2009) and Müller et al. (1997) for details on the HAR model, and Corsi, Pirino and Reno (2012) and Sokolinskiy and van Dijk (2011) for the usefulness of HAR models for forecasting volatility. Consider the following specification:

$$\varphi_{t+1} = \omega + \beta\varphi_t + \alpha s_t + \gamma_{RC}RCorr_t + \gamma_W \frac{1}{5} \sum_{i=1}^5 RCorr_{t+1-i} + \gamma_M \frac{1}{22} \sum_{i=1}^{22} RCorr_{t+1-i} \quad (11)$$

The results for this model are presented in Table 4. We see that the coefficient on lagged realized correlation is again statistically significant across all nine specifications. The coefficients on the past 5-day and 22-day average realized correlation are generally not significant, indicating that these additional lags of realized correlation do not much improve the fit of the model (perhaps due to the persistence that the GAS part of the model already captures). The penultimate row tests the restriction that all realized measures have coefficients equal to zero, and we are unable to reject this in any specification, indicating that high frequency information is not useful.

**Table 4**  
GRAS with HAR realized correlations.

	Microsoft and Apple			Exxon and Chevron			Celgene and J&J		
	Normal	R. Gumbel	Student's t	Normal	R. Gumbel	Student's t	Normal	R. Gumbel	Student's t
$\omega$	0.037 (0.089)	-0.748 (0.417)	0.054 (0.079)	0.535 (0.250)	-0.249 (0.153)	0.475 (0.214)	0.002 (0.048)	-0.026 (0.947)	0.002 (0.046)
$\alpha$	0.059 (0.038)	0.097 (0.068)	0.089 (0.036)	0.153 (0.030)	0.168 (0.037)	0.183 (0.040)	0.023 (0.033)	0.046 (0.097)	0.023 (0.032)
$\beta$	0.570 (0.234)	0.628 (0.211)	0.643 (0.175)	0.553 (0.182)	0.538 (0.235)	0.540 (0.179)	0.968 (0.322)	0.990 (0.366)	0.970 (0.319)
$\nu^{-1}$			0.110 (0.021)			0.127 (0.019)			0.036 (0.021)
$\gamma_{RC}$	0.569** (0.254)	0.646* (0.357)	0.582*** (0.214)	0.536** (0.268)	0.361* (0.217)	0.524*** (0.225)	0.463*** (0.181)	0.821* (0.485)	0.440** (0.189)
$\gamma_W$	-0.313 (0.550)	-0.112 (0.589)	-0.268 (0.463)	-0.517 (0.571)	-0.227 (0.523)	-0.366 (0.500)	-0.410 (0.547)	-0.724 (0.829)	-0.386 (0.519)
$\gamma_M$	0.734 (0.644)	0.533 (0.696)	0.507 (0.555)	0.805* (0.448)	0.572 (0.494)	0.836* (0.463)	0.035 (0.659)	-0.047 (1.270)	0.030 (0.709)
log L	329.189	332.347	354.534	1510.10	1491.20	1553.20	121.710	117.400	123.470
$p - val_{GAS}$	0.270	0.480	0.360	0.260	0.370	0.230	0.840	0.670	0.860
$p - val_{GRAS}$	0.310	0.580	0.410	0.280	0.410	0.240	0.870	0.770	0.860

Notes: This table presents the estimated parameters from the GRAS model presented in Eq. (6), for the copula of three pairs of assets. The coefficient on realized correlation is denoted  $\gamma_{RC}$ . Estimates of  $\gamma_{RC}$ ,  $\gamma_W$ , or  $\gamma_M$  that are significant at the 1%, 5%, and 10% levels are denoted with 3, 2 and 1 asterisks respectively. The penultimate row of this table presents the  $p$ -value from a test that all coefficients on realized measures are equal to zero. The bottom row presents the  $p$ -value from a test that all coefficients on realized measures except for that on the lagged realized correlation are equal to zero.

This contradicts our findings in the previous section, where high frequency data was found to be strongly significant, and is explained by the last row of Table 4, where we test the restriction that the coefficients on the weekly and monthly realized correlation lags are zero. We fail to reject this null in all cases, and we conclude that by including weekly and monthly, as well as daily, realized correlation, we “dilute” its significance. Overall, it appears that just the one-day lag of realized correlation is useful for modeling the dynamic copula.

Finally, we consider a specification that decomposes realized correlation into a “continuous” and a “jump” component, as in Bollerslev et al. (2013). These labels come from an assumption that the continuous time diffusion generating the observed returns is subject to occasional jumps.<sup>15</sup> We construct these measures using the procedure in Bollerslev et al. (2013), described in detail in the supplemental appendix. In brief, five-minute returns are categorized as “jump” or “non-jump” using standard methods, and any five-minute window where at least one of the two assets had a “jump” is excluded from the computation of a “continuous realized correlation.” The difference between the usual realized correlation and the continuous component is defined as the jump realized correlation. This leads us to the specification:

$$\varphi_{t+1} = \omega + \varphi f_t + \alpha s_t + \gamma_{CTS} RCORR_t^{(CTS)} + \gamma_{jump} RCORR_t^{(jump)} \quad (12)$$

This model nests the baseline GRAS model when  $\gamma_{CTS} = \gamma_{jump}$ , and a test of that restriction is presented in the bottom row of Table 5 (denoted “ $p - val_{GRAS}$ ”). We see that for the Microsoft–Apple pair that this restriction is rejected for all three copula specifications. For Exxon–Chevron it is rejected for just the Student's  $t$  copula, and for the Celgene–J&J pair it is not rejected in any case. Thus the additional explanatory power of jump versus continuous correlations is mixed across these assets.

As general conclusions, for all three asset pairs and for all copula specifications, we find that exploiting information from realized correlation significantly improves the fit of the model. Some additional explanatory power can be gained for certain assets by also including realized volatilities, or the continuous and jump components of realized correlations separately, however our baseline GRAS model, which includes just the one-day lag of realized correlation, generally performs very well across these assets. In the remainder of the paper we focus on this baseline GRAS specification.

### 3.4. The choice of copula functional form

Up to this stage of the analysis, we have focused attention on the specification of the *dynamics* of the model for the conditional copula (constant, GAS, or GRAS), and paid little attention to the choice of *shape*. The estimation results presented above also enable us to shed light on the best-fitting copula shape across these three pairs of asset returns, and we discuss these results here.

In Table 6 we compare different specifications of the shape of the copula model, holding the model for the dynamics fixed. The comparison of the rotated Gumbel with the Normal or the Student's  $t$  involves testing non-nested models, and we use the test of Rivers and Vuong (2002), see Patton (2013) for a discussion of its implementation for comparing copula models. Given the tests of

<sup>15</sup> In the presence of jumps, the above analyses based on realized correlation can be interpreted as the correlation implied by the quadratic covariation matrix, which combines both the continuous and jump components.

**Table 5**  
GRAS with continuous and jump realized correlations.

	Microsoft and Apple			Exxon and Chevron			Celgene and J&J		
	Normal	R. Gumbel	Student's t	Normal	R. Gumbel	Student's t	Normal	R. Gumbel	Student's t
$\omega$	0.066 (0.060)	−0.441 (0.236)	0.083 (0.080)	0.504 (0.265)	−0.192 (0.082)	0.441 (0.220)	0.010 (0.019)	−0.447 (0.372)	0.005 (0.020)
$\alpha$	0.043 (0.027)	0.076 (0.055)	0.079 (0.033)	0.146 (0.036)	0.150 (0.046)	0.173 (0.035)	0.025 (0.024)	0.044 (0.086)	0.030 (0.023)
$\beta$	0.758 (0.102)	0.765 (0.145)	0.728 (0.121)	0.586 (0.162)	0.608 (0.165)	0.601 (0.131)	0.855 (0.067)	0.825 (0.156)	0.930 (0.073)
$\nu^{-1}$			0.110 (0.019)			0.128 (0.022)			0.038 (0.022)
$\gamma_{CTS}$	0.405*** (0.158)	0.577* (0.345)	0.495*** (0.201)	0.796*** (0.260)	0.593*** (0.189)	0.858*** (0.208)	0.390** (0.171)	0.718 (0.592)	0.201 (0.156)
$\gamma_{Jump}$	0.798*** (0.282)	0.766*** (0.283)	0.716*** (0.250)	0.418 (0.331)	0.375 (0.237)	0.499* (0.288)	0.381 (0.301)	1.136 (0.947)	0.147 (0.386)
log L	327.421	331.661	353.130	1508.10	1489.30	1552.00	120.600	114.250	122.270
$p - val_{CAS}$	0.010	0.070	0.020	0.000	0.000	0.000	0.330	0.310	0.780
$p - val_{GRAS}$	0.010	0.030	0.020	0.250	0.150	0.080	0.390	0.390	0.820

Notes: This table presents the estimated parameters from the GRAS model presented in Eq. (6), for the copula of three pairs of assets. The coefficients on realized correlation is denoted  $\gamma_{CTS}$  and  $\gamma_{Jump}$ . Estimates of  $\gamma_{CTS}$  or  $\gamma_{Jump}$  that are significant at the 1%, 5%, and 10% levels are denoted with 3, 2 and 1 asterisks respectively. The penultimate row of this table presents the  $p$ -value from a test that all coefficients on realized measures are equal to zero. The bottom row presents the  $p$ -value from a test that  $\gamma_{CTS} = \gamma_{Jump}$ , in which case the model simplifies to the baseline GRAS model.

are non-nested, we use a two-sided test for this comparison. The Student's  $t$  copula nests the Normal copula, and we compare these models via a simple (one-sided) Wald test on the estimated degrees of freedom parameter.<sup>16</sup>

The top panel of Table 6 reveals that the rotated Gumbel model generally out-performs the Normal model, across all three asset pairs and all three models for dynamics, revealed through the positive  $t$ -statistics, however in no case is the difference significant. The lower two panels of this table reveal the improved fit provided by the Student's  $t$  copula: it significantly beats the Normal copula in eight out of nine comparisons, and significantly beats the rotated Gumbel in six out of nine comparisons. This finding confirms the importance of allowing for joint fat tails in multivariate models of asset returns. It is apparent evidence against the importance of allowing for asymmetric dependence (e.g., for dependence being higher during market downturns than market upturns), given that the (symmetric) Student's  $t$  copula out-performs the (asymmetric) rotated Gumbel copula, however this interpretation is muddled by the additional flexibility the Student's  $t$  copula attains from it having two free parameters compared with just one for the Gumbel copula.

#### 4. Out-of-sample forecast comparisons

The previous section showed that high frequency data significantly improves the fit of dynamic copula models for daily asset returns. However, there remains the concern that improved in-sample fit does not guarantee an improvement in out-of-sample forecast performance. In this section we investigate whether high frequency information also leads to gains in out-of-sample forecast performance. We do so using two applications. Firstly, we consider a density forecasting application, and evaluate performance using a metric related to the Kullback–Leibler information criterion (KLIC). Secondly, we consider an illustrative portfolio decision problem, and evaluate the competing models using the realized utility from a portfolio optimized using the predictive density.

##### 4.1. Multivariate density forecasting

We use the approach of Diks et al. (2012), which provides a means of comparing density forecasts over the entire support, as well as in particular regions of the support. Motivated by its connection with the Kullback–Leibler information criterion (KLIC), both of these approaches are based on the out-of-sample log-likelihood of the density forecast. Note that since all of our multivariate models use the same marginal distribution models (means, variances, and standardized residual densities), comparing the full multivariate density forecasts from two models reduces to simply comparing the copula density forecasts.

We use the period from January 2000 to December 2005 as the in-sample period, and January 2006 to December 2010 as our out-of-sample period. Given the computational complexity of the models estimated here, we use a standard rolling window estimation scheme for the marginal distribution parameters, but a fixed window estimation scheme for the copula parameters.<sup>17</sup> In all cases,

<sup>16</sup> Specifically, we test that the estimated inverse degree of freedom parameter is equal to zero. This is on the boundary of the parameter space, which means that the  $t$ -statistic will not be standard Normally distributed, however the right-tail critical values, which are the ones we require, remain applicable.

<sup>17</sup> The asymptotic testing framework of Giacomini and White (2006), which is used by Diks et al. (2012) to implement their test, can handle both of these schemes and so this raises no theoretical issues. The out-of-sample fit of the copula models is presumably worse using a fixed window scheme than using a rolling window scheme, but should not affect our main conclusions on relative performance.

**Table 6**  
Comparison of copula functional forms.

		Microsoft and Apple	Exxon and Chevron	Celgene and J&J
	Constant	0.698	0.545	0.591
Normal vs. R. Gumbel	GAS	0.557	0.166	0.248
	GRAS	0.350	−0.469	−1.478
Normal vs. Student's <i>t</i>	Constant	4.928	7.216	2.249
	GAS	5.981	5.565	1.682
	GRAS	4.489	5.690	1.143
R. Gumbel vs. Student's <i>t</i>	Constant	1.732	3.260	0.312
	GAS	2.174	3.388	0.553
	GRAS	2.311	3.767	2.023

Notes: This table presents *t*-statistics for the comparison of models of copula functional form, for a given choice of model for the copula dynamics (constant, GAS or GRAS). The top and bottom comparisons involve non-nested copula models, and we employ a two-sided Rivers and Vuong (2002) test. The middle panel involves a nested comparison, and we use a standard, one-sided, Wald test, based on the bootstrap standard errors presented in Table 2. In all cases a positive *t*-statistic indicates that the second model listed in the panel label is preferred to the first model.

the density forecast for day *t* is based only on information up until day *t* − 1. Diks et al. (2012) propose a conditional likelihood test to compare density forecasts, where the performance of model *A* in the region  $[0, q] \times [0, q]$  is based on:

$$S_t^A(\mathbf{u}_t) = (\log \mathbf{c}_t^A(\mathbf{u}_t) - \log \mathbf{C}_t^A(\mathbf{q})) \mathbf{1}_{\{\mathbf{u}_t \leq \mathbf{q}\}} \quad (13)$$

where  $\mathbf{q} = [q, q]$  and  $\mathbf{u}_t = [u_{1t}, u_{2t}]$ . Note that when  $q = 1$  this expression simplifies to the log copula density,  $\log \mathbf{c}_t^A(\mathbf{u}_t)$ , and for  $q < 1$  it can be interpreted as the log-likelihood of the model conditional on the observation lying in the region  $[0, q] \times [0, q]$ . The null hypothesis of equal predictive accuracy of models *A* and *B* in the region  $[0, q] \times [0, q]$  is:

$$\begin{aligned} H_0 : E[S_t^A(\mathbf{u}_t) - S_t^B(\mathbf{u}_t)] &= 0 \\ \text{vs. } H_1 : E[S_t^A(\mathbf{u}_t) - S_t^B(\mathbf{u}_t)] &\neq 0 \end{aligned} \quad (14)$$

Diks et al. (2012) propose testing this hypothesis based on a test on the average difference in the conditional likelihoods. Let:

$$d_t = S_t^A(\mathbf{u}_t) - S_t^B(\mathbf{u}_t) \quad (15)$$

then a test that  $E[d_t] = 0$  can be used to test the null hypothesis above. This series may be serially correlated and heteroskedastic, and so robust standard errors are required.<sup>18</sup>

We consider five regions for analysis of the copula density forecasts: the joint 1%, 5% 10% and 25% lower tails, as well as the entire support. The results of these tests are presented in Table 7. The right-most column of this table presents the results for the entire support, and shows that across all three pairs of assets, and across three copula models (Normal, rotated Gumbel, and Student's *t*) our proposed GRAS model significantly beats the constant copula model in all but one case. This shows that the dynamics captured by our copula model based on high frequency data are useful out-of-sample: the predictive copula densities from the GRAS model provide a better out-of-sample fit than those from simple, constant parameter copula models. This conclusion holds true across all joint lower tail regions but one, providing strong support for our model.

A tougher hurdle for our GRAS model is to significantly out-perform the GAS model. The GAS model provides a parsimonious way of capturing dynamics in the conditional copula, and our proposed GRAS model will only out-perform it if the high frequency information we exploit (realized correlation) has additional explanatory power out-of-sample. When looking at the entire support we find that the GRAS beats the GAS model (the *t*-statistics are positive in 8 out of 9 comparisons), but the out-performance is generally not statistically significant. When we use the method of Diks et al. (2012) to zoom in on the joint tails of the support, we find that the GRAS model uniformly out-performs the GAS model, and in a majority of comparisons the difference is statistically significant. Thus it appears that high frequency information is particularly useful for dynamic copula models when interest is focused on the tails of the distribution.

<sup>18</sup> Diks et al. (2012) also propose a "censored likelihood" test as an alternative to the conditional likelihood test. The results from that test are very similar to those discussed here, and are presented in the supplemental appendix.

**Table 7**

Out-of-sample comparison of density forecasts.

		Lower tail region probability				
		0.01	0.05	0.10	0.25	1.00
<i>Microsoft and Apple</i>						
GRAS vs. Const	Normal	1.338*	2.126**	3.232***	3.142***	1.536*
	R. Gumbel	1.291*	1.799**	2.974***	3.279***	2.000**
	Student's <i>t</i>	1.724**	1.923**	3.242***	3.149***	2.014**
GRAS vs. GAS	Normal	1.100	1.471*	2.790***	3.164***	1.192
	R. Gumbel	0.855	0.980	2.086**	2.472***	0.997
	Student's <i>t</i>	1.703**	0.948	2.303**	2.601***	0.932
<i>Exxon and Chevron</i>						
GRAS vs. Const	Normal	2.649**	6.363***	7.845***	8.857***	2.770***
	R. Gumbel	0.871	3.819***	3.523***	3.678***	3.530***
	Student's <i>t</i>	2.598**	6.174***	7.387***	8.934***	3.236***
GRAS vs. GAS	Normal	2.624**	6.120***	7.312***	7.725***	1.081
	R. Gumbel	0.132	2.546***	1.538*	1.143	1.475*
	Student's <i>t</i>	2.297**	5.562***	6.240***	7.033***	1.618*
<i>Celgene and J&amp;J</i>						
GRAS vs. Const	Normal	1.328*	3.130***	4.520***	8.260***	3.100***
	R. Gumbel	1.287*	2.982***	4.191***	7.302***	1.240
	Student's <i>t</i>	1.300*	3.123***	4.524***	8.308***	2.586**
GRAS vs. GAS	Normal	1.384*	2.894***	3.999***	7.153***	0.623
	R. Gumbel	1.397*	2.912***	4.107***	6.584***	−0.375
	Student's <i>t</i>	1.405*	2.912***	4.084***	7.203***	0.138

Notes: This table presents *t*-statistics from pair-wise comparisons of the out-of-sample likelihoods of competing density forecasts. We consider five regions of support over which to compare the competing density forecasts: the joint lower 0.01, 0.05, 0.10 and 0.25 tails, as well as the entire support. For a given copula specification (Normal, rotated Gumbel and Student's *t*) we compare specifications of the dynamics: Constant versus GRAS and GAS versus GRAS. Test statistics that are significant at the 1%, 5%, and 10% levels are denoted with 3, 2 and 1 asterisks respectively.

#### 4.2. A portfolio decision problem

Previous work has documented evidence of two types of asymmetries in the joint distribution of stock returns, namely skewness in the distribution of individual returns, and asymmetry in the dependence between asset returns.<sup>19</sup> With this motivation, we use a portfolio application to gain insights into the economic significance of using high frequency data in dynamic copula models. We approach the portfolio allocation problem following the methodology proposed in Patton (2004) and Jondeau and Rockinger (2012). The latter paper proposes approximating a CRRA utility function using a fourth-order polynomial for tractability:

$$E_t[\mathcal{U}(W_{t+1})] = \varphi_0 + \varphi_1 m_{t+1}^{(1)} + \varphi_2 m_{t+1}^{(2)} + \varphi_3 m_{t+1}^{(3)} + \varphi_4 m_{t+1}^{(4)} \quad (16)$$

with  $\varphi_k = (k!)^{-1} \partial^k \mathcal{U}(W) / dW^k|_{W=1}$ , which are known functions of *k* and the degree of risk aversion, and  $m_{t+1}^{(j)} = E_t[W_{t+1}^j]$  is the uncentered *j*th moment, with  $W_{t+1}$  the value of the portfolio at time *t* + 1.<sup>20</sup> We consider five different levels of relative risk aversion, *RRA* = 1, 3, 7, 10 and 20, similar to the range considered in Campbell and Viceira (1999) and At-Sahalia and Brandt (2001). Optimal portfolio weights are found by maximizing the expected utility under the predictive density:

$$\begin{aligned} \omega_{t+1}^* &= \arg \max_{\omega \in \Upsilon} E_t[\mathcal{U}(1 + \omega' \mathbf{Y}_{t+1})] \\ &= \arg \max_{\omega \in \Upsilon} \int \int \mathcal{U}(1 + \omega_1 y_1 + \omega_2 y_2) \hat{f}_{1,t+1}(y_1) \hat{f}_{2,t+1}(y_2) \hat{c}_{t+1}(\hat{F}_{1,t+1}(y_1), \hat{F}_{2,t+1}(y_2)) dy_1 dy_2 \end{aligned} \quad (17)$$

where  $\Upsilon$  is a compact subset of  $\mathbb{R}^2$  for the unconstrained investor, and the two-dimensional unit simplex<sup>21</sup> for the short-sales constrained investor. For simplicity we take the return on the risk-free asset to be zero. The time series of estimated optimal weights for the intermediate level of risk aversion are presented in Figure 3.

We use the same in-sample and out-of-sample periods as in the previous section, but in this section we evaluate the competing models using the average utility of the portfolio returns formed using weights optimized according to the model. To make the units used in this performance study interpretable, we convert the average utility into a “management fee,” which is the fixed amount,

<sup>19</sup> See Erb et al. (1994), Longin and Solnik (2001), Ang and Chen (2002), and Patton (2004) for example.

<sup>20</sup> The optimal portfolio weights under a CRRA utility function do not depend on the level of initial wealth, and so we set the initial wealth to 1. Thus the end-of-period wealth is equal to the gross return on the portfolio.

<sup>21</sup> That is, for short-sales constrained investors,  $\Upsilon = \{(\omega_1, \omega_2) \in [0, 1]^2 : \omega_1 + \omega_2 \leq 1\}$ .

**Table 8**  
Management fees and realized portfolio return performance.

		Relative risk aversion				
		1	3	7	10	20
<i>Microsoft and Apple</i>						
GRAS vs. Const	Normal	−1.195	−0.235	−0.486	−0.305	−0.193
	R. Gumbel	−5.868	−0.010	−0.061	0.021	0.122
	Student's <i>t</i>	43.003*	10.286**	3.984**	2.781**	1.327**
GRAS vs. GAS	Normal	−8.879	−2.402	−1.434	−1.039	−0.547
	R. Gumbel	−16.682*	−4.594	−1.700	−1.156	−0.465
	Student's <i>t</i>	−11.907	−6.609**	−3.206***	−2.295***	−1.207***
<i>Exxon and Chevron</i>						
GRAS vs. Const	Normal	−3.291	−0.451	−0.325	−0.178	−0.118
	R. Gumbel	6.752	0.688	0.350	0.279	0.109
	Student's <i>t</i>	8.969	1.910	0.776	0.505	0.205
GRAS vs. GAS	Normal	−1.936	0.118	−0.368	−0.179	−0.081
	R. Gumbel	1.472	−1.176	−0.057	−0.057	0.010
	Student's <i>t</i>	17.922	4.923	1.732	1.237	0.552
<i>Celgene and J&amp;J</i>						
GRAS vs. Const	Normal	0.626	1.395	0.385	0.187	0.113
	R. Gumbel	−9.207	−3.549	−1.477	−1.109	−0.534
	Student's <i>t</i>	26.663***	8.015***	3.362**	2.316**	1.178***
GRAS vs. GAS	Normal	1.008	0.962	0.271	0.169	0.100
	R. Gumbel	0.710	−0.237	−0.049	−0.077	−0.030
	Student's <i>t</i>	12.026*	3.907*	1.496*	1.123*	0.549*

Notes: This table presents the “management fee,” in annualized percent returns, that an investor with risk aversion given in the column titles would be willing to pay to switch from the Constant or GAS model to the GRAS model. A block bootstrap is used to ascertain the significance of these fees, and estimates that are significant at the 1%, 5%, and 10% levels are denoted with 3, 2 and 1 asterisks respectively.

$\vartheta$ , that could be charged (or paid, if needed) each period to an investor to switch from model  $\mathcal{A}$  to model  $\mathcal{B}$ . Alternatively, it can be interpreted as the amount that could be deducted from the daily return on portfolio  $\mathcal{B}$  over the out-of-sample period and leave the investor indifferent between portfolio  $\mathcal{A}$  and portfolio  $\mathcal{B}$ . The management fee is the solution to the following equation:

$$\frac{1}{P} \sum_{t=R+1}^{R+P} U(1 + \omega_{A,t+1}^* \mathbf{Y}_{t+1}) = \frac{1}{P} \sum_{t=R+1}^{R+P} U(1 + \omega_{B,t+1}^* \mathbf{Y}_{t+1} - \vartheta) \quad (18)$$

where  $R$  is the length of the in-sample period and  $P$  is the length of the out-of-sample period.

Table 8 presents the management fees estimates, and for all comparisons we consider switching from a smaller model (Constant or GAS) to the larger model (GRAS).<sup>22,23</sup> Consider first the comparison of the constant copula model and the GRAS model. Across all asset pairs, copula shape specifications, and levels of risk aversion, we see that the GRAS model is generally preferred: the management fees are positive in a majority of cases, and are significantly positive for two of the asset pairs (Microsoft–Apple and Celgene–J&J) when combined with the Student's  $t$  copula. For example, for the Microsoft–Apple pair, the annualized management fees that could be charged to switch from a constant  $t$  copula to a GRAS  $t$  copula range from 43.0% (RRA = 1) to 1.3% (RRA = 20), and is 4.0% for the intermediate level of risk aversion.<sup>24</sup> In no case is the management fee negative and significant. This is strong out-of-sample support for our GRAS model, at least against the simple benchmark of a constant copula.

As in the previous section, a tougher benchmark for the GRAS model is the GAS specification. In this case the results are more mixed: for Celgene–J&J we find that the GRAS copula is preferred to the GAS copula, particularly when combined with the Student's  $t$  shape specification. For Exxon–Chevron, the GRAS model is generally preferred, with the performance fees ranging from −1.9% to +17.9%, but these are not statistically significant in any case, indicating that the GAS and GRAS models are equally good out-of-sample, according to this metric. For the Microsoft–Apple pair the GRAS model is beaten by the GAS model, with the performance fees being negative and often significant. Thus the out-of-sample gains from a GRAS model relative to a GAS model from a portfolio decision perspective are somewhat sensitive to the choice of assets and the copula shape specification: in some cases positive and significant, in others negative and significant, and in yet others not different from zero. Further work on the determinants of this variability is clearly desirable, as is a more comprehensive and higher-dimension portfolio decision application.

<sup>22</sup> The supplemental appendix presents some summary statistics on the portfolio returns, as well as results for the short-sales constrained investor, which yields conclusions that are consistent with those for the unconstrained investor. The supplemental appendix also presents plots the optimal portfolio weights for some representative models. For lower levels of risk aversion, these plots reveal substantial portfolio turnover, and in practice it may be desirable to impose turnover constraints on the optimization problem. We do not attempt this here.

<sup>23</sup> We use a stationary bootstrap with average block length of 60 days to determine the significance of the “management fees.” Note that the quantities being bootstrapped here are asset returns, not cumulative wealth, and so issues of non-stationarity or integrated processes do not arise.

<sup>24</sup> In almost all cases, we find that the absolute value of the management fee decreases with the level of risk aversion. This reflects the fact that more risk averse investors shrink their weights on risky assets towards zero, making the differences in portfolio weights from different models smaller.

## 5. Conclusion

This paper proposes a new class of dynamic copula models for daily asset returns that exploits information from high frequency (intra-daily) data. This class of models bridges the gap between the existing time-varying copula models, which have almost exclusively used lower frequency data, and models from the volatility and correlation forecasting literature, which have successfully used high frequency data. We accomplish this by augmenting the generalized autoregressive score (GAS) model of Creal et al. (2013) with high frequency measures, such as realized correlation, to obtain a “realized GAS,” (GRAS) model. While measures of dependence based in high frequency data cannot generally be interpreted as unbiased or consistent estimators of dependence and lower frequencies, our approach builds on the intuition that by better measuring the current degree of dependence, broadly defined, we can build better forecasts of dependence in the future.

We employ equity return data on six U.S. firms over the period 2000 to 2010, and show that including high frequency information in a dynamic model of daily equity return dependence leads to improvements in both goodness-of-fit and out-of-sample forecasts. In sample, we find that the information in realized correlation leads to significant improvements across all assets and copula specifications considered, and we find some evidence that realized volatility, and a measure of correlation that is related to “co-jumps,” improves our baseline GRAS model even more for certain assets. Out of sample, we find that density forecasts based on our GRAS specification out-perform forecasts based on models that ignore high frequency information, and this out-performance is strongly significant when attention is focused on the tail of the joint distribution. We also find some gains in out-of-sample portfolio construction based on these density forecasts, though the gains vary across the assets under consideration.

While this paper provides some initial light on the value of high frequency information for lower frequency copula modeling, some important questions remain open. Firstly, it would be interesting to extend the analysis in this paper to high dimension copula models, along the lines of Christoffersen et al. (2012) and Oh and Patton (2013) for example. Estimates of large-dimension realized covariance matrices have recently attracted attention, see Hautsch et al. (2012) for example, and these could prove useful in an extension of GRAS models to high dimensions. Furthermore, extending and elaborating the portfolio decision application presented above would enable a better understanding of the economic value of high frequency information for dynamic copula modeling. We leave these important extensions for future research.

## Appendix A. Supplementary data

Supplementary data to this article can be found online at <http://dx.doi.org/10.1016/j.jempfin.2014.11.008>.

## References

- Andersen, T.G., Bollerslev, T., Labys, P., 2003. Modeling and forecasting realized volatility. *Econometrica* 71, 3–29.
- Andersen, T.G., Bollerslev, T., Christoffersen, P.F., Diebold, F.X., 2006. Volatility and correlation forecasting. In: Elliott, Graham, Clive, W.J., Granger, W.J., Timmermann, Allan (Eds.), *Handbook of Economic Forecasting*. Elsevier Science B.V, Amsterdam.
- Ang, A., Chen, J., 2002. Asymmetric correlations of equity portfolios. *J. Financ. Econ.* 63, 443–494.
- At-Sahalia, Y., Brandt, M.W., 2001. Variable selection for portfolio choice. *J. Financ.* 56, 1297–1355.
- Barndorff-Nielsen, O.E., Shephard, N., 2007. Variation, jumps, market frictions and high frequency data in financial econometrics. In: Blundell, R., Torsten, P., Newey, W.K. (Eds.), *Advances in Economics and Econometrics Theory and Applications*, Econometric Society Monographs. Cambridge University Press, Cambridge, pp. 328–372.
- Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., Shephard, N., 2009. Realized kernels in practice: trades and quotes. *Econ. J.* 12, 1–32.
- Blasques, F., Koopman, S.J., Lucas, A., 2012. Stationarity and ergodicity of univariate generalized autoregressive score processes. *Tinbergen Institute Discussion Paper* (TI 2012-059/4).
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *J. Econ.* 31, 307–327.
- Bollerslev, T., 2009. Glossary to ARCH (GARCH). In: Bollerslev, Tim, Russell, Jeffrey R., Watson, Mark W. (Eds.), *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*. Oxford University Press, Oxford.
- Bollerslev, T., Li, S.Z., Todorov, V., 2013. Jump tails, extreme dependencies and the distribution of stock returns. *J. Econ.* 172, 307–324.
- Campbell, J.Y., Viceira, L.M., 1999. Consumption and portfolio decisions when expected returns are time varying. *Q. J. Econ.* 114, 433–495.
- Christoffersen, P., Errunza, V.R., Jacobs, K., Langlois, H., 2012. Is the potential for international diversification disappearing? A dynamic copula approach. *Rev. Financ. Stud.* 25, 3711–3751.
- Corsi, F., 2009. A simple approximate long-memory model of realized volatility. *J. Financ. Econ.* 7, 174–196.
- Corsi, F., Pirino, D., Reno, R., 2012. In: Bauwens, L., Hafner, C., Laurent, S. (Eds.), *Handbook of Volatility Models and Their Applications*. New Jersey, Wiley.
- Creal, D., Koopman, S.J., Lucas, A., 2013. Generalized autoregressive score models with applications. *J. Appl. Econ.* 28, 777–795.
- Davis, R.A., Dunsmuir, W.T.M., Streett, S., 2003. Observation driven models for Poisson counts. *Biometrika* 90, 777–790.
- Diks, C., Panchenko, V., Sokolinskiy, O., van Dijk, D., 2012. Comparing the accuracy of copula-based multivariate density forecasts in selected regions of support. *working paper*.
- Engle, R.F., 1982. Autoregressive conditional heteroscedasticity with estimates of the variance of UK inflation. *Econometrica* 50, 987–1007.
- Engle, R.F., 2002. Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *J. Bus. Econ. Stat.* 20, 339–350.
- Engle, R.F., Russell, J.R., 1998. Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica* 66, 1127–1162.
- Erb, C.B., Harvey, C.R., Viskanta, T.E., 1994. Forecasting international equity correlations. *Financ. Anal. J.* 50, 32–45.
- Fengler, M.R., Okhrin, O., 2012. *Realized copula*. SFB 649 Discussion Papers.
- Giacomini, R., White, H., 2006. Tests of conditional predictive ability. *Econometrica* 74, 1545–1578.
- Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *J. Financ.* 48, 1779–1801.
- Gonçalves, S., White, H., 2004. Maximum likelihood and the bootstrap for nonlinear dynamic models. *J. Econ.* 119, 199–219.
- Gonçalves, S., Hounyo, U., Patton, A.J., Sheppard, K., 2013. Bootstrapping two-stage extremum estimators. *working paper*. Oxford-Man Institute of Quantitative Finance.
- Hafner, C.M., Manner, H., 2012. Dynamic stochastic copula models: estimation, inference and applications. *J. Appl. Econ.* 27, 269–295.
- Han, H., 2015. Asymptotic properties of GARCH-X processes. *J. Financ. Econ.* 13, 188–221.
- Hansen, B.E., 1994. Autoregressive conditional density estimation. *Int. Econ. Rev.* 35, 705–730.

- Hansen, P.R., Huang, Z., Shek, H.H., 2011. Realized GARCH: a joint model for returns and realized measures of volatility. *J. Appl. Econ.* 27, 877–906.
- Hansen, P.R., Lunde, A., Voev, V., 2014. Realized beta GARCH: a multivariate GARCH model with realized measures of volatility. *J. Appl. Econ.* 29, 774–799.
- Harvey, A.C., 2013. *Dynamic Models for Volatility and Heavy Tails*, Econometric Society Monograph 52. Cambridge University Press, Cambridge.
- Harvey, A.C., Sucarrat, G., 2014. EGARCH models with fat tails, skewness and leverage. *Comput. Stat. Data Anal.* 76, 320–338.
- Hautsch, N., Kyj, L.M., Oomen, R.C.A., 2012. A blocking and regularization approach to high dimensional realized covariance estimation. *J. Appl. Econ.* 27, 625–645.
- Joe, H., 2005. Asymptotic efficiency of the two-stage estimation method for copula-based models. *J. Multivar. Anal.* 94, 401–419.
- Jondeau, E., Rockinger, M., 2006. The copula-GARCH model of conditional dependencies: an international stock market application. *J. Int. Money Financ.* 25, 827–853.
- Jondeau, E., Rockinger, M., 2012. On the importance of time variability in higher moments for asset allocation. *J. Financ. Econ.* 10, 84–123.
- Li, S.Z., 2013. Continuous beta, discontinuous beta, and the cross-section of expected stock returns. working paper. Department of Economics, Duke University.
- Longin, F., Solnik, B., 2001. Extreme correlation of international equity markets. *J. Financ.* 56, 649–676.
- Müller, U.M., Dacorogna, R.D., Olsen, R., Pictet, O., Von Weizsäcker, J., 1997. Volatilities of different time resolutions—analysing the dynamics of market components. *J. Empir. Financ.* 4, 213–239.
- Noureldin, D., Shephard, N., Sheppard, K., 2012. Multivariate high-frequency-based volatility (HEAVY) models. *J. Appl. Econ.* 27, 907–933.
- Oh, D.H., Patton, A.J., 2013. Modelling dependence in high dimensions using factor copulas. working paper. Duke University.
- Okimoto, T., 2008. New evidence of asymmetric dependence structure in international equity markets. *J. Financ. Quant. Anal.* 43, 787–815.
- Patton, A.J., 2004. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. *J. Financ. Econ.* 2, 130–168.
- Patton, A.J., 2006a. Modelling asymmetric exchange rate dependence. *Int. Econ. Rev.* 47, 527–556.
- Patton, A.J., 2006b. Estimation of multivariate models for time series of possibly different lengths. *J. Appl. Econ.* 21, 147–173.
- Patton, A.J., 2013. Copula methods for forecasting multivariate time series. In: Elliott, G., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*. vol. 2. Elsevier, Oxford.
- Politis, D.N., Romano, J.P., 1994. The stationary bootstrap. *J. Am. Stat. Assoc.* 89, 1303–1313.
- Rivers, D., Vuong, Q., 2002. Model selection tests for nonlinear dynamic models. *Econ. J.* 5, 1–39.
- Rodríguez, J.C., 2007. Measuring financial contagion: a copula approach. *J. Empir. Financ.* 14 (3), 401–423.
- Shephard, N., Sheppard, K., 2010. Realising the future: forecasting with high frequency based volatility (HEAVY) models. *J. Appl. Econ.* 25, 197–231.
- Sklar, A., 1959. Fonctions de répartition à  $n$  dimensions et leurs marges. 8. Publications de l'Institut Statistique de l'Université de Paris, pp. 229–231.
- Sokolinskiy, O., van Dijk, D., 2011. Forecasting volatility with copula-based time series models. Tinbergen Institute Discussion Paper.
- White, H., 1994. Estimation, inference and specification analysis. *Econometric Society Monographs* No. 22. Cambridge University Press, Cambridge, U.K.