

Koopman, Siem Jan; Lit, Rutger; Lucas, André

Working Paper

Intraday Stock Price Dependence using Dynamic Discrete Copula Distributions

Tinbergen Institute Discussion Paper, No. 15-037/III/DSF90

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Koopman, Siem Jan; Lit, Rutger; Lucas, André (2015) : Intraday Stock Price Dependence using Dynamic Discrete Copula Distributions, Tinbergen Institute Discussion Paper, No. 15-037/III/DSF90, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/111716>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Duisenberg school of finance – Tinbergen Institute Discussion Paper

TI 15–037/III/ DSF90

Intraday Stock Price Dependence using Dynamic Discrete Copula Distributions

Siem Jan Koopman

Rutger Lit

André Lucas

*Faculty of Economics and Business Administration, VU University Amsterdam, and
Tinbergen Institute, the Netherlands.*

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and VU University Amsterdam.

More TI discussion papers can be downloaded at <http://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 1600

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Duisenberg school of finance is a collaboration of the Dutch financial sector and universities, with the ambition to support innovative research and offer top quality academic education in core areas of finance.

DSF research papers can be downloaded at: <http://www.dsf.nl/>

Duisenberg school of finance
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 525 8579

Intraday Stock Price Dependence Using Dynamic Discrete Copula Distributions

Siem Jan Koopman^(a,b,c), *Rutger Lit*^(a,b) and *André Lucas*^(a,b)

^(a)VU University Amsterdam, The Netherlands ^(b)Tinbergen Institute Amsterdam

^(c)CREATES, Aarhus University, Denmark

March 18, 2015

Abstract

We investigate the intraday dependence pattern between tick data of stock price changes using a new time-varying model for discrete copulas. We let parameters of both the marginal models and the copula vary over time using an observation driven autoregressive updating scheme based on the score of the conditional probability mass function with respect to the time-varying parameters. We apply the model to high-frequency stock price changes expressed as discrete tick-size multiples for four liquid U.S. financial stocks. Our modeling framework is based on Skellam densities for the marginals and a range of different copula functions. We find evidence of intraday time-variation in the dependence structure. After the opening and before the close of the stock market, dependence levels are lower. We attribute this finding to more idiosyncratic trading at these times. The introduction of score driven dynamics in the

dependence structure significantly increases the likelihood values of the time-varying copula model. By contrast, a fixed daily seasonal dependence pattern clearly fits the data less well.

Key words: time-varying copulas, dynamic discrete data, score driven models, Skellam distribution, dynamic dependence.

JEL Classification: C32, G11.

Acknowledgments: We thank István Barra for useful comments on an earlier version of this paper. Lit and Lucas acknowledge the financial support of the Dutch National Science Foundation (NWO, grant VICI453-09-005). Koopman acknowledges support from CREATES, the Center for Research in Econometric Analysis of Time Series (DNRF78) at Aarhus University, Denmark, funded by the Danish National Research Foundation.

1 Introduction

A key empirical finding from many analyses of intraday tick data is that stock price volatility is higher during opening hours than during the rest of the day; see, for example, Andersen and Bollerslev (1997) and Tsay (2005). Much less is known about the intraday pattern of the dependence structure between stock price changes. The dependence structures are of direct importance for intraday risk management, for example, when managing a book of multiple stocks that are traded repeatedly over the course of the day. This is common practice for many advanced players in today's markets. In our study we investigate the pattern of intraday dependence dynamics (beyond correlation structures) for a number of U.S. financial stocks observed at the tick-by-tick frequency. In earlier studies, lower-frequency data (5 minutes) is typically analyzed using standard correlation models; see, for example, Allez and Bouchaud (2011). We account for the discreteness of stock price changes in our analysis of high frequency tick-by-tick data and adopt a flexible dynamic copula framework for the modeling of the dependence structure.

There are at least two reasons why one can expect the dependence structure between price changes to vary within the day. First, news may have accumulated overnight. As many of the common macro announcement are scheduled during normal trading hours, a relatively higher percentage of idiosyncratic, firm-specific news is impounded in stock prices during the first minutes after the opening. Such increased information flows are known to affect intraday volatilities upwards immediately after the opening of the exchange; see for example Wood, McInish, and Ord (1985) and Admati and Pfleiderer (1988). Given the relatively higher fraction of idiosyncratic information after the opening, price changes are also likely to exhibit lower dependence during the first minutes after the opening compared

to the rest of the day. Towards the end of the day, idiosyncratic components may also play an important role and thus result in lower levels of dependence. In particular, we expect many players to unwind inventory positions that are built up over the course of the day in order to limit the (overnight) risk. The positions at the end of the day are therefore likely to contain relatively higher idiosyncratic components. Hence the expected dependence between price changes at the end of the trading day is lower.

We study intraday dynamics in price changes using tick-by-tick data observed at the one-second frequency over the year of 2012 for four financial stocks that are heavily traded. As the tick-size for our stocks is 1 dollar cent, prices as well as price changes move on a discrete grid. It is widely established that intraday price changes are subject to time-varying volatility and hence a time-varying marginal distribution. Many econometric challenges arise in the modeling of the dependence structure between discrete variables in case both the marginal distributions and the dependence structure are allowed to vary over time. The main methodological contribution of our current paper is that we provide a novel framework to address this issue in a way that is congruent with the empirical data. In particular, the dynamic parameters in our model, including stock return volatilities and dependence parameters, are updated using an observation driven, autoregressive updating function based on the score of the conditional observation probability mass function; for an introduction to the score driven approach, see Creal, Koopman, and Lucas (2011, 2013) and Harvey (2013), and for successful applications see, for example, De Lira Salvatierra and Patton (2013), Lucas, Schwaab, and Zhang (2014), Harvey and Luati (2014), and Creal, Schwaab, Koopman, and Lucas (2014). As is known from the literature, score driven models have three main advantages: (i) they possess information theoretic optimality properties, see Blasques, Koopman,

and Lucas (2015); (ii) they have similar forecasting performance as their parameter driven counterparts, even when the latter constitute the true data generating process, see Koopman, Lucas, and Scharth (2015); and (iii) as score driven models are observation driven rather than parameter driven in the classification of Cox (1981), the model's static parameters can be estimated in a straightforward way using maximum likelihood methods.

We adopt a dynamic Skellam distribution to model the tick-size price changes on the grid $\dots, -2, -1, 0, 1, 2, \dots$; see Irwin (1937) and Skellam (1946). The Skellam distribution has also been used to model price change series in other recent contributions; see Barndorff-Nielsen, Pollard, and Shephard (2012) for Skellam Lévy processes, and Shahtahmassebi (2011) for a Bayesian analysis of a parameter driven Skellam model. Rather than only having a dynamic generalization of the Skellam distribution for the marginal models, our main focus is on formulating a time-varying specification for the dependence structure in discrete data based on a copula framework. Discrete copulas and, in particular, dynamic discrete copulas pose a number of challenges. First, copulas for discrete marginals are not unique over the entire domain of the unit hypercube. Second, the copula density is no longer well-defined for discrete marginals, but is replaced by a copula probability mass function. Third, given the time-varying nature of the marginal distributions, the grid that defines the copula uniquely changes from one time period to the next. We address these issues by using a parametric copula specification that parsimoniously describes the copula surface. This function should cover grid points over which the copula at the current time point is uniquely defined but also grid points that may become relevant at future time points given the time-varying nature of the marginal distributions. We further allow for time-variation in the dependence structure by endowing the copula parameters with autoregressive dynamics that are a function of the

score of the copula probability mass function. In a Monte Carlo study, we show that our dynamic copula approach works well in uncovering the true parameter dynamics if the model is correctly specified; we can extract the path of the dynamic parameters with high precision. But also when the model is not correctly specified, we show that our approach accurately extracts the correct parameter path.

In our empirical study, we investigate the dependence in tick-by-tick price changes for a selection of four U.S. financial stocks traded on the NYSE. We present evidence that significant intraday time-variation in the dependence structure of these four stocks is present. The intraday dependence in all trading days of 2012 quickly increases during the first 30 minutes after the opening. After the first 30 minutes, the average intraday dependence remains relatively constant until, say, 15 minutes before the close when a sharp decrease in the dependence takes place for the six stock pairs considered in our analysis. As an alternative approach, we can specify the intraday pattern of the dependence as a fixed intraday seasonality pattern based on some flexible spline function. However, we show that our score driven time-varying copula methodology significantly outperforms the alternative spline-based approach. This indicates that time-variation in the intraday dependence is captured accurately by the score driven model and varies substantially between days. Furthermore it suggests that substantial day-to-day deviations from the average intraday pattern occur regularly.

The remainder of this paper is organized as follows. We introduce the model in Section 2. Section 3 presents simulation results on the model's adequacy. Our empirical analysis is presented in Section 4, while Section 5 contains the conclusions. The Appendix gathers a number of the more technical background expressions for the score-dynamics of the different marginals and copulas used in this paper.

2 Score driven dynamic discrete copula model

Consider a d -dimensional integer-valued vector $y_t = (y_{1,t}, \dots, y_{d,t})' \in \mathbb{Z}^d$ with time-varying conditional marginal distributions $F_i(y_{i,t} \mid \mathcal{F}_{t-1}; \theta_{i,t}^m)$ for $i = 1, \dots, d$ and $t = 1, \dots, T$, where $\theta_{i,t}^m$ is a time-varying parameter vector for the i th marginal distribution, and $\mathcal{F}_t = \{y_t, y_{t-1}, \dots\}$. The elements of y_t may for instance consist of counts, such as Poisson or binomial counts, or alternatively of changes in counts, such as the Skellam distributed discrete (tick-size) price changes in our empirical application in Section 4. The mean and variance of the Skellam distribution for stock i are then part of $\theta_{i,t}^m$. We characterize the dependence structure by a parametric conditional d -dimensional copula function

$$C \left[F_1(y_{1,t} \mid \mathcal{F}_{t-1}; \theta_{1,t}^m), \dots, F_d(y_{d,t} \mid \mathcal{F}_{t-1}; \theta_{d,t}^m) \mid \mathcal{F}_{t-1}; \theta_t^c \right], \quad (1)$$

where θ_t^c is the parameter vector defining the copula function C ; see Sklar (1959). The time-varying nature of θ_t^c allows us to study settings where the dependence structure changes over time. For example, in Section 4 we study the intraday dependence between discrete stock price changes. For notational simplicity, we suppress the dependence on the conditioning set \mathcal{F}_{t-1} and write the marginal distributions as

$$F_i := F_i(y_{i,t}; \theta_{i,t}^m) \equiv F_i(y_{i,t} \mid \mathcal{F}_{t-1}; \theta_{i,t}^m), \quad i = 1, \dots, d,$$

and copula function as $C \left[F_1(y_{1,t}; \theta_{1,t}^m), \dots, F_d(y_{d,t}; \theta_{d,t}^m); \theta_t^c \right]$. The dynamic specifications of the parameter vectors $\theta_{i,t}^m$ and θ_t^c are provided below. The dynamic conditional copula formulation presented in equation (1) is obtained from Patton (2002, 2006).

A discrete data analysis based on dynamic copulas faces several challenges; see also the review of Genest and Nešlehová (2007) on the use of static copulas for discrete marginals. For example, standard summary dependence measures such as Kendall's τ or Spearman's ρ are no longer guaranteed to lie in the $[-1, 1]$ interval and need to be used with caution in a discrete setting. In addition, we can no longer guarantee the uniqueness of the copula function in the standard Sklar (1959) representation of a distribution in terms of its marginal distributions and a copula function. The copula is only uniquely determined on the set $\text{Ran } F_1 \times \dots \times \text{Ran } F_d$, where $\text{Ran } F_i$ denotes the range of the cumulative distribution function (cdf) F_i , $i = 1, \dots, d$. This stands in sharp contrast with the case of continuous marginal distributions, where the copula function is unique over the entire unit hypercube $[0, 1]^d$.

Despite its non-uniqueness, discrete copulas can still be usefully applied in an empirical setting; see, for example, Zimmer and Trivedi (2006). At one extreme, we can model the value of the copula function at each point of its domain separately. This method can work in simple settings where the discrete data only takes a small number of different values; for example, in case of Bernoulli variables. This approach becomes infeasible, however, when the copula is defined over many different points as is the case in the empirical setting of Section 4. First, the price changes in our empirical example take values on \mathbb{Z} , and are therefore defined on (countably) infinitely many points. Second, and most importantly, the marginal distributions are time-varying. As a result, also the ranges $\text{Ran } F_i(\cdot; \theta_{i,t}^m)$ and therefore the domain over which C is uniquely identified are time-varying. Consequently, it is no longer feasible to estimate the value of the copula function over all points in the domain across all time periods, as there will be infinitely many of them. A possible solution is to model the copula in a parsimonious way. For example, we can use a parametric copula function defined

over the entire $[0, 1]^d$ space even though uniqueness is only guaranteed over a set of discrete points. This is the approach we will adopt in our analyses below.

The dynamic specifications for $\theta_{i,t}^m$ and θ_t^c in (1) are based on the score driven approach of Creal et al. (2011, 2013) and Harvey (2013). We collect the time-varying parameters in $\theta'_t = (\theta_{1,t}^{m'}, \dots, \theta_{d,t}^{m'}, \theta_t^{c'})$. The score driven model represents a class of models in which the update of θ_t over time is formulated as a function of past data y_{t-1}, y_{t-2}, \dots and past realized parameter values $\theta_{t-1}, \theta_{t-2}, \dots$. At time t we can write the update function as

$$\theta_{t+1} = \theta_{t+1}(y_t, y_{t-1}, \dots, \theta_t, \theta_{t-1}, \dots; \psi),$$

where ψ is an unknown parameter vector that contains the update coefficients and the remaining static parameters of the marginal distributions and the copula function. It follows that θ_t is \mathcal{F}_{t-1} -measurable and the approach is observation driven in the classification of Cox (1981). The estimation of the static parameter vector ψ is typically carried out by the method of maximum likelihood in a straightforward manner. A score driven model updates θ_t in the direction of the steepest increase of the log conditional probability mass function (pmf) at time t given the past information set \mathcal{F}_{t-1} . Updating θ_t in this way possesses information theoretic optimality properties as shown by Blasques et al. (2015).

Let $p(y_t | \mathcal{F}_{t-1}; \theta_t)$ denote the pmf of y_t , which we again write in short-hand notation as $p(y_t; \theta_t)$, suppressing its dependence on the parameter vector ψ . Using the so called ‘inclusion-exclusion’ formula, we obtain from equation (1) that

$$p(y_t; \theta_t) = \sum_{j_1=0,1} \dots \sum_{j_d=0,1} (-1)^{j_1+\dots+j_d} \times C [F_1(y_{1t} - j_1; \theta_{1,t}^m), \dots, F_d(y_{dt} - j_d; \theta_{d,t}^m); \theta_t^c]. \quad (2)$$

For instance, for the bivariate case ($d = 2$), the pmf becomes

$$p(y_t; \theta_t) = C [F_1(y_{1,t}; \theta_{1,t}^m), F_2(y_{2,t}; \theta_{2,t}^m); \theta_t^c] - C [F_1(y_{1,t} - 1; \theta_{1,t}^m), F_2(y_{2,t}; \theta_{2,t}^m); \theta_t^c] - \quad (3)$$

$$C [F_1(y_{1,t}; \theta_{1,t}^m), F_2(y_{2,t} - 1; \theta_{2,t}^m); \theta_t^c] + C [F_1(y_{1,t} - 1; \theta_{1,t}^m), F_2(y_{2,t} - 1; \theta_{2,t}^m); \theta_t^c],$$

where the evaluation of equation (2) requires 2^d evaluations of the copula function, for any t , and is feasible for low values of d as in (3). The evaluation of (2) clearly becomes more challenging for larger values of d ; see, for example, Panagiotelis, Czado, and Joe (2012). The score-based update function for θ_t takes the form

$$\theta_{t+1} = \omega + A\nabla_t + B\theta_t, \quad \nabla_t = \frac{\partial \log p(y_t; \theta_t)}{\partial \theta_t}, \quad (4)$$

where ∇_t is the score vector of the (predictive) density $p(y_t; \theta_t)$ in (2), ω is a vector of constants, and A and B are fixed coefficient matrices. These coefficients are functions of the parameter vector ψ that also includes the unknown parameters of the marginal distributions F_i and the copula function C in (2). Since $p(y_t; \theta_t)$ relies on ψ , it follows that ∇_t is also a function of ψ . The derivative ∇_t in (4) is straightforward to obtain because the pmf is typically differentiable in the time-varying parameters θ_t . The updating equation (4) corresponds to the unit scaling option of Creal et al. (2013) and can be generalized in different ways; for example, by adding more lagged values of θ_t and ∇_t .

The time-varying parameter vector θ_t is initialized at θ_1 , which we include in the parameter vector ψ . In the case of a bivariate copula, the individual components of θ_t consist of two marginal parameter vectors and one copula dependence parameter. To introduce further parsimony, we assume diagonal matrices for A and B , such that each element

of θ_t is updated by its own score function only. The static parameter vector ψ becomes $\psi = \{\theta'_1, \omega', \text{diag}(A)', \text{diag}(B)'\}$, where $\text{diag}(M)$ denotes the vector of the diagonal elements of any matrix M . A more parsimonious specification is obtained by having θ_1 as to the unconditional mean of θ_t , that is $\theta_1 = \omega \oslash (1 - \text{diag}(B))$ where \oslash denotes the Hadamard division (pointwise division).

The score function for $\nabla'_t = (\nabla'_{1,t}, \dots, \nabla'_{d,t}, \nabla'^c_t)$ has an analytical solution that is given by the elements

$$\nabla^m_{k,t} = \frac{\partial \log p(y_t; \theta_t)}{\partial \theta^m_{k,t}} = \frac{\sum_{j_1=0,1} \dots \sum_{j_d=0,1} (-1)^{j_1+\dots+j_d} \frac{\partial C(u_{1,t}, \dots, u_{d,t}; \theta_t^c)}{\partial u_{k,t}} \cdot \frac{u_{k,t}}{\partial \theta^m_{k,t}}}{\sum_{j_1=0,1} \dots \sum_{j_d=0,1} (-1)^{j_1+\dots+j_d} C(u_{1,t}, \dots, u_{d,t}; \theta_t^c)}, \quad (5)$$

$$\nabla^c_t = \frac{\partial \log p(y_t; \theta_t)}{\partial \theta_t^c} = \frac{\sum_{j_1=0,1} \dots \sum_{j_d=0,1} (-1)^{j_1+\dots+j_d} \partial C(u_{1,t}, \dots, u_{d,t}; \theta_t^c) / \partial \theta_t^c}{\sum_{j_1=0,1} \dots \sum_{j_d=0,1} (-1)^{j_1+\dots+j_d} C(u_{1,t}, \dots, u_{d,t}; \theta_t^c)}, \quad (6)$$

for $k = 1, \dots, d$, and with $u_{i,t} = F_i(y_{i,t} - j_i; \theta^m_{i,t})$, for $j_i \in \{0, 1\}$, $i = 1, \dots, d$, and $t = 1, \dots, T$.

The denominators in (5) and (6) are equal to the pmf as given in (2). In case of the Gaussian copula as well as the commonly encountered copulas from the Archimedean class, analytical expressions for ∇_t are available. We refer to Appendix A for further details and to Schepsmeier and Stöber (2014) for expressions for a range of derivatives of bivariate copulas.

Given that θ_t is \mathcal{F}_{t-1} -measurable and our model specification is observation driven in the classification of Cox (1981), we obtain the likelihood function in closed form by a standard prediction error decomposition,

$$L(y; \psi) = \sum_{t=1}^T \log p(y_t; \theta_t), \quad (7)$$

with $y = (y_1, \dots, y_T)$. We define the corresponding maximum likelihood estimator (MLE)

of ψ as $\hat{\psi} = \arg \max_{\psi} L(y; \psi)$. In practice we obtain the MLE of ψ via the direct numerical maximization of $L(y; \psi)$ with respect to ψ .

Example: Frank copula with Skellam marginals

As a concrete example, consider the bivariate Frank copula with Skellam marginals. This combination of copula and marginals is used to perform the simulation study in Section 3.

The Frank copula is a symmetric copula given by

$$C_{\text{Fr}}(u_{1,t}, u_{2,t}; \theta_t^c) = \frac{1}{\theta_t^c} \log \left[1 + \frac{(\exp(-\theta_t^c u_{1,t}) - 1)(\exp(-\theta_t^c u_{2,t}) - 1)}{\exp(-\theta_t^c) - 1} \right], \quad (8)$$

with $\theta_t^c \in \mathbb{R} \setminus \{0\}$; see Frank (1979) and Nelsen (2006). When $\theta_t^c \rightarrow 0$, the Frank copula converges to the independence copula $C_{\text{Fr}}(u_{1,t}, u_{2,t}; 0) = u_{1,t}u_{2,t}$.

A Skellam pmf with location parameter μ_t and scale parameter σ_t^2 is given by

$$\Pr(Y_t = y_t; \mu_t, \sigma_t^2) = \exp(-\sigma_t^2) \left(\frac{\mu_t + \sigma_t^2}{\sigma_t^2 - \mu_t} \right)^{y_t/2} I_{|y_t|} \left(\sqrt{\sigma_t^4 - \mu_t^2} \right), \quad (9)$$

with $y_t \in \mathbb{Z}$ and where $I_{|y_t|}(\cdot)$ is the modified Bessel function of order $|y_t|$. The shape of the Skellam distribution depends on μ_t and σ_t^2 and is symmetric for $\mu_t = 0$, skewed right when $\mu_t > 0$, and left-skewed for $\mu_t < 0$. The excess kurtosis of the Skellam pmf is $1/\sigma_t^2$, and it has the Gaussian distribution as a limiting case. The Skellam distribution was originally derived by Irwin (1937) and Skellam (1946) as a distribution for the difference between two Poisson variables. Our parameterization in equation (9) is a reparameterization of the original version and can be transformed back by substituting $\mu_t = \lambda_{1,t} - \lambda_{2,t}$ and $\sigma_t^2 = \lambda_{1,t} + \lambda_{2,t}$ in (9), where $\lambda_{1,t}$ and $\lambda_{2,t}$ are the means of the underlying Poisson distributions; see also Alzaid and Omair

(2010).

The mean μ_t and variance σ_t^2 in the full model in equation (9) are time-varying. In our application of Section 4, however, the mean turns out to be insignificantly different from zero and not time-varying, while the variance remains time-varying. In this case, equation (9) simplifies to

$$\Pr(Y_t = y_t; \sigma_t^2) = \exp(-\sigma_t^2) I_{|y_t|}(\sigma_t^2). \quad (10)$$

We then obtain a Frank copula function $C(u_{1t}, u_{2t}; \theta_t^c)$ with $\theta_{i,t}^m = \sigma_{i,t}^2$ and

$$u_{i,t} = \Pr(Y_{i,t} \leq k; \sigma_{i,t}^2) = \exp(-\sigma_{i,t}^2) \sum_{j=-\infty}^k I_{|j|}(\sigma_{i,t}^2), \quad i = 1, 2. \quad (11)$$

3 Simulation study

To investigate the properties of our model in a controlled setting, we carry out two simulation studies. In our first study, we assume that the score driven model of equations (2) and (4) is the true data generating process and verify the finite sample behaviour of the maximum likelihood estimates for the parameter vector ψ . In the second study, we consider a misspecified model setting. We assume that the marginal parameters and the dependence parameter come from some exogenous dynamic patterns that do not rely on the score function. We then verify to what extent the score driven framework is able to recover the true underlying dynamics of the time-varying parameter vector θ_t . In both simulation studies, we focus on a positive dependence between two series, that is $\theta_t^c \in \mathbb{R}^+$. We specify $\bar{\theta}_t^c = \log(\theta_t^c)$ as the time-varying parameter rather than θ_t^c itself. We adopt the same specification for the variance of the Skellam distribution, that is $\bar{\theta}_{i,t}^m = \log(\theta_{i,t}^m) = \log(\sigma_{i,t}^2)$ and $\bar{\theta}_{i,t}^m$ varies

over time. The score function ∇_t in (5) and (6) adapt to this reparameterization into $\bar{\theta}_t$ by pre-multiplying ∇_t by $\partial\theta'_t/\partial\bar{\theta}_t$. This reparameterization yields an estimation procedure that is numerically more stable. In both simulation studies, the observation series are simulated from a bivariate Frank copula with Skellam marginals as discussed in Section 2.

3.1 Estimating parameters when model is correctly specified

We simulate $S = 500$ series of correlated Skellam observations. The length of the sample is set to $T \in \{250, 1000, 3000\}$. To generate the data, we apply the algorithm of Nelsen (2006, p.41) using a numerical inverse cdf of the Skellam distribution. For the log-transforms of the dynamic parameters θ_t , we consider equation (4). The estimates of the parameter vector ψ are obtained via the numerical maximization of the loglikelihood function (7) using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm.

Table 1 presents the results. The method of maximum likelihood is able to estimate the parameters in ψ accurately, even for the small sample size $T = 250$. For $T \in \{1000, 3000\}$, the maximum likelihood estimates for the unconditional mean $\bar{\theta}_1$ and the score loadings $(a_1, a_2, a_3)'$ in the updating equations are virtually equal to the corresponding true parameters. In the case of $T = 250$, the persistence parameters b_1, b_2, b_3 are underestimated, which matches small sample biases encountered in similar studies for standard linear time series models. The biases disappear for larger sample sizes. In the case of $T = 3000$, the b_1, b_2, b_3 parameters are estimated close to their true values. Finally, we can conclude from the average computing times $t(s)$ for estimation, also reported in Table 1, that the score driven methodology applied to the bivariate Frank copula is quite fast. The computing times for parameter estimation ranges from less than 18 seconds, on average, for $T = 250$, to approx-

Table 1: Simulation results under correct model specification

This table reports simulation averages of maximum likelihood estimates of the static parameters for the Skellam-Frank score driven model of Section 2. The results use $S = 500$ replications of time series of length $T \in \{250, 1000, 3000\}$. The intercepts ω in (4) are set to $(I - B)\bar{\theta}_1 = \omega$, such that $\bar{\theta}_1$ is the unconditional mean of $\bar{\theta}_t$, where $\bar{\theta}_t$ contains the logs of the elements of θ_t . The matrices A and B are diagonal with elements (a_1, a_2, a_3) and (b_1, b_2, b_3) , respectively. Standard deviations of the estimates over the Monte Carlo simulations are in parentheses. The column $t(s)$ denotes the average computation time (in seconds) for finding the maximum of the log likelihood function. Computations are carried out on a i7-2600, 3.40 GHz desktop PC using four cores.

T	$\bar{\theta}_{1,1}$	$\bar{\theta}_{2,1}$	$\bar{\theta}_{3,1}$	a_1	a_2	a_3	b_1	b_2	b_3	$t(s)$
true	1.00	1.00	2.00	0.10	0.05	0.10	0.90	0.95	0.98	-
250	1.00 (0.14)	1.01 (0.14)	2.03 (0.28)	0.11 (0.08)	0.06 (0.06)	0.11 (0.08)	0.74 (0.28)	0.77 (0.27)	0.87 (0.21)	17.76 -
1000	1.00 (0.07)	1.00 (0.07)	2.00 (0.15)	0.10 (0.03)	0.05 (0.03)	0.10 (0.03)	0.87 (0.09)	0.91 (0.09)	0.97 (0.03)	59.98 -
3000	1.00 (0.04)	1.00 (0.04)	2.00 (0.06)	0.10 (0.02)	0.05 (0.02)	0.10 (0.02)	0.89 (0.04)	0.94 (0.08)	0.98 (0.01)	108.30 -

imately 108 seconds, for $T = 3000$. The computations are carried out by an i7-2600, 3.40 GHz desktop PC using four cores.

3.2 Estimating time-varying paths when model is misspecified

Next we deviate from the assumption that the score driven model (2) and (4) is the data generation process. In our second Monte Carlo study, the time-varying Skellam variances and the time-varying dependence parameter are generated as sinusoidal patterns with different periods and amplitudes. We investigate to what extent our misspecified score driven framework is able to identify these time-varying patterns. We generate $S = 500$ time series of length $T \in \{250, 1000, 3000\}$ and estimate the parameters in vector ψ by the method of maximum likelihood.

In Figure 1, we present the true time-varying parameters σ_{1t}^2 , σ_{2t}^2 and θ_t^c together with their estimated counterparts $\hat{\sigma}_{1t}^2$, $\hat{\sigma}_{2t}^2$ and $\hat{\theta}_t^c$, respectively. The results for $T = 1000$ and $T = 3000$ show that the score driven model is able to capture the true paths of the time-




Figure 1: Simulation results under mis-specification

The figure presents the point wise Monte Carlo averages (solid fat) over 500 replications of the Skellam variances $\sigma_{1,t}^2$ and $\sigma_{2,t}^2$, and of the Frank copula parameter θ_t^c . All three parameters are parametrized in log form in the score driven specification. Each panel also contains the true time varying parameter (solid thin) and a band of two times the point wise standard deviations (dotted). From left to right the panels show time series length of $T = 250, 1000, 3000$, respectively.

varying parameters accurately, despite its misspecification. Only in the case of the small sample size $T = 250$ and the rapidly changing parameter paths for $\sigma_{1,t}^2$, the estimates are less accurate. In our empirical study in Section 4, we have more than 40,000 observations per month. Hence we expect the score driven model to perform sufficiently accurately in our empirical study.

4 Dependence between discrete price changes

The dependence measures between price changes of individual stocks or assets are the key ingredients in, for instance, portfolio risk management. In our empirical study, we establish the intraday dependence structure in high-frequency price changes. Whereas most studies concentrate on the intraday dynamics of volatility, our study is, to the best of our knowledge, the first to concentrate on the intraday dynamics of the dependence structure using a copula approach in a tick-by-tick data analysis.

The data sets consist of price changes of stocks traded at the New York Stock Exchange (NYSE). The resulting series consist of discrete, integer multiples of the tick-size of one dollar cent. The observations take values in \mathbb{Z} . We model the discrete tick-size price changes instead of the returns. Münnix, Schäfer, and Guhr (2010) argue that the discrete nature of the price grid affects the empirical distribution of returns severely. This distribution concentrates around the actual tick-sizes, is severely multi-modal and, consequently, highly non-Gaussian.

Several models for data in \mathbb{Z} are available in the literature. For example, the model of Rydberg and Shephard (2003) decomposes stock price movements into activity, direction of moves, and size of the moves. Freeland (2010), Alzaid and Omair (2014) and Andersson and Karlis (2014) extend the integer autoregressive (INAR) model for \mathbb{N} variables to the case of \mathbb{Z} variables. They propose the Skellam distribution and use static Skellam parameters. Barndorff-Nielsen et al. (2012) analyze Skellam Lévy processes for intraday price changes. Shahtahmassebi (2011) present a Bayesian analyses based on a Skellam model for \mathbb{Z} variables. The dynamic Skellam model for time series observations in \mathbb{Z} is developed by Koopman, Lit, and Lucas (2014) based a non-Gaussian state space analysis. In our current framework, we

adopt the Skellam distribution for the marginals and allow the corresponding parameters to vary over time using the score driven model of Section 2.

Although the Skellam distribution is an important ingredient of our analysis, our main focus is on the dependence structure as this feature has received much less attention in other related studies so far. Our analysis proceeds in two steps. First, we study the dependence characteristics between price changes of four major NYSE listed financials over a period of one trading month for a variety of copulas. We consider Bank of America Corporation (BAC), Citigroup Inc. (C), JPMorgan Chase & Co. (JPM) and Wells Fargo & Company (WFC). Based on our initial findings for these four stocks, we select the best copula for the second part of our analysis: an analysis of the intraday dependence dynamics over the long time span of an entire calendar year.

4.1 Data description

We first analyze intraday stock prices obtained from the TAQ database for April 2012. We clean the high-frequency data by following the standard procedures described in Brownlees and Gallo (2006) and Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008) for TAQ data. This database has a time stamp precision of 1 second so that for many seconds we obtain a number of transactions with the same time stamp. It is common practice to merge these transactions and to replace them by the median price rounded to the nearest tick.

Figure 2 presents the intraday tick price changes for our four selected stocks. We present the results for a typical trading day, April 23, 2012. We find that more trades with relatively large price changes occur at the beginning of the day and a quiet period with small or no price changes takes place during lunch-time. Appendix B contains additional descriptive

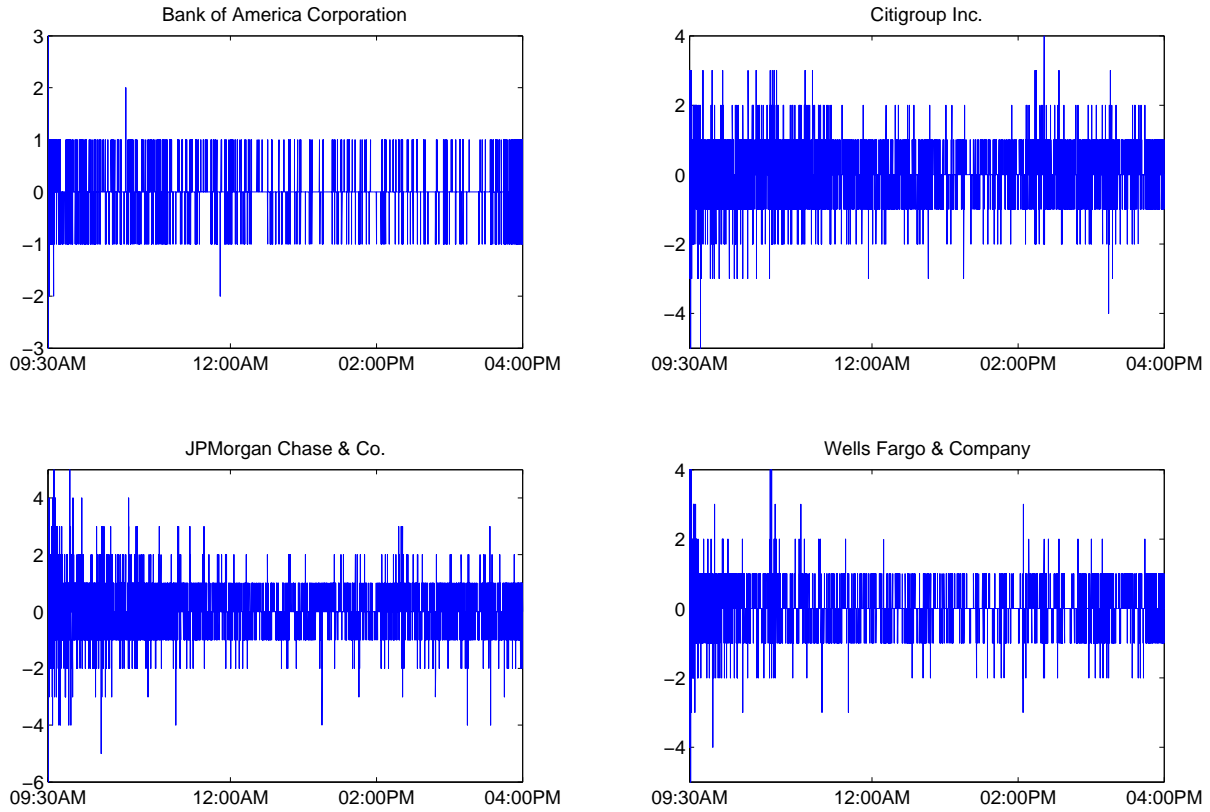


Figure 2: Tick price changes of Bank of America Corporation (BAC), Citigroup Inc. (C), JPMorgan Chase & Co. (JPM) and Wells Fargo & Company (WFC) on April 23, 2012.

plots of the data.

Table 2 presents descriptive statistics of the tick-size price changes. We find that Citigroup and JPMorgan are the most liquid stocks in terms of the number of trades, followed by Wells Fargo and Bank of America. The absolute price level has a clear impact on the tick-size volatility: the minimum and maximum tick-size changes as well as the tick-size variance are substantially lower for Bank of America than for the other three institutions. We account for this effect by using different parameters in the marginal models for each stock.

Table 2: Descriptive statistics

The table reports company name, ticker symbol (Code), the number of trades (#Trades), the opening price at 9:30 AM of the first trading day in the sample (P_{open}), the closing price at 16:00 PM of the last trading day in the sample (P_{close}), the largest up-tick (\uparrow) measured in multiples of the tick-size, the largest down-tick (\downarrow), the variance (Var) and mean (Mean) of the tick-size price changes, and the percentage of 0-trades (%0).

Company	Code	#Trades	P_{open}	P_{close}	\uparrow	\downarrow	Var	Mean	%0
Apr 2012									
Bank of America Corp.	BAC	41,640	9.53	8.09	7	-6	0.242	-0.004	76.84
Citigroup Inc.	C	93,872	36.34	33.03	8	-11	0.753	-0.004	55.93
JPMorgan Chase & Co.	JPM	90,936	45.79	42.95	8	-8	0.747	-0.001	54.12
Wells Fargo & Company	WFC	64,529	33.85	33.40	8	-9	0.575	0.000	60.77
Jan 2012 - Dec 2012									
Bank of America Corp.	BAC	560,102	5.76	11.62	7	-6	0.232	0.001	77.30
Citigroup Inc.	C	1,084,943	27.20	39.59	11	-15	0.663	0.001	57.87
JPMorgan Chase & Co.	JPM	1,029,844	34.10	44.00	20	-16	0.725	0.001	55.30
Wells Fargo & Company	WFC	766,712	28.00	34.22	13	-14	0.510	0.001	63.30

4.2 Missing values

Our observation driven model is formulated for a time frequency in seconds. Since we do not observe a trade for every second during the trading day, we encounter many missing observations. We distinguish four situations that can occur at second t during a day.

Situation 1: At time t , stock 1 trades while stock 2 does not trade so that the price change for series 2 is missing at time t . The copula dependence parameter cannot be updated as we require two observations to update the parameter related to instantaneous dependence. Furthermore, the marginal variance $\sigma_{1,t}^2$ cannot be updated by taking derivatives from the copula mass function in (6) since both observations from series 1 and 2 are needed as input. In this case, variance $\sigma_{1,t}^2$ is updated by only using the score of the marginal Skellam log pmf in (10). No score updating takes place for $\sigma_{2,t}^2$ and θ_t^c and hence these parameters mean revert by setting $\nabla_{2,t}^m$ and ∇_t^c to zero in (5) and (6). The contribution to the likelihood at time t is given by the logarithm of the pmf in (10) with $\sigma_{1,t}^2$ and $y_{1,t}$ as input.

Situation 2: At time t , stock 1 does not trade while stock 2 is traded. This is the converse of Situation 1 and has an analogous solution.

Situation 3: At time t , both stocks trade. The whole time-varying parameter vector θ_t is updated according to (4), where the score is obtained by taking derivatives from the copula mass function in equation (6). The contribution to the likelihood at time t is made by the logarithm of the copula mass function in (5).

Situation 4: At time t neither stock 1 nor stock 2 trades. In this case, none of the parameters is updated and there is no contribution to the likelihood.

For the purpose of estimating a dependence parameter, situation 3 has clearly the most impact. We therefore present in Figure 3 the number of simultaneous trades per half hour of the trading day. The numbers are averaged over all 250 trading days of the year 2012. Figure 3 reveals more joint trades at the beginning and the end of the day compared to the middle of the day. We may therefore expect more information in the data on the dependence parameter θ_t^c at the start and at the end of the day. Figure B.2 in the Appendix reveals that the same increased trading intensity at the start and end of the trading day occurs for other stock combinations as well.

4.3 Copula selection

We take the independence copula as a benchmark and verify for a range of copulas whether they improve the model fit. The model fits are compared by means of the Bayesian Information Criterion (BIC) for both static dependence θ^c and time-varying dependence θ_t^c . For all models considered, the marginal parameters of the Skellam distribution, σ_{1t}^2 and σ_{2t}^2 are allowed to vary over time. Our selection of copula functions includes the independence cop-




Figure 3: The figure displays the number of simultaneous trades per half hour of the trading day as well as the the number of trades if only CitiGroup or JPMorgan trade. The numbers are averaged over all 250 trading days of the year 2012.

ula (Indep), the symmetric Ali-Mikhail-Haq (AMH), Frank, and Gaussian copulas, and the asymmetric Clayton (lower tail dependence), Gumbel (upper tail dependence), Joe (upper tail dependence), and Symmetrized Joe Clayton (SJC) copula (upper and lower tail dependence); see, for instance, Nelsen (2006) and Patton (2006) for the functional specifications of the these copulas.

For each day, the vector of time-varying parameters θ_t is initialized at θ_1 which is estimated as part of the vector of static parameters ψ . Table 3 presents the model selection results for all trading days in April 2012. Entries indicate the number of points by which the corresponding copula outperforms the BIC of the independence copula. Higher entries are thus preferred.

From Table 3 we learn that dynamic dependence is preferred over static dependence for

Table 3: BIC Improvements Compared to the Independence Copula over April 2012

The table reports the difference in Bayesian information criterion for the independence copula vis-à-vis the Gaussian, Ali-Mikhail-Haq (AMH), Frank, Clayton, Gumbel, Joe, and Symmetrized-Joe-Clayton (SJC) copulas: $DB_s = BIC_s^{Indep} - BIC_s^\tau$, with $s \in \{st, dy\}$ and where τ denotes the copula under consideration. The data are tick price change series for Bank of America (BAC), Citigroup (C), JPMorgan (JPM), and Wells Fargo (WFC), observed during April 2012. $\#_{st}$ and $\#_{dy}$ denote the number of parameters in the case of static and dynamic dependence model, respectively. The marginal Skellam distributions are always dynamic. The largest difference in BIC compared to the independence copula is boxed for static dependence and highlighted in gray for dynamic dependence.

Copula	$\#_{st}$	$\#_{dy}$	BAC/C		BAC/JPM		BAC/WFC	
			DB_{st}	DB_{dy}	DB_{st}	DB_{dy}	DB_{st}	DB_{dy}
Gaussian	9	12	367.72	492.44	454.74	430.19	309.00	320.44
AMH	9	12	348.47	456.91	428.41	407.31	288.07	294.79
Frank	9	12	338.51	465.47	416.51	398.64	283.25	283.45
Clayton	9	12	284.95	398.11	368.14	337.74	257.40	255.70
Gumbel	9	12	253.44	369.16	322.37	301.15	222.86	216.10
Joe	9	12	151.58	260.74	190.49	168.69	136.11	119.80
SJC	10	16	268.74	407.11	350.56	353.04	262.81	233.91

Copula	$\#_{st}$	$\#_{dy}$	C/JPM		C/WFC		JPM/WFC	
			DB_{st}	DB_{dy}	DB_{st}	DB_{dy}	DB_{st}	DB_{dy}
Gaussian	9	12	4545.80	4793.97	3593.87	3771.53	3929.64	4108.64
AMH	9	12	4264.01	4421.15	3336.53	3441.86	3660.14	3770.96
Frank	9	12	4469.20	4694.16	3593.59	3751.12	3895.69	4029.04
Clayton	9	12	3447.26	3680.29	2653.37	2836.15	3027.43	3200.86
Gumbel	9	12	3868.24	4083.63	3174.97	3311.71	3468.41	3612.73
Joe	9	12	2693.09	2889.16	2294.63	2410.84	2474.92	2621.57
SJC	10	16	4227.30	4413.54	3411.32	3500.54	3733.42	3835.32

five out of the six pairwise data sets based on the BIC. The symmetric copulas, Gaussian, AMH, and Frank, are generally preferred over the asymmetric ones. It confirms the somewhat symmetric patterns in the pairwise up and down tick movements encountered in the scatter plots of the data; see the Appendix for more evidence. The main conclusion of our first analysis is clear: both for static dependence as well as for dynamic dependence, the Gaussian copula fits the data best for all stock pairs. The Gaussian copula exhibits zero tail dependence. Given that copula functions with upper and/or lower tail dependence, such as Clayton, Gumbel, Joe, and Symmetrized Joe Clayton copulas, fit the data less well, we infer that tail dependence is not a dominant feature in tick-size price change series.

4.4 Full year results

In this section we extend our analysis over the entire year 2012. Descriptive statistics for this larger time span were given in Table 2. The characteristics of the data for all trading days in 2012 are broadly similar to those for the trading days in April 2012 only. Therefore, we use the Gaussian copula as our best fitting specification based on our preliminary analysis in Section 4.3. For the Gaussian copula correlation parameter ρ_t , we use the time-varying parameter θ_t^c , with

$$\rho_t = \theta_t^c / \sqrt{1 + (\theta_t^c)^2}. \quad (12)$$

This parameterization of ρ_t via θ_t^c ensures that the copula dependence parameter is always within the appropriate interval, that is $\rho_t \in (-1, 1)$. The likelihood for a full year of tick price changes is maximized in approximately 4 to 15 hours (depending on starting values and data sets) on a i7-2600, 3.40 GHz desktop PC using four cores. The parameter estimates are presented in the Appendix.

We are mainly interested in the intraday pattern of the copula dependence parameter. Therefore, we first compute the point-wise sample mean of the intraday path of the copula dependence parameter over all 250 trading days of 2012. Figure 4 presents these sample means together with the confidence bands based on the corresponding sample variances. We compare our estimates of the intraday Gaussian dependence with an adjusted version of Spearman's rank correlation coefficient. This non-parametric rank correlation measure is computed for a rolling window of 600 seconds using only the observations with simultaneous trades. The observations are ordered while ties in ranks are corrected in the usual way by averaging the ranks. The resulting ranks are divided by 1 plus the number of observations.

Finally we transform the ranks through the inverse normal cdf. The Pearson correlation between these transformed ranks are presented in Figure 4.

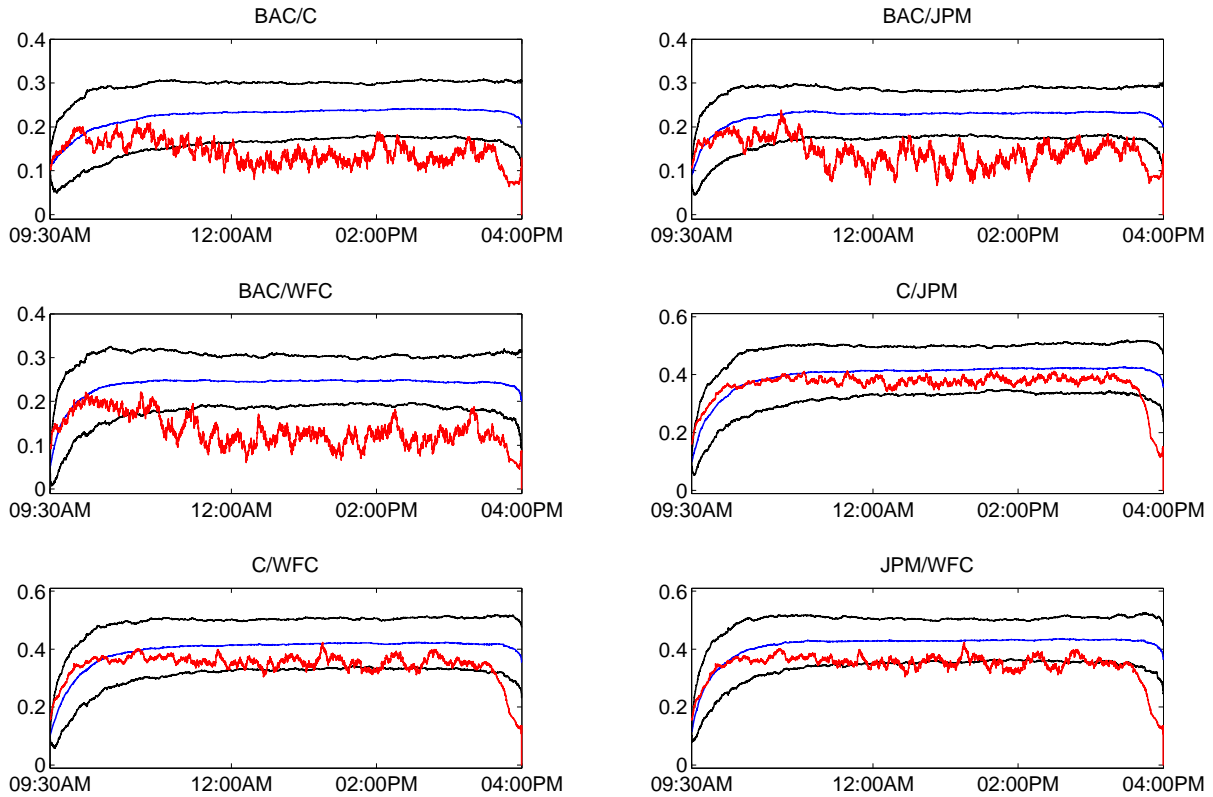


Figure 4: Point wise mean copula dependence intraday patterns over the 250 trading days in 2012 based on the Gaussian copula with Skellam marginals (smooth line). The smooth bounds are based on two sample standard deviations. The noisy series is the adjusted version of Spearman’s non-parametric rank based estimator.

We find that the dependence between tick-size price changes exhibits a clear daily pattern across all stocks. We see that the trading day starts with a relatively small positive dependence level. But within the first hour of trading, the average dependence increases to a higher level where it remains throughout the trading day. Only during the last 15 minutes of trading, the dependence drops abruptly to a somewhat lower value. This pattern is found across all stock pairs. The point wise sample mean of the non-parametric rank-based dependence measure is much less smooth than our model-based measure. We also observe that the rank-based measure is significantly lower than the score driven dependence implied

by the Gaussian copula, which is partly due to the problems with rank-based statistics such as Spearman's rho for discrete data. We may conclude that our copula framework uses the data more efficiently. We emphasize that the estimated dependence patterns are not due to a lack of observation pairs at the end of the day. By contrast, Figure 3 shows that the number of joint observations is relatively higher at the start and at the end of the day.

The empirical intraday pattern for the dependence parameter can be expected given the flow of information over the 24 hour cycle. Throughout the trading day, information becomes available and can immediately be processed and impounded into stock prices due to active trading. The accumulated overnight information can only be impounded at the opening of the trading day. While most of the common macro announcements are made during the trading day, most major firm-specific information is revealed after the active trading hours. The information available at the opening may therefore have a relatively larger idiosyncratic component. This causes the lower dependence level at the start of the trading day. Interestingly, the lower level in dependence at the opening mirrors the typically higher levels of intraday volatility during the opening.

It is likely that the short, sudden drop in dependence at the end of the day is related to the unwinding of open positions by market participants built up over the trading day. Such unwinding may be spurred by the need to satisfy overnight risk constraints. Hence it comprises a relatively larger idiosyncratic component and therefore also results in a decrease in the dependence parameter.

4.5 Comparison with intraday spline

The smooth patterns for the estimated intraday dependence across all stock pairs may prompt the question whether we can alternatively consider a smooth function to capture intraday dependence. We therefore compare our score driven updating function for the copula dependence parameter ρ_t with a basic cubic spline function to account for the intraday seasonal pattern. The width of the confidence bands around the sample averages of the intraday dependence estimates presented in Figure 4 indicate that there exists considerable variation in the dependence parameter across the 250 trading days of 2012. For example, according to the 95% confidence bands the dependence parameter can vary between 0.1 and 0.3 at Noon.

To investigate whether a spline suffices to model the dependence parameter, we keep our score driven approach for the marginal Skellam distributions, but model the copula dependence path by a cubic spline regression function as proposed by Poirier (1973). For the cubic spline regression, we specify the copula parameter by $\theta_t^c = \kappa' W_t$ where κ is a $q \times 1$ vector of parameters associated with the location of the q spline knots, and W_t is the t -th column of the weight matrix W as constructed in Poirier (1973). We have considered different numbers of knots and different locations for the knots in order to control for the possible sensitivity of the approach. The elements of κ become part of the parameter vector ψ and are jointly estimated by the method of maximum likelihood.

Table 4 presents the results for a range of different models. We report the loglikelihood gains and BIC reductions (in parentheses) for the considered spline model compared to the dynamic score driven Skellam-Gaussian copula model. For almost all combinations, the loglikelihood gains are reported to be negative, indicating that the score driven model

Table 4: Model comparison: intraday dependence spline versus score driven dynamics

The entries reflect the gain in log likelihood points (and improvements in BIC in parentheses) of the spline model compared to the dynamic score driven Skellam-Gaussian copula model. The time points between braces are the positions of the spline knots. $\#_{\psi}$ denotes the number of estimated parameters, i.e., the dimension of ψ . Stocks are Bank of America (BAC), Citi (C), JPMorgan (JPM), and Wells Fargo (WFC).

Model description	$\#_{\psi}$	BAC/C	BAC/JPM	BAC/WFC
Spline {09:30, 12:00, 16:00}	11	-60.13 (106.06)	-54.61 (95.05)	-31.27 (48.53)
Spline {09:30, 10:00, 16:00}	11	-54.97 (95.72)	-45.84 (77.50)	-19.97 (25.94)
Spline {09:30, 10:00, 13:00, 16:00}	12	-54.76 (109.52)	-44.23 (88.45)	-18.61 (37.23)
Spline {09:30, 10:00, 12:00, 14:00, 16:00}	13	-13.89 (41.99)	4.88 (4.41)	25.89 (-37.77)
		C/JPM	C/WFC	JPM/WFC
Spline {09:30, 12:00, 16:00}	11	-845.33 (1676.28)	-604.49 (1194.70)	-571.95 (1129.65)
Spline {09:30, 10:00, 16:00}	11	-768.95 (1523.50)	-538.29 (1062.30)	-504.24 (994.22)
Spline {09:30, 10:00, 13:00, 16:00}	12	-736.52 (1473.03)	-513.78 (1027.57)	-470.66 (941.32)
Spline {09:30, 10:00, 12:00, 14:00, 16:00}	13	-525.16 (1064.71)	-343.89 (702.06)	-254.51 (523.27)

outperforms the spline-based dynamic copula model in terms of fit. Although the models are not nested, the loglikelihood reductions are considerable. It comes as no surprise therefore that when we compare the models in terms of BIC reductions, we conclude that a fixed intraday spline does not capture the intraday dependence dynamics between discrete price changes as accurately as a model with a time-varying score driven dependence parameter. The score driven approach provides a better statistical description of our high-frequency data. To provide further evidence, we graphically display the dynamic copula parameter in Figure 5, for three randomly chosen trading days in 2012. These graphs also reveal that the daily pattern of θ_t^c may deviate substantially from the average intraday seasonal pattern.

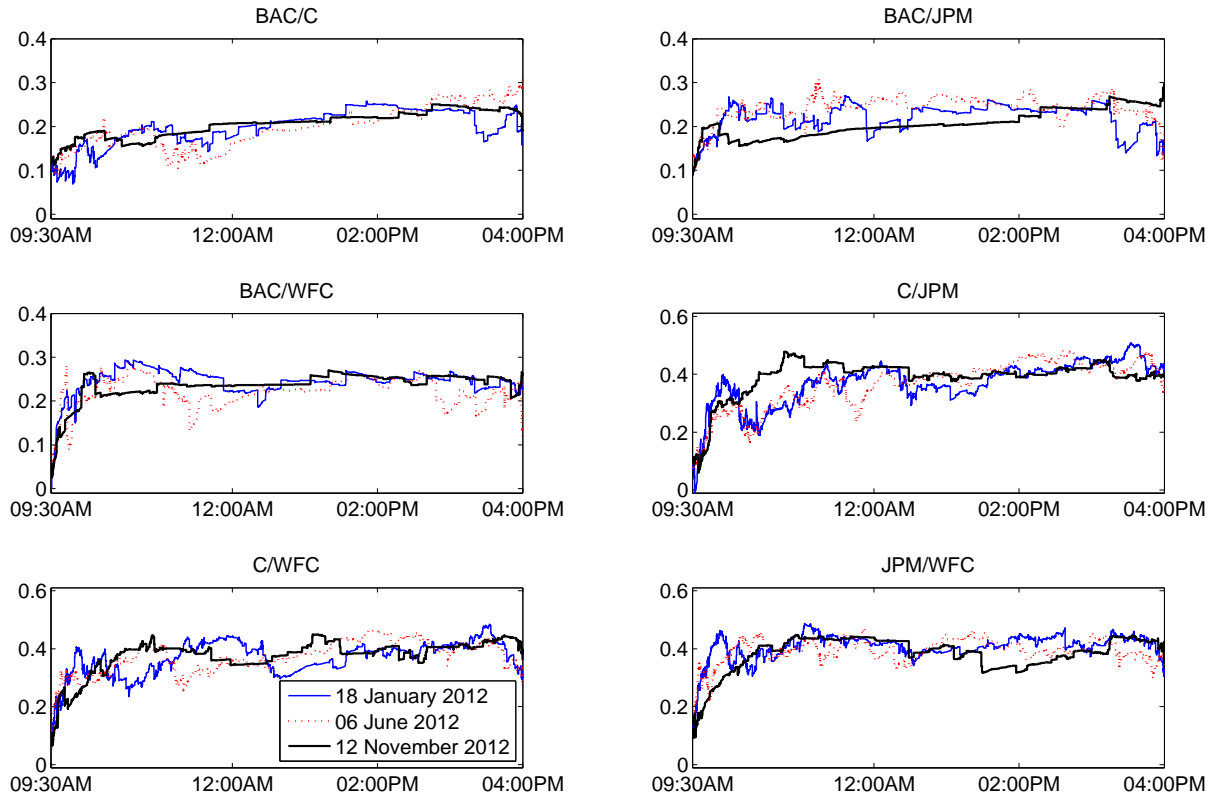


Figure 5: Copula dependence intraday patterns for a random selection of three days in 2012 based on the Gaussian copula with Skellam marginals. The selected days are 18 January 2012, 6 June 2012 and 12 November 2012. The panels show that the dependence pattern of a single day can be substantially less smooth than the point wise mean copula dependence path as presented in Figure 4.

5 Conclusions

Many empirical studies have concentrated on extracting high-frequency intraday volatility measures using tick-by-tick data. Here we have extended this literature to capture the intraday dynamic features of *dependence* using an observation driven model-based copula approach with discrete marginals. We have developed a new model to capture the intraday seasonal pattern of dependence between discrete tick-size price changes of different stocks. The complete dependence model is composed of dynamic Skellam marginal distributions for the discrete price changes combined with a time-varying copula structure. The dynamic specifications rely on the score of the predictive loglikelihood with respect to the relevant

dynamic parameters. The model performs well both in a controlled Monte Carlo setting and in an empirical study using high-frequency data. For four liquid U.S. financial stocks we found that the dependence structure varies over time during the trading day. There is a steep increase in dependence within the first hour of trading, and a steep decrease within the last 15 minutes of trading. We attribute these changes in dependence to the existence of more idiosyncratic risk components in the discrete price changes during the opening and closing hours of trading, in particular overnight firm-specific information accumulation when the market opens and the unwinding of inventory positions when the market closes.

References

- Admati, A. R. and P. Pfleiderer (1988). A theory of intraday patterns: Volume and price variability. *The Review of Financial Studies* 1(1), 3–40.
- Allez, R. and J. P. Bouchaud (2011). Individual and collective stock dynamics: intra-day seasonalities. *New Journal of Physics* 13, 1–12.
- Alzaid, A. and M. A. Omair (2010). On the Poisson difference distribution inference and applications. *Bulletin of the Malaysian Mathematical Sciences Society* 33(1), 17–45.
- Alzaid, A. A. and M. A. Omair (2014). Poisson difference integer valued autoregressive model of order one. *Bulletin of the Malaysian Mathematical Sciences Society*, forthcoming.
- Andersen, T. G. and T. Bollerslev (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115–158.
- Andersson, J. and D. Karlis (2014). A parametric time series model with covariates for integers in Z . *Statistical Modelling* 14(2), 135–156.
- Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2008). Realised kernels in practice: Trades and quotes. *Econometrics Journal* 4, 1–33.
- Barndorff-Nielsen, O. E., D. G. Pollard, and N. Shephard (2012). Integer-valued Lévy processes and low latency financial econometrics. *Quantitative Finance* 12(4), 587–605.
- Blasques, F., S. J. Koopman, and A. Lucas (2015). Information theoretic optimality of observation driven time series models for continuous responses. *Biometrika*, forthcoming.
- Brownlees, C. T. and G. M. Gallo (2006). Financial econometric analysis at ultra-high frequency: Data handling concerns. *Computational Statistics & Data Analysis* 51, 2232–2245.
- Cox, D. R. (1981). Statistical analysis of time series: some recent developments. *Scandinavian Journal of Statistics* 8, 93–115.
- Creal, D., B. Schwaab, S. J. Koopman, and A. Lucas (2014). Observation driven mixed-measurement dynamic factor models. *Review of Economics and Statistics* 96(5), 898–915.
- Creal, D. D., S. J. Koopman, and A. Lucas (2011). A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Business and Economic Statistics* 29, 552–563.

- Creal, D. D., S. J. Koopman, and A. Lucas (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics* 28(5), 777–795.
- De Lira Salvatierra, I. and A. J. Patton (2013). Dynamic copula models and high frequency data. Discussion paper, Duke University.
- Frank, M. J. (1979). On the simultaneous associativity of $f(x, y)$ and $x + y - f(x, y)$. *Aequationes Mathematicae* 19, 194–226.
- Freeland, R. K. (2010). True integer value time series. *AStA-Advances in Statistical Analysis* 94, 217–229.
- Genest, C. and J. Nešlehová (2007). A primer on copulas for count data. *Astin Bulletin* 37(2), 475–515.
- Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Econometric Series Monographs. Cambridge: Cambridge University Press.
- Harvey, A. C. and A. Luati (2014). Filtering with heavy tails. *Journal of the American Statistical Association* 109(507), 1112–1122.
- Irwin, J. O. (1937). The frequency distribution of the difference between two independent variates following the same Poisson distribution. *Journal of the Royal Statistical Society, Series A* 100(3), 415–416.
- Koopman, S. J., R. Lit, and A. Lucas (2014). The dynamic Skellam model with applications. Discussion paper, VU university and Tinbergen Institute.
- Koopman, S. J., A. Lucas, and M. Scharth (2015). Predicting time-varying parameters with parameter-driven and observation-driven models. *Review of Economics and Statistics*, forthcoming.
- Lucas, A., B. Schwaab, and X. Zhang (2014). Conditional euro area sovereign default risk. *Journal of Business and Economic Statistics* 32(2), 271–284.
- Münnix, M. C., R. Schäfer, and T. Guhr (2010). Impact of the tick-size on financial returns and correlations. *Physica A* 389(21), 4828–4843.
- Nelsen, R. B. (2006). *An Introduction to Copulas*. New York: Springer.
- Panagiotelis, A., C. Czado, and H. Joe (2012). Pair copula constructions for multivariate discrete data. *Journal of the American Statistical Association* 107, 1063–1072.

- Patton, A. J. (2002). Applications of copula theory in financial econometrics. Ph.D. dissertation, University of California, San Diego.
- Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review* 47(2), 527–556.
- Plackett, R. L. (1954). A reduction formula for Normal multivariate integrals. *Biometrika* 41(3,4), 351–360.
- Poirier, D. J. (1973). Piecewise regression using cubic spline. *Journal of the American Statistical Association* 68(343), 515–524.
- Rydberg, T. H. and N. Shephard (2003). Dynamics of trade-by-trade price movements: decomposition and models. *Journal of Financial Econometrics* 1(1), 2–25.
- Schepsmeier, U. and J. Stöber (2014). Derivatives and Fisher information of bivariate copulas. *Statistical Papers* 55, 525–542.
- Shahtahmassebi, G. (2011). *Bayesian Modelling of Ultra High-Frequency Financial Data*. Doctoral thesis, Research with Plymouth University. University of Plymouth.
- Skellam, J. G. (1946). The frequency distribution of the difference between two Poisson variates belonging to different populations. *Journal of the Royal Statistical Society* 109(3), 296.
- Sklar, A. (1959). *Fonctions de répartition à n dimensions et leurs marges*. Publications de l’Institut Statistique de l’Université de Paris, 8, 229-231.
- Tsay, R. S. (2005). *Analysis of financial time series* (2nd ed.). New Jersey: Wiley-Interscience.
- Wood, R. A., T. H. McInish, and J. K. Ord (1985). An investigation of transactions data for nyse stocks. *Journal of Finance* 40(3), 723–739.
- Zimmer, D. M. and P. K. Trivedi (2006). Using trivariate copulas to model sample selection and treatment effects: Application to family health care demand. *Journal of Business and Economic Statistics* 24(1), 63–76.

APPENDICES

A Derivation of the score vector

The derivations presented here focus on bivariate copulas but can easily be extended to higher dimensions. We assume a time-varying factor θ_t that consist of three elements, where the first two elements correspond to the marginal parameters and the third element corresponds to the copula dependence parameter. We have $\theta_t = (\theta_{1,t}^m, \theta_{2,t}^m, \theta_t^c)'$. The derivative of a bivariate copula with respect to $\theta_{1,t}^m$ is given by

$$\frac{\partial C(u_{1,t}, u_{2,t}; \theta_t^c)}{\partial \theta_{1,t}^m} = \frac{\partial C(u_{1,t}, u_{2,t}; \theta_t^c)}{\partial u_{1,t}} \cdot \frac{\partial u_{1,t}}{\partial \theta_{1,t}^m}. \quad (\text{A.1})$$

We observe that for the continuous parametric copula functions used in this paper, the first component on the right hand side of (A.1) can be written as a conditional copula

$$P(U_{2,t} \leq u_{2,t} | U_{1,t} = u_{1,t}) = \frac{\partial C(u_{1,t}, u_{2,t}; \theta_t^c)}{\partial u_{1,t}}. \quad (\text{A.2})$$

The second component on the right hand side of (A.1) is the derivative of the first marginal cdf, $u_{1,t} = F_1(y_{1,t}; \theta_{t,1}^m)$, with respect to $\theta_{1,t}^m$. The derivative of a bivariate copula with respect to θ_t^c is denoted by $\frac{\partial C(u_{1,t}, u_{2,t}; \theta_t^c)}{\partial \theta_t^c}$.

As a concrete example, consider a bivariate Gaussian copula with Skellam marginals, where $\theta_{i,t}^m = \log(\sigma_{i,t}^2)$, and $\rho_t = \theta_t^c / \sqrt{1 + (\theta_t^c)^2}$. This combination of copula, marginals, and parameterization is used in the application of Section 4. The Skellam distribution is discussed in Section 2. The bivariate Gaussian copula is given by

$$C_{\text{Ga}}(u_{1,t}, u_{2,t}; \rho_t) = \Phi_2 \left(\Phi^{-1}(u_{1,t}), \Phi^{-1}(u_{2,t}); \rho_t \right), \quad (\text{A.3})$$

where Φ_2 is a bivariate standard normal cdf, Φ^{-1} a univariate inverse standard normal cdf, and $\rho_t \in (-1, 1)$ is a correlation parameter. The first expression on the right hand side of (A.1) follows directly from a bivariate normal cdf, we have

$$\frac{\partial C_{\text{Ga}}(u_{1,t}, u_{2,t}, \rho_t)}{\partial u_{1,t}} = \Phi \left(\frac{\Phi^{-1}(u_{2,t}) - \rho_t \Phi^{-1}(u_{1,t})}{\sqrt{1 - \rho_t^2}} \right). \quad (\text{A.4})$$

A probably less well-known, but very useful result is given by Plackett (1954). It states that

for a bivariate standard Gaussian cdf, we have

$$\frac{\partial \Phi_2(x, y; \rho)}{\partial \rho} = (2\pi)^{-1}(1 - \rho^2)^{-1/2} \exp\left(-\frac{(x^2 - 2\rho xy + y^2)}{2(1 - \rho^2)}\right), \quad (\text{A.5})$$

where we can substitute $x = \Phi^{-1}(u_{1,t})$, $y = \Phi^{-1}(u_{2,t})$ and $\rho = \rho_t$ to obtain the appropriate expression for

$$\frac{\partial C(u_{1,t}, u_{2,t}; \theta_t^c)}{\partial \theta_t^c} = \frac{\partial C_{\text{Ga}}(u_{1,t}, u_{2,t}; \rho_t)}{\partial \rho_t} \cdot \frac{\partial \rho_t}{\partial \theta_t^c} = (1 + \theta_t^c)^{-3/2} \cdot \frac{\partial C_{\text{Ga}}(u_{1,t}, u_{2,t}; \rho_t)}{\partial \rho_t}. \quad (\text{A.6})$$

The first derivatives of the marginal Skellam cdfs in (A.1) are given by

$$\frac{\partial u_{i,t}}{\partial \sigma_{i,t}^2} = \exp(-\sigma_{i,t}^2) \sum_{\nu=-\infty}^k \left[\left(\frac{\nu}{\sigma_{i,t}^2} - 1 \right) I_{|\nu|}(\sigma_{i,t}^2) + I_{|\nu+1|}(\sigma_{i,t}^2) \right], \quad (\text{A.7})$$

with

$$\frac{\partial u_{i,t}}{\partial \theta_{i,t}^m} = \frac{\partial u_{i,t}}{\partial \sigma_{i,t}^2} \frac{\partial \sigma_{i,t}^2}{\partial \theta_{i,t}^m} = \sigma_{i,t}^2 \cdot \frac{\partial u_{i,t}}{\partial \sigma_{i,t}^2}, \quad (\text{A.8})$$

for $i = 1, 2$.

B Further tables and figures

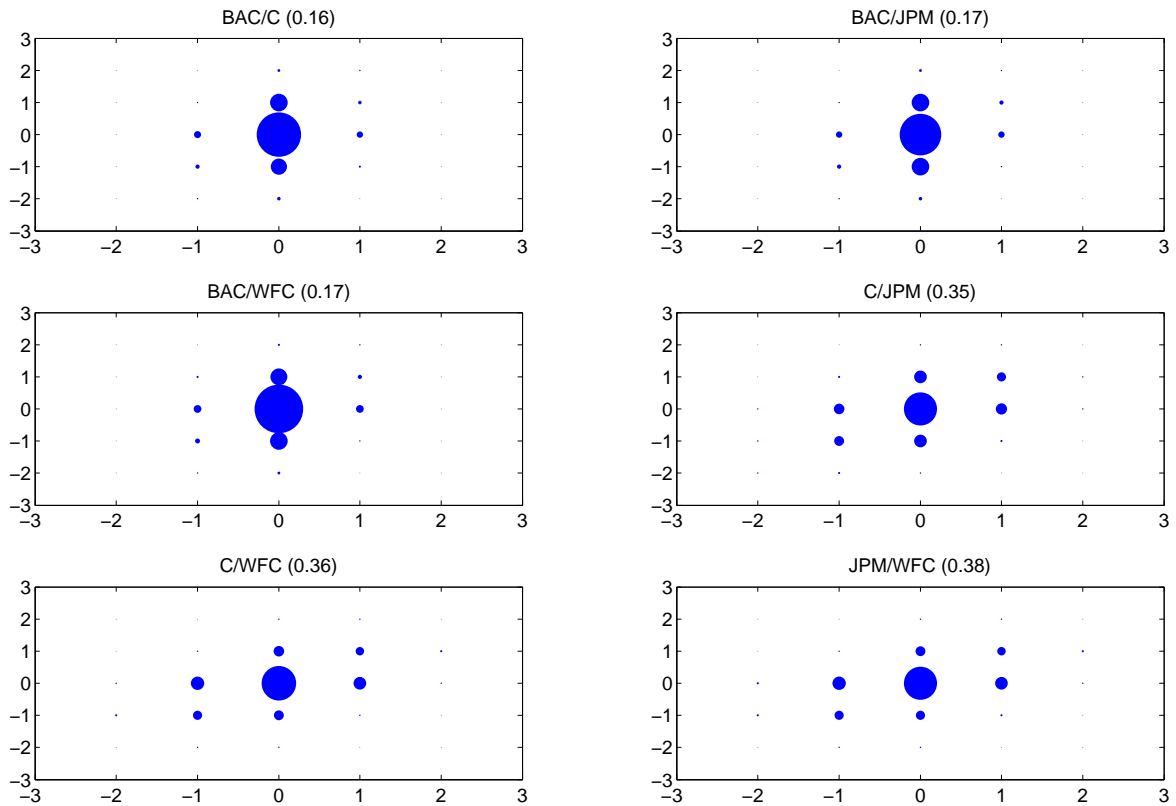


Figure B.1: The figure shows discrete scatter plots of the bivariate tick price change series of 20 trading days in April 2012 analysed in Section 4. The diameter of the circle represents the bivariate observation frequency in the data. We emphasize that the panels only show the situation where both price change series have a trade at time t . Both axes are for the interval $[-3, 3]$ since tick price changes outside this interval do not occur frequently enough to become visible in this plot. The reported value between parenthesis in the panel header is Pearson's linear correlation between the series.

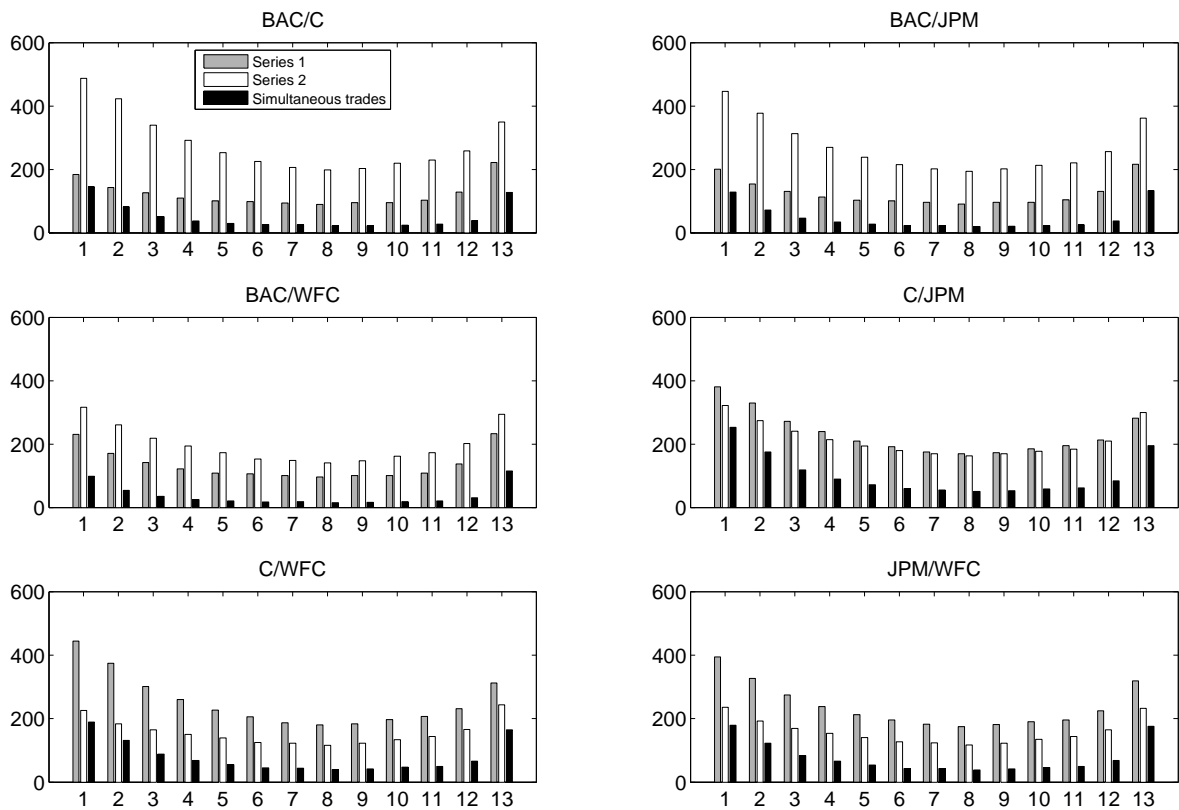


Figure B.2: The figure displays the number of simultaneous trades per half hour of the trading day as well as the the number of trades if only series 1 or series 2 trade. The numbers are averaged over all 250 trading days of the year 2012. The panels show the six combinations of stocks under consideration. The numbers on the x-axis represent the number of half hours in a trading day (13 in total).

Table B.1: The table reports maximum likelihood estimates of the Gaussian copula with Skellam marginals fitted to a data set of tick price changes for the year 2012. Standard errors are in parenthesis and are calculated using a numerically obtained Hessian matrix. The parameters are obtained in approximately four to 15 hours (depending on starting values and data sets) on a i7-2600, 3.40 GHz desktop PC using four cores.

Series	$\hat{\theta}_{1,1}$	$\hat{\theta}_{2,1}$	$\hat{\theta}_{3,1}$	$\hat{\omega}_1$	$\hat{\omega}_2$	$\hat{\omega}_3$	\hat{a}_1	\hat{a}_2	\hat{a}_3	\hat{b}_1	\hat{b}_2	\hat{b}_3
BAC/C	-0.263 (0.032)	0.844 (0.034)	0.116 (0.032)	-0.115 (0.002)	-0.002 (0.001)	2.54E-04 (0.001)	0.709 (0.011)	0.060 (0.001)	0.030 (0.003)	0.910 (0.003)	0.997 (1.24E-04)	0.999 (2.01E-04)
BAC/JPM	-0.249 (0.032)	1.167 (0.034)	0.089 (0.032)	-0.116 (0.002)	-0.002 (0.001)	4.97E-04 (0.001)	0.695 (0.010)	0.063 (0.001)	0.036 (0.006)	0.910 (0.002)	0.996 (1.32E-04)	0.998 (0.001)
BAC/WFC	-0.305 (0.030)	0.916 (0.031)	0.049 (0.030)	-0.132 (0.002)	-0.004 (0.001)	0.001 (0.001)	0.677 (0.009)	0.083 (0.002)	0.050 (0.004)	0.899 (0.003)	0.994 (2.03E-04)	0.997 (3.59E-04)
C/JPM	0.831 (0.008)	1.141 (0.008)	0.103 (0.008)	-0.002 (1.88E-05)	-0.002 (1.88E-05)	0.001 (1.88E-05)	0.057 (0.001)	0.061 (0.001)	0.030 (0.001)	0.997 (9.61E-05)	0.997 (1.07E-04)	0.999 (7.87E-05)
C/WFC	0.832 (0.008)	0.898 (0.008)	0.100 (0.008)	-0.002 (2.15E-05)	-0.003 (2.15E-05)	0.001 (2.15E-05)	0.057 (0.001)	0.080 (0.001)	0.033 (0.002)	0.997 (1.10E-04)	0.996 (2.40E-04)	0.999 (2.63E-04)
JPM/WFC	1.145 (0.008)	0.872 (0.008)	0.116 (0.008)	-0.002 (2.05E-05)	-0.003 (2.05E-05)	0.001 (2.05E-05)	0.062 (0.001)	0.076 (0.001)	0.034 (0.002)	0.996 (1.19E-04)	0.996 (2.06E-04)	0.998 (2.16E-04)