Modelling and Predicting Crypto–Currencies Financial Time–Series

Leopoldo Catania

Aarhus BSS and CREATES

NBP Workshop on Forecasting
(NBP 2018)

Warsaw, Poland
November 26–27, 2018

Based on: Catania and Grassi (2017), Catania, Grassi, and Ravazzolo (2018a) and Catania, Grassi, and Ravazzolo (2018b)

Motivation

- Bitcoin is the first decentralized cryptocurrency created in 2009 and documented in Nakamoto (2009).
- Since its inception, it gained a growing attention from the media, academics, and finance industry
- In recent months the global interest in Bitcoin and crypto–currencies has spiked dramatically:
  - Japan has recognized Bitcoin as a legal method of payment;
  - some central banks are exploring the use of the crypto–currencies;
  - a large number of companies and banks created the Enterprise Ethereum Alliance to make use of the crypto–currencies and the related technology called blockchain, Forbes (2017).
  - Finally the Chicago Mercantile Exchange (CME) stated the negotiation of Bitcoin futures on 18th of December 2017, see Exchange (2017). Nasdaq and Tokyo Financial Exchange will follow late in 2018, see Bloomberg (2017).
- All this interest has been reflected on the crypto–currencies market capitalization that exploded from around 19 billion in February 2017 to around 850 billion in January 2018.

Motivation

- The dynamic of those series is quite complex displaying:
  - extreme observations;
  - asymmetries;
  - nonlinear characteristics which are difficult to model.
- New econometric models to study the characteristics of those series are needed.
- As this new and unexplored market will develop, those models will be important for asset allocation, risk management, and pricing of derivative securities.

Bitcoin

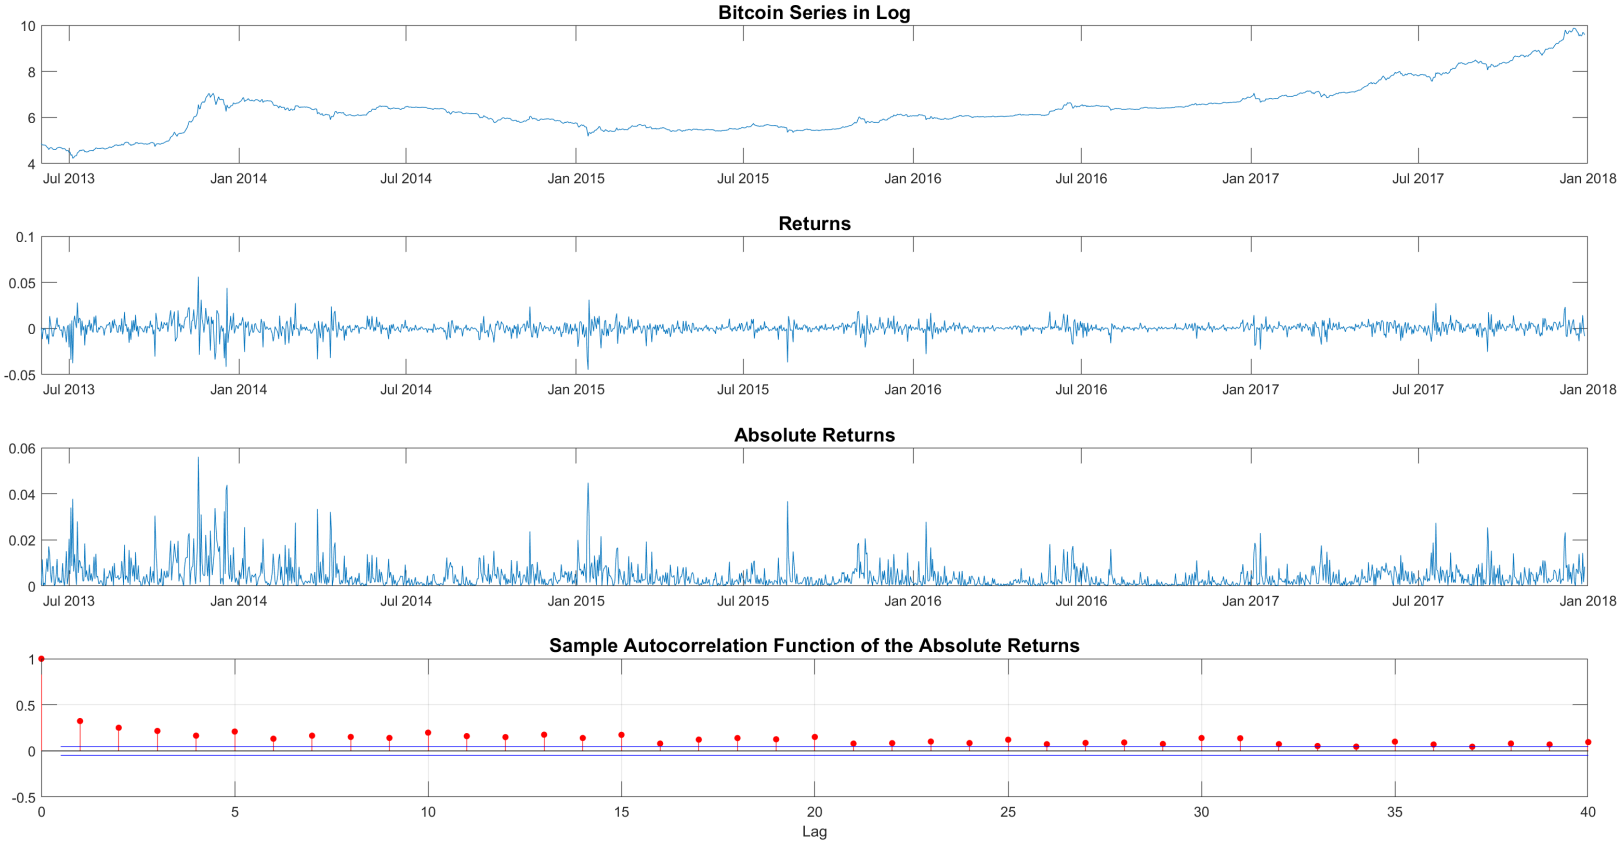
Bitcoin Series in Log
10
8
6
4
Jul 2013
Jan 2014
Jul 2014
Jan 2015
Jul 2015
Jan 2016
Jul 2016
Jan 2017
Jul 2017
Jan 2018
Returns
0.1
0.05
0
-0.05
Absolute Returns
0.06
0.04
0.02
0
Sample Autocorrelation Function of the Absolute Returns
1
0.5
0
-0.5
0
5
10
15
20
25
30
35
40
Lag

Literature review: Cryptocurrency

- Although Bitcoin can be considered to be relatively new, there has already been some initial analysis into the crypto–currency.
- Hencic and Gourieroux (2014) applied a non–causal autoregressive model to detect the presence of bubbles in the Bitcoin/USD exchange rate.
- Sapuric and Kokkinaki (2014) measure volatility of Bitcoin exchange rate against six major currencies.
- Caporale et. al (2017) examines the persistence in Crytocurrency.
- Chu et al. (2015) provide a statistical analysis of the log–returns of the exchange rate of Bitcoin versus the USD. They found that the Generalized Hyperbolic distribution seems to be the most appropriate choice to model the unconditional distribution of crypto–currencies time–series.

Literature review: GAS model

- Generalised Autoregressive Score (GAS) models have been recently introduced by Creal et al. (2013) and Harvey (2013).
- These are observation driven models in the sense of Cox (1981) which implies that the likelihood function is available in closed form, *i.e.*, there is no need to resort to computational intensive simulation methods.
- Their general specification assume:

$$
\begin{aligned}
y_t|\mathcal{F}_{t-1} &\sim \mathcal{D}(\theta_t, \psi) \\
\theta_t &= \Lambda(\widetilde{\theta}_t) \\
\widetilde{\theta}_t &= \kappa + A\widetilde{u}_{t-1} + B\widetilde{\theta}_{t-1}
\end{aligned}
$$

where

$$
\widetilde{u}_t \propto \frac{\partial \log p(y_t; \theta, \psi)}{\partial \theta}\Big|_{\theta=\theta_t},
$$

i.e.: $\widetilde{u}_t$ is proportional to the score with respect to the θ_t parameter.

GAS or GARCH: Different volatility reaction

- The plain GARCH(1,1) model:

$$y_t = \sigma_t \varepsilon_t, \quad \varepsilon_t \overset{iid}{\sim} (0,1)$$
$$\sigma_t^2 = \omega + \alpha y_{t-1}^2 + \beta \sigma_{t-1}^2$$

  can be written as:

$$\sigma_t^2 = \omega + \phi \sigma_{t-1}^2 + \alpha \sigma_{t-1}^2 u_{t-1}$$

  where $\phi = \alpha + \beta$ and $u_t = y_t^2/\sigma_t^2 - 1$ is a Martingale Difference (MD).
- A GAS model with Student's t innovations with ν degrees of freedom would imply:

$$u_t = \frac{(\nu+1)y_t^2}{(\nu-2)\sigma_t^2 + y_t^2} - 1$$

Does it really matter? Yes.

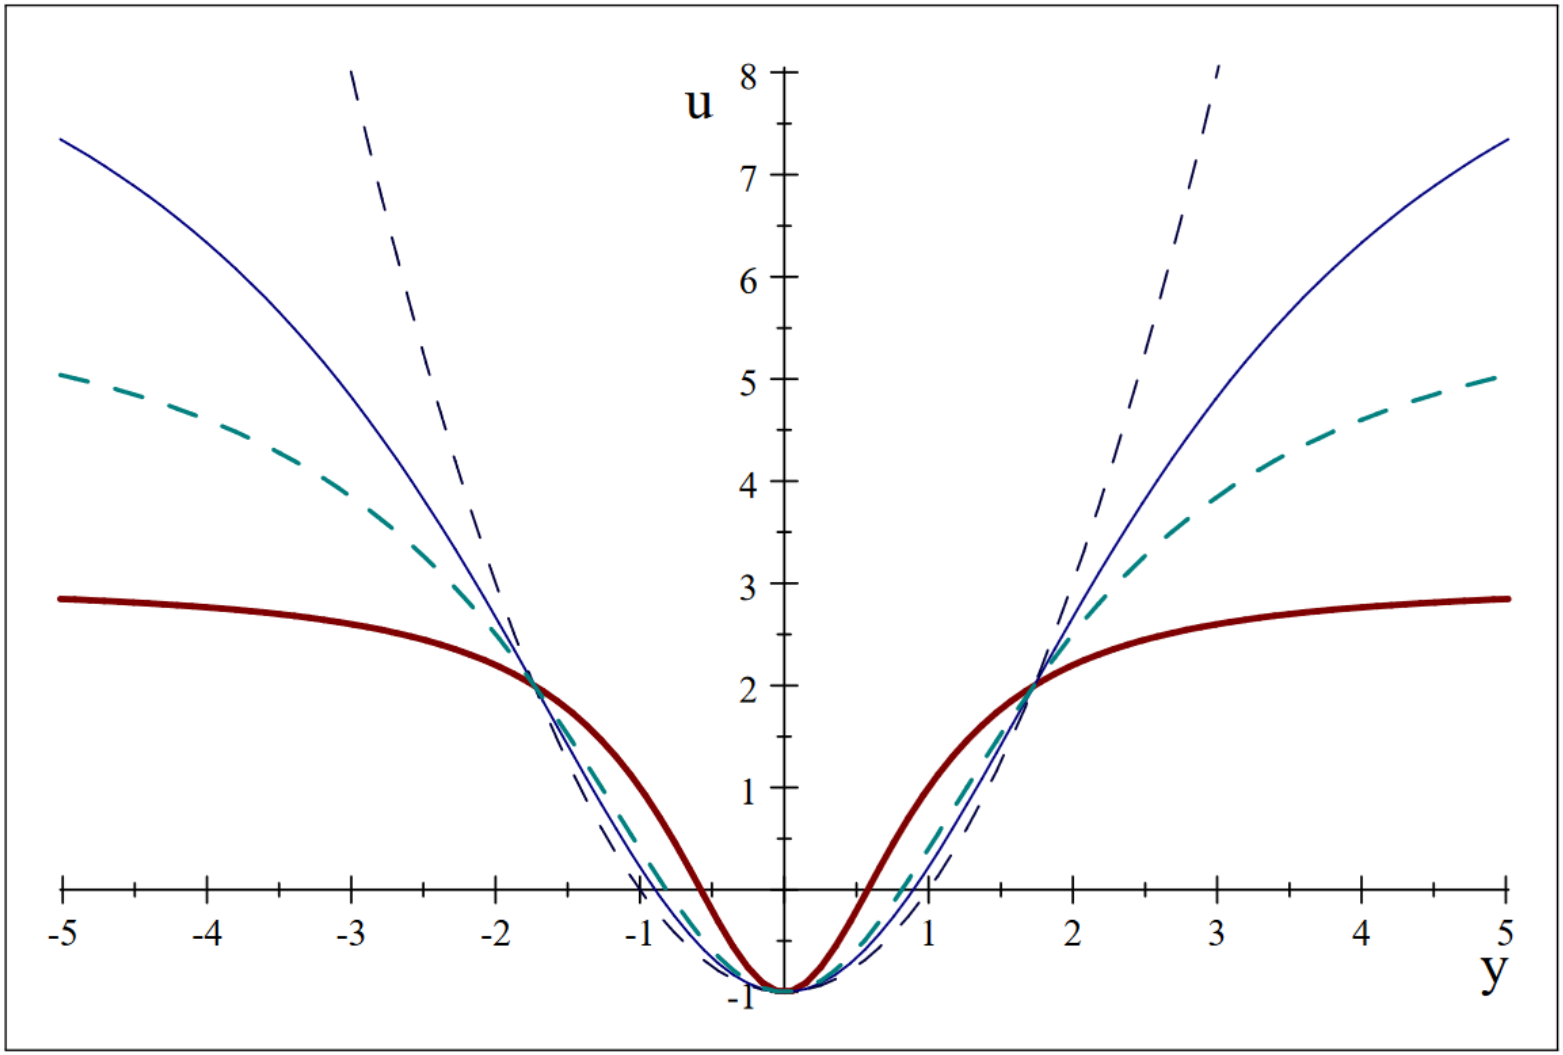

Impact of u_t for $\nu = 3$ (thick), $\nu = 6$ (medium dashed), $\nu = 10$ (thin) and $\nu = \infty$ (dashed).

GAS or GARCH ?

Taken from Harvey (2013)

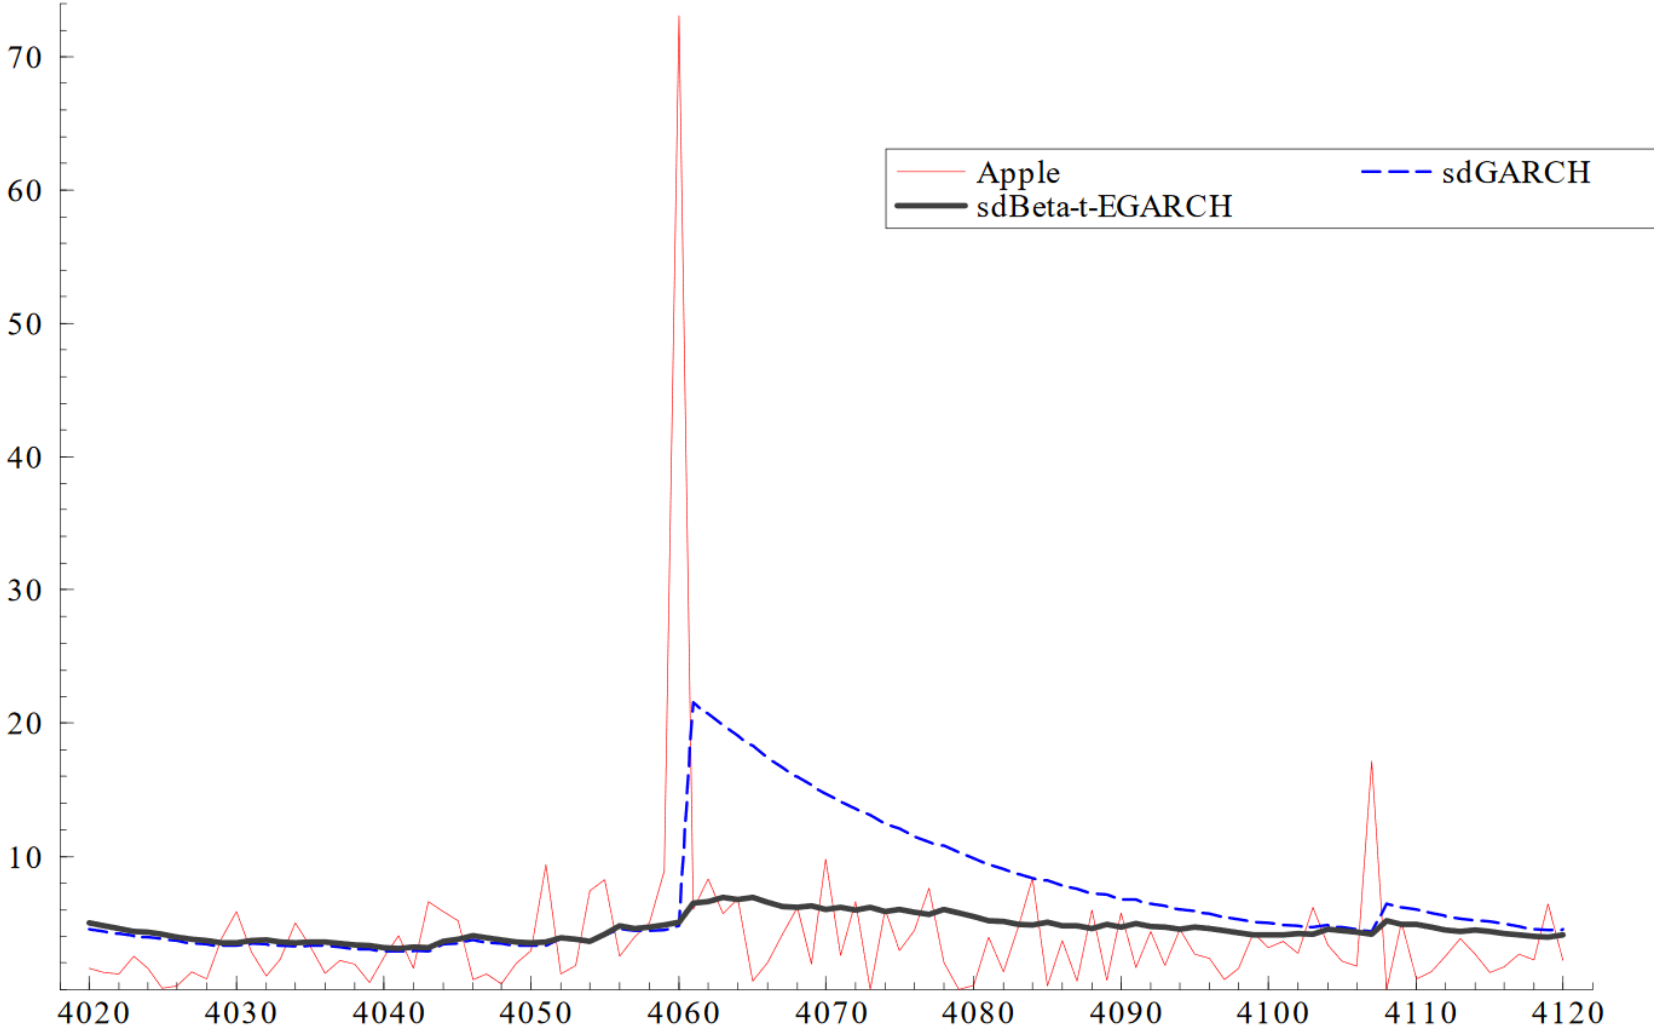

Absolute Apple returns and estimated volatility for GARCH and GAS with Student's t distribution.

GAS or GARCH ?

Taken from www.gasmodel.com

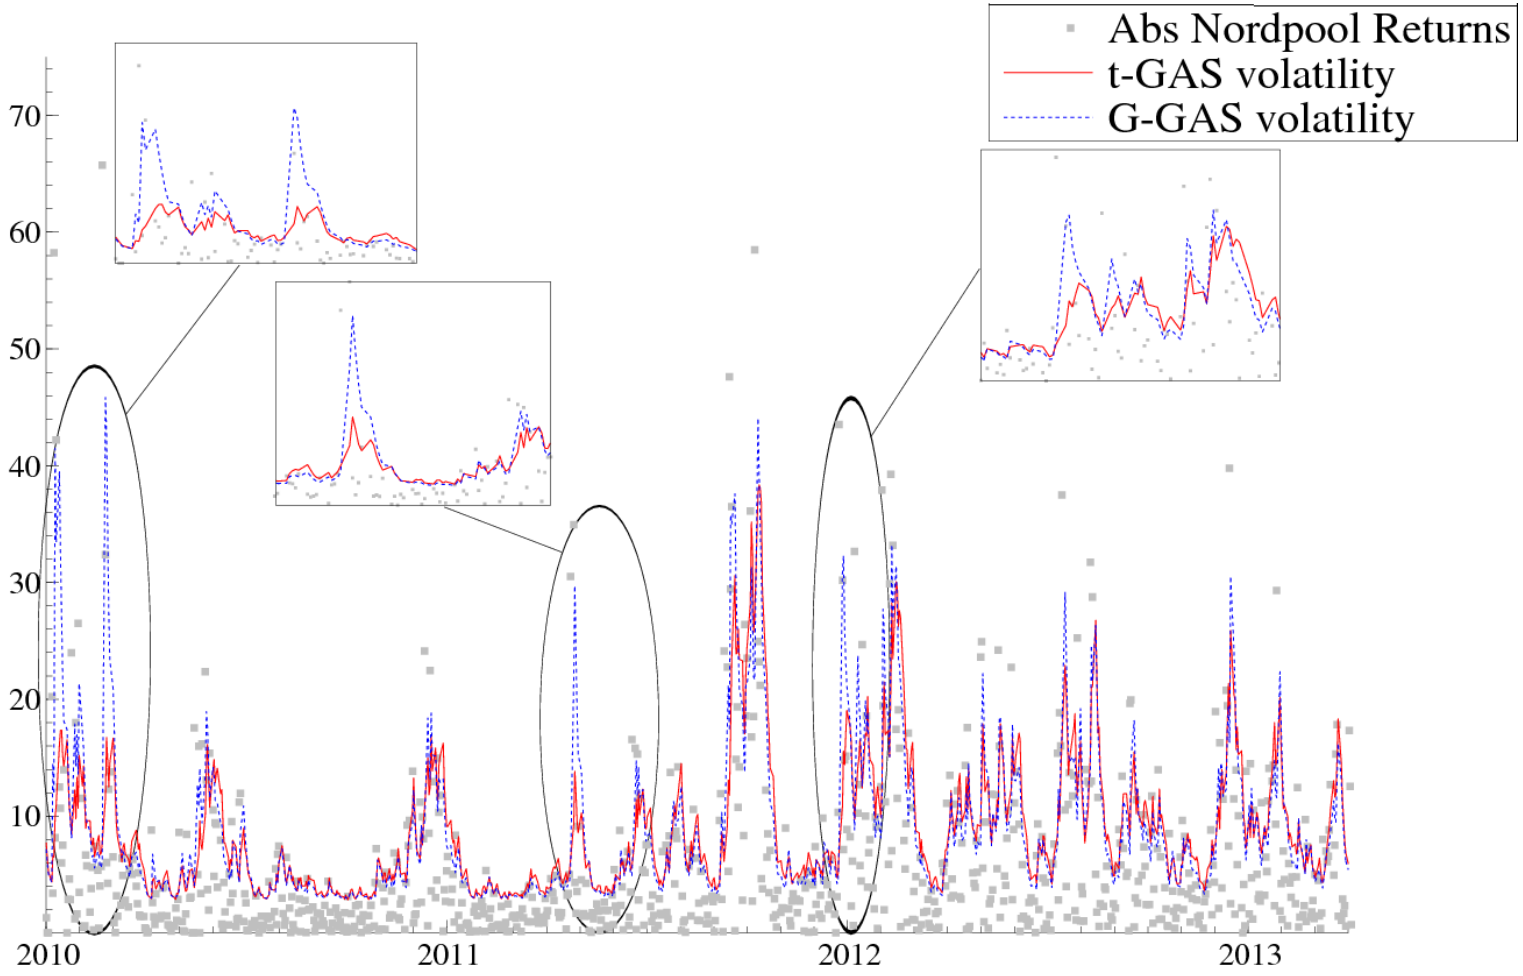

GAS estimated volatility for Nordpool electricity prices. Gaussian GAS is equivalent to Gaussian GARCH.

Methodological contributions of the paper

- In Catania and Grassi (2017) we propose a new Generalized Autoregressive Score driven model where the conditional distribution of the data, is assumed to follow the Generalized Hyperbolic Skew Student's t law (GAS-GHSKT);
- Moreover we extend Lucas et al. (2014) and Lucas et al. (2017) in several directions:
  - we derive expressions in a univariate framework, these cannot be easily recovered from Lucas et al. (2014),
  - we investigate in deep the characteristics of the filter implied for the volatility update using asymmetric tail behaviour of the GHSKT distribution.
  - we extend the model in order to account for several features such as:
    - leverage effects;
    - long memory properties of the volatility process;
    - time–varying higher order moments.

Empirical contributions of the paper

- We examine the dynamic properties (not the unconditional distribution) of crypto–currencies financial time–series.
- We also investigate a large set of crypto–currencies (289 series), which to the best of our knowledge has not been done yet.
- We investigate for the presence of leverage effects; long memory and time–varying higher order moments.
- We provide a forecasting exercise.

Model

- Let y_t be the realization of the series of interest (e.g. Bitcoin) at time t, we assume:

$$Y_t = \varsigma + \exp(h_t)\varepsilon_t, \qquad \varepsilon_t \overset{iid}{\sim} GHSKT(0,1,\boldsymbol{\theta}), \tag{1}$$

$$h_t = \phi(h_{t-1}, y_{t-1}; \boldsymbol{\phi}), \tag{2}$$

  where ς is a constant, h_t is the log–volatility obtained as a function of the past log–volatility level, h_{t-1}, and the realization y_{t-1}, via the known function $\phi : \Re \times \Re \to \Re^+$.

- The function ϕ constitutes the filter for the volatility and characterises the model specification, for example, when

$$\phi(h_{t-1}, y_{t-1}, \psi) = (1/2)\log(\phi_0 + \phi_1(y_{t-1} - \varsigma)^2 + \phi_2 \exp(2h_{t-1}))$$

  where $\psi = (\phi_0, \phi_1, \phi_2, \varsigma)'$, we obtain the popular GARCH specification of Bollerslev (1986).

- The only source of error affecting the dynamic system, ε_t, is assumed to be distributed independently as a standardized GHSKT.

The filter ϕ, in Generalized Autoregressive Score specification

Following Creal et al. (2013) and Harvey (2013) calculating the score of the conditional distribution $(Y_t|\mathbf{y}_1^{t-1})$, with respect to the log–volatility h_t we get the following model:

$$h_t = \phi(h_{t-1}, y_{t-1}), \tag{3}$$

$$= \kappa^h(1-\beta^h) + \alpha^h u_{t-1}^h + \beta^h h_{t-1}, \tag{4}$$

where:

- u_t^h is the score of the conditional distribution $(Y_t|\mathbf{y}_1^{t-1})$ with respect to h_t;
- $\kappa = \mathbb{E}[h_t]$ (since u_t is a Martingale Difference with respect to $(Y_t|\mathbf{y}_1^{t-1})$ for all t) is the unconditional log–volatility level;
- $|\beta^h| < 1$ is an autoregressive parameter controlling the persistence of the log–volatility process;
- $\alpha^h > 0$ is the parameter associated to the score which determines the step of the log–volatility update. Positiveness of α^h is assumed in order to not distort the signal given by the score, see Creal et al. (2013).

Monte Carlo experiment

	κ^h	α^h	β^h	ς	$\bar{\beta}$	ν
			$T = 500$			
BIAS	0.003	0.001	-0.077	0.000	1.099	1.763
RMSE	0.043	0.019	0.205	0.043	4.109	5.130
SD	0.042	0.019	0.190	0.043	3.959	4.817
			$T = 750$			
BIAS	0.001	0.001	-0.042	0.000	0.477	0.925
RMSE	0.017	0.015	0.139	0.035	1.440	2.723
SD	0.017	0.015	0.133	0.035	1.359	2.562
			$T = 1000$			
BIAS	0.000	0.001	-0.026	0.000	0.318	0.635
RMSE	0.010	0.012	0.100	0.030	0.990	2.027
SD	0.010	0.012	0.096	0.030	0.938	1.926
			$T = 2000$			
BIAS	0.000	0.000	-0.009	0.000	0.147	0.307
RMSE	0.004	0.008	0.037	0.021	0.520	1.176
SD	0.004	0.008	0.036	0.021	0.499	1.135

- The coefficients' value are similar to those found in the empirical study on crypto–currencies: $\kappa^h = 0$, $\alpha^h = 0.04$, $\beta^h = 0.94$, $\varsigma = 0$, $\bar{\beta} = 1.5$, $\nu = 7.0$.

First extension of GAS-GHSKT: Long Memory

- Long memory of the volatility process is a very well–known feature of financial returns, see Andersen et al. (2001).
- The time–series literature developed three main methodologies to incorporate this feature in dynamic models which are:
  i) fractional integration (Palma, 2007)
  ii) multiple components models (Andersen et al., 2006)
  iii) via the so–called heterogeneous autoregressive framework introduced by Corsi (2009).
- Here we use multiple components rewriting the filtering equation as:

$$h_t = \kappa^h + \sum_{l=1}^{L} h_{t,l}, \tag{5}$$

$$h_{t,l} = \beta_l^h h_{t-1,l} + \alpha_l^h u_{t-1}^h, \tag{6}$$

where

- L is the number of components;
- $h_{1,l} = 0$ and $\beta_l^h > \beta_{l-1}^h$, are imposed for identifiability purposes. Namely, the log–volatility process is now defined as a sum of autoregressive processes with different autoregressive parameters and same forcing variable u_t^h.
- The process $\{h_t, t > 0\}$ is still weekly stationary (provided the existence of $E[u_t^{h^2}]$).

Long Memory plot by components

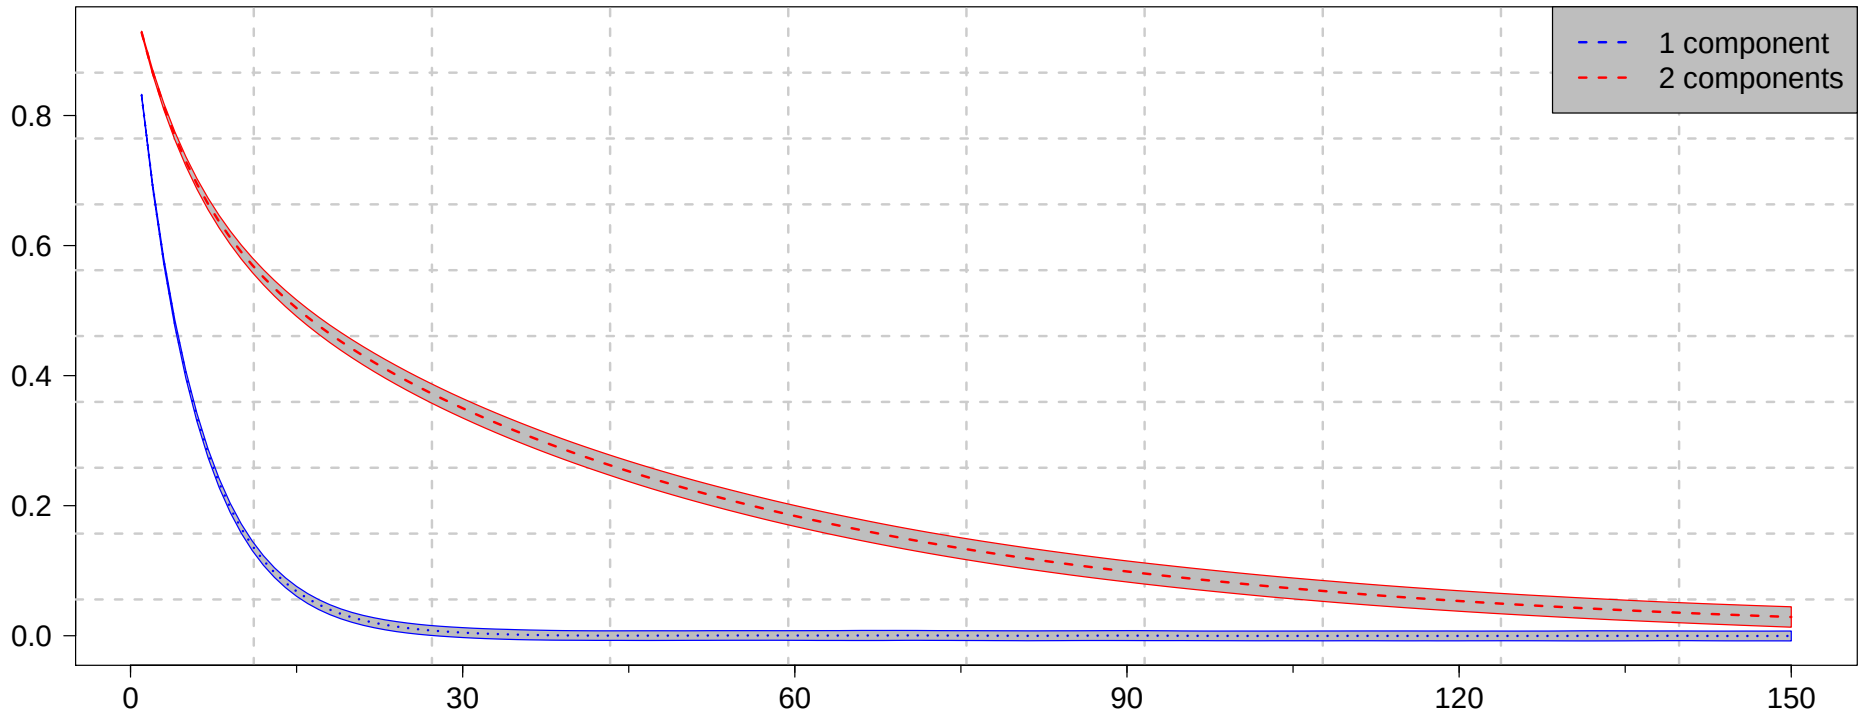

Autocorrelation function evaluated for one (dotted, blue) and two (dashed, red) components ($L = 2$) GHSKT volatility model. One standard deviation confidence bands are reported in gray. Parameter values driving the volatility dynamic are: $\kappa^h = 0$, $\alpha^h = 0.06$, $\beta^h = 0.84$ for the one component model and: $\kappa^h = 0$, $\alpha_1^h = 0.04$, $\alpha_2^h = 0.06$, $\beta_1^h = 0.98$, and $\beta_2^h = 0.84$ for the two component model.

Second extension of GAS-GHSKT: Leverage effect

- Leverage effects refer to the asymmetric reaction of the volatility process to past positive and negative returns, a stylised fact which dates back to Crane (1959).
- A Score–Driven volatility model with the inclusion of leverage effect has first been introduced by Harvey (2013).
- Following Harvey's implementation, we rewrite the main GAS equation as:

$$h_t = \kappa^h(1-\beta^h) + \alpha^h u_{t-1}^h + \beta^h h_{t-1} + \gamma^h \mathrm{sgn}(-(y_{t-1}-\varsigma))(u_{t-1}+1),$$

  where γ^h is an additional parameter and $\mathrm{sgn}(\cdot)$ is the sign function.
- Since we want to keep the moment condition $\mathbb{E}[h_t] = \kappa^h$, the filter is modified as:

$$h_t = \kappa^h(1-\beta^h-\gamma^h\varpi) + \alpha^h u_{t-1}^h + \beta^h h_{t-1} + \gamma^h \mathrm{sgn}(-(y_{t-1}-\varsigma))(u_{t-1}+1),$$

  where $\varpi = 2P(\varepsilon_t < 0) - 1$.

News impact curve

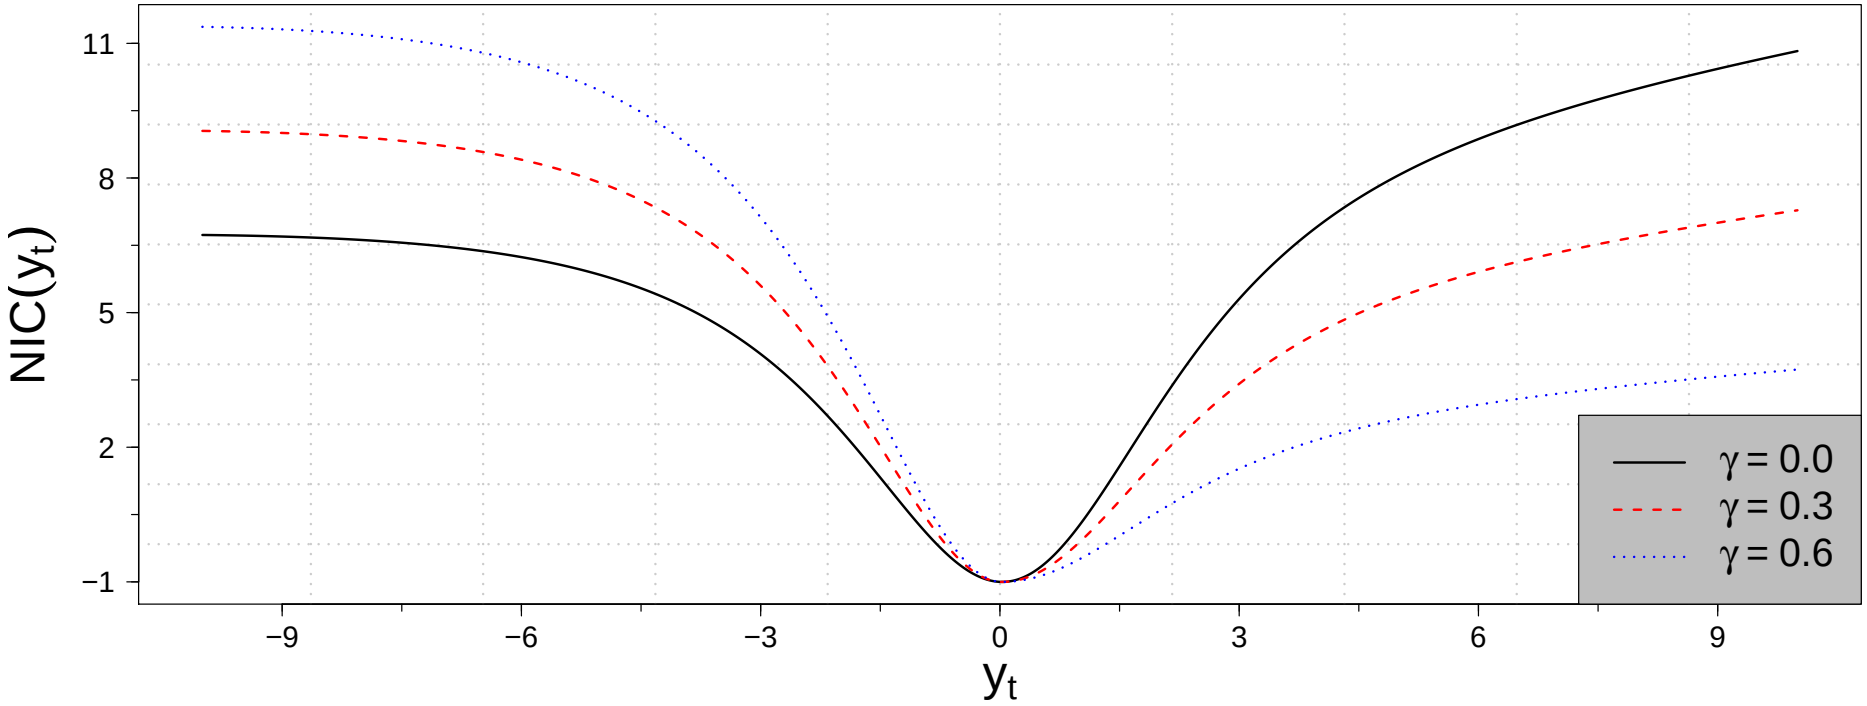

NIC with the inclusion of the leverage effect component for different values of γ: solid (black) line $\gamma = 0$; dashed (red) line $\gamma = 0.3$; dotted (blue) line $\gamma = 0.6$. The other parameters as set as follows: $\alpha^h = 1$, $\kappa^h = 0$, $\bar{\beta} = -0.5$ and $\nu = 7.0$.

Third extension of GAS-GHSKT: Time–varying higher order moments

- The inclusion of time–varying higher order moments for financial time–series dates back to the Autoregressive Conditional Density model of Hansen (1994).
- Before the development of Score–Driven models, the modeling of higher order moments was performed via ad hoc procedures. The usual methodology was that of including arbitrary transformations of past returns
- In the Score–Driven methodology the inclusion of time–varying higher order moments for financial time–series is easy see Lucas and Zhang (2016) and Bernardi and Catania (2016).
- Time–varying higher order are obtained by letting the skewness ($\bar{\beta}$) and shape (ν) parameters to be updated according to a GAS process.

Empirical Experiment

- The data used in this study are the crypto–currencies closing prices. Those data are freely available from https://coinmarketcap.com/
- Since the introduction of Bitcoin in 2009, hundreds of other crypto–currencies have been created and, as of Janualy 2018, 1440 crypto–currencies exist.
- Some of them, do not represent serious attempts at establishing a foothold in the market, others are quite young with just a few useful observations.
- For this reason, we focus on coins with at least 700 observations, which correspond to almost 2 years of data, this leaves us with 289 series.

Empirical Experiment

- We group each series into three categories based on the market capitalization that are:
  i) Large Cap (top 20%)
  ii) Mid Cap (between 40% and 60%)
  iii) Small Cap (lower 20%).
- We report results at the aggregate level (All), for the three categories (Large, Mid and Small Cap), and for the four major crypto–currencies:
  i) Bitcoin
  ii) Ethereum
  iii) Ripple
  iv) Litecoin.

Four principal crypto–currencies

Coin	Bitcoin	Ethereum	Ripple	Litecoin
Nr. Obs.	1497	702	1412	1437
Created	03-Jan-09	01-Aug-14	01-Jul-13	01-Nov-13
Purpose	Payment System	Smart Contract	Currency exchange	Payment System
Supply	21 Millions Total	18 Millions Yearly	100 Billions Total	84 Millions Total
Market Cap	200 Billions	44 Billions	9.5 Billions	5.5 Billions

In sample results

Score–Driven GHSKT vs GARCH–GHSKT

- The GARCH model with the same conditional distributional assumption, is defined as

$$y_t = \varsigma + \sigma_t \varepsilon_t, \quad \varepsilon_t \sim \mathcal{GHSKT}(0, 1, \bar{\beta}, \nu), \tag{7}$$

$$\sigma_t^2 = \zeta_0 + \zeta_1 (y_{t-1} - \varsigma)^2 + \zeta_2 \sigma_{t-1}^2, \tag{8}$$

  where $\zeta_i > 0$, for $i = 0, 1, 2$ and $\zeta_1 + \zeta_2 < 1$, see Bollerslev (1986).
- It is worth noting that the NIC implied by the GARCH filter is quadratic in the residual $y_t - \varsigma$, and hence different from the one implied by the Score–Driven filter; it is not robust to extreme observations.
- The number of parameters for the GARCH and Score–Driven model is the same; model comparison can be performed by a simple likelihood comparison.

Kernel density estimated on the likelihood percentage increase between GHSKT Score–Driven and GARCH–GHSKT model

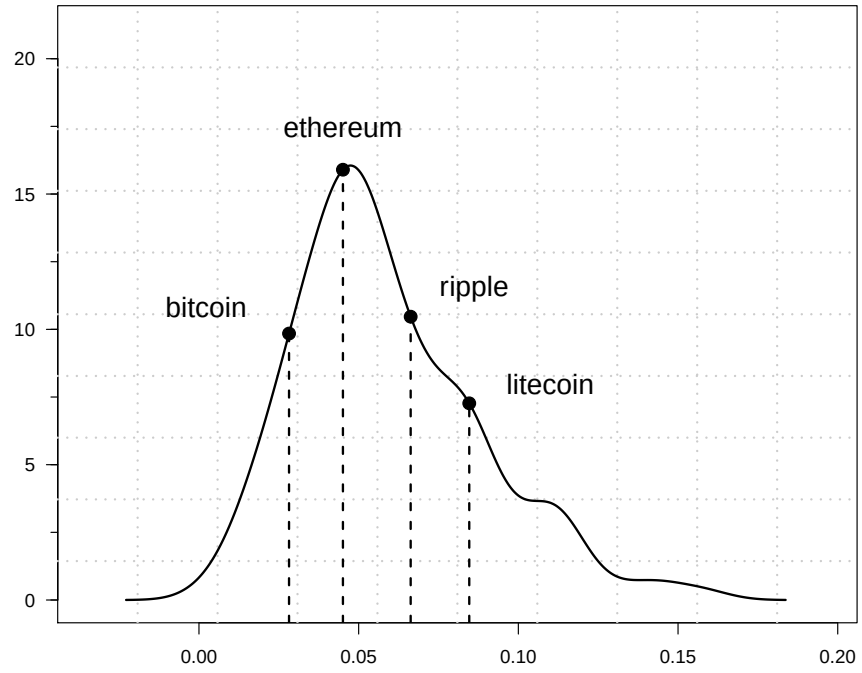

(a) All

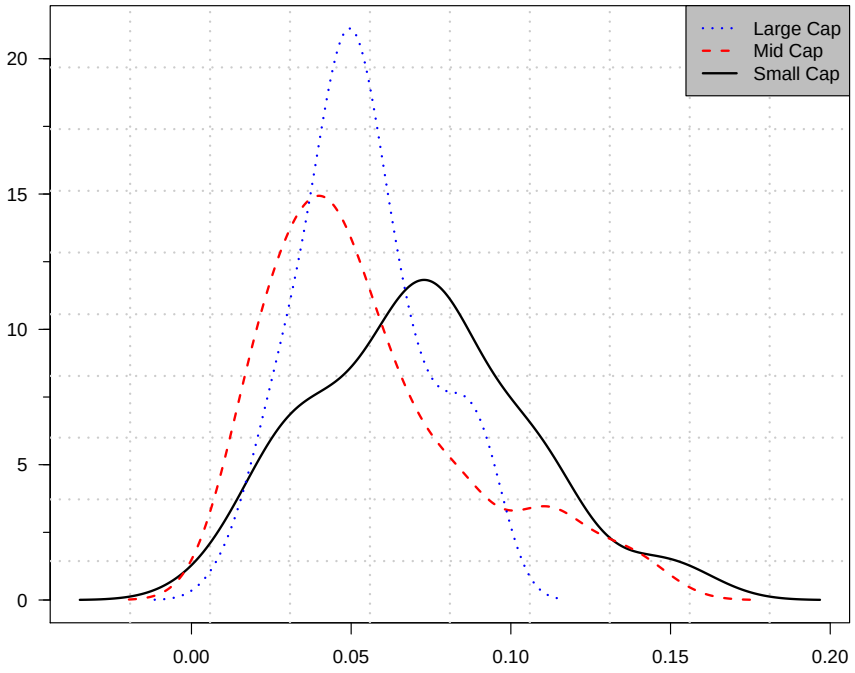

(b) Subgroups

Kernel density estimated on the likelihood percentage increase between the GHSKT Score–Driven model and the GARCH–GHSKT model. Positive values indicate outperformance of the GHSKT Score–Driven with respect to the GARCH–GHSKT alternative.

Filtered volatility between GHSKT Score–Driven and GARCH–GHSKT model

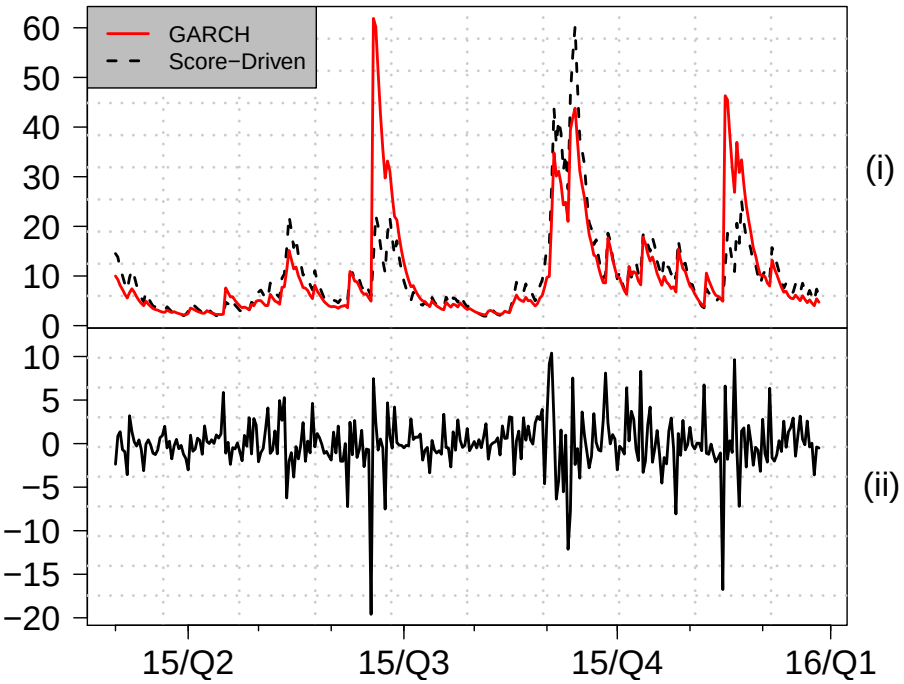

(c) Bitcoin

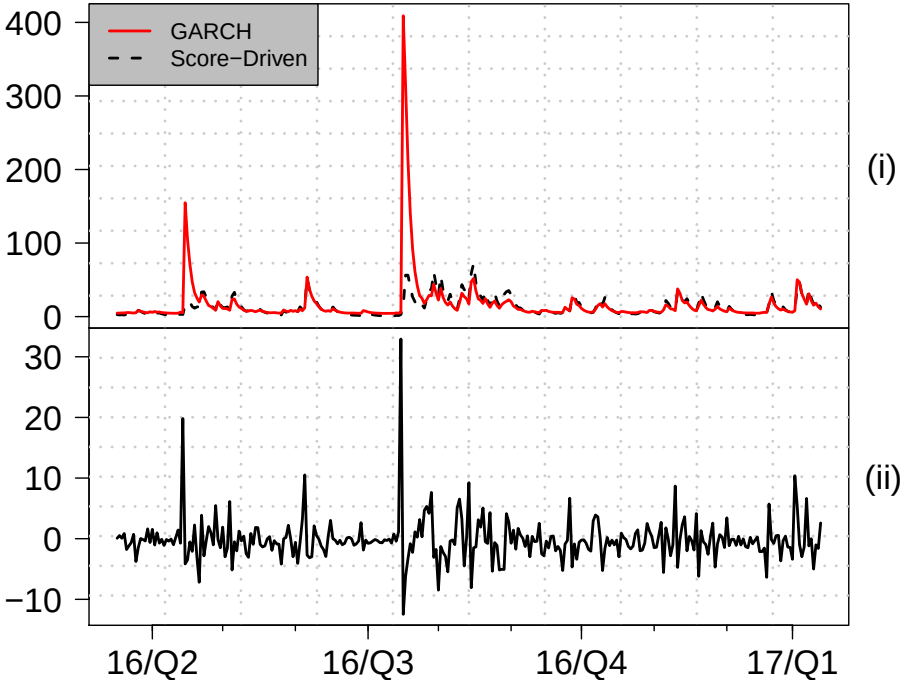

(d) Ripple

Filtered conditional variances for Bitcoin (left) and Ripple (right) delivered by the Score–Driven GHSKT model (black dashed line) and the GARCH (red solid line).

Long memory component

#	Bitcoin	Ethereum	Ripple	Litecoin	Small Cap	Mid Cap	Large Cap	All
Panel (a): AIC								
1	x	x	–	–	35	33	25	28
2	–	–	–	x	58	61	61	65
3	–	–	x	–	7	4	14	7
4	–	–	–	–	0	2	0	0
Panel (b): BIC								
1	x	x	–	x	75	74	65	66
2	–	–	x	–	25	26	35	34
3	–	–	–	–	0	0	0	0
4	–	–	–	–	0	0	0	0
Panel (c): HQ								
1	x	x	–	–	51	56	37	44
2	–	–	x	x	49	44	61	56
3	–	–	–	–	0	0	2	1
4	–	–	–	–	0	0	0	0

Leverage component

	Bitcoin	Ethereum	Ripple	Litecoin	Small Cap	Mid Cap	Large Cap	All
AIC	–	x	x	x	54	46	72	57
BIC	–	–	x	–	26	19	44	30
HQ	–	–	x	–	42	32	58	42

Comparison between models with and without the leverage component. Comparison is made according to the three information criteria: AIC, BIC, and HQ.
The label x indicates preference with respect to the model with the leverage component, the label $--$ indicates the opposite. Columns 5 to 8 report the percentage of times the model with the inclusion of the leverage component is preferred for the three groups: Small Cap, Mid Cap, and Large Cap, as well as over the full sample (All).

Time–varying Higher Order Moments

	Bitcoin	Ethereum	Ripple	Litecoin	Small Cap	Mid Cap	Large Cap	All
Panel (a): Time–varying Skewness								
AIC	–	x	–	–	35	35	49	41
BIC	–	–	–	–	11	7	26	14
HQ	–	–	–	–	25	23	39	28
Panel (b): Time–varying Kurtosis								
AIC	–	–	–	–	0	0	0	0
BIC	–	–	–	–	0	0	0	0
HQ	–	–	–	–	0	0	0	0

Time–varying Higher Order Moments

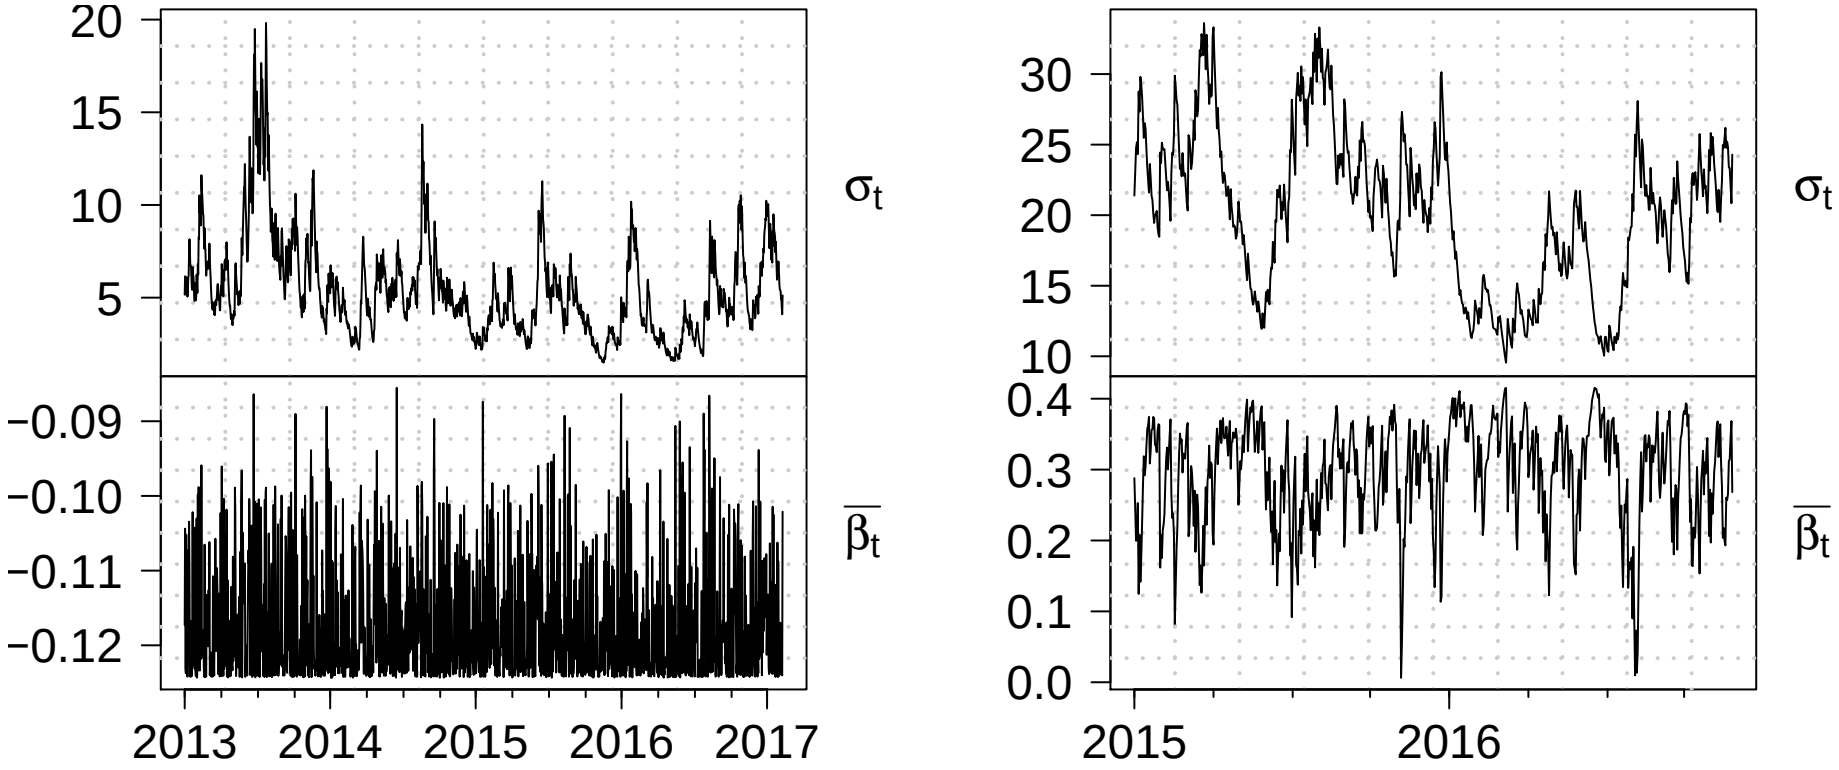

(e) Bitcoin (f) Ethereum

Figure: Filtered volatility (σ_t) and skewness parameter ($\bar{\beta}_t$) for Bitcoin and Ethereum over the full sample.

Time–varying Higher Order Moments

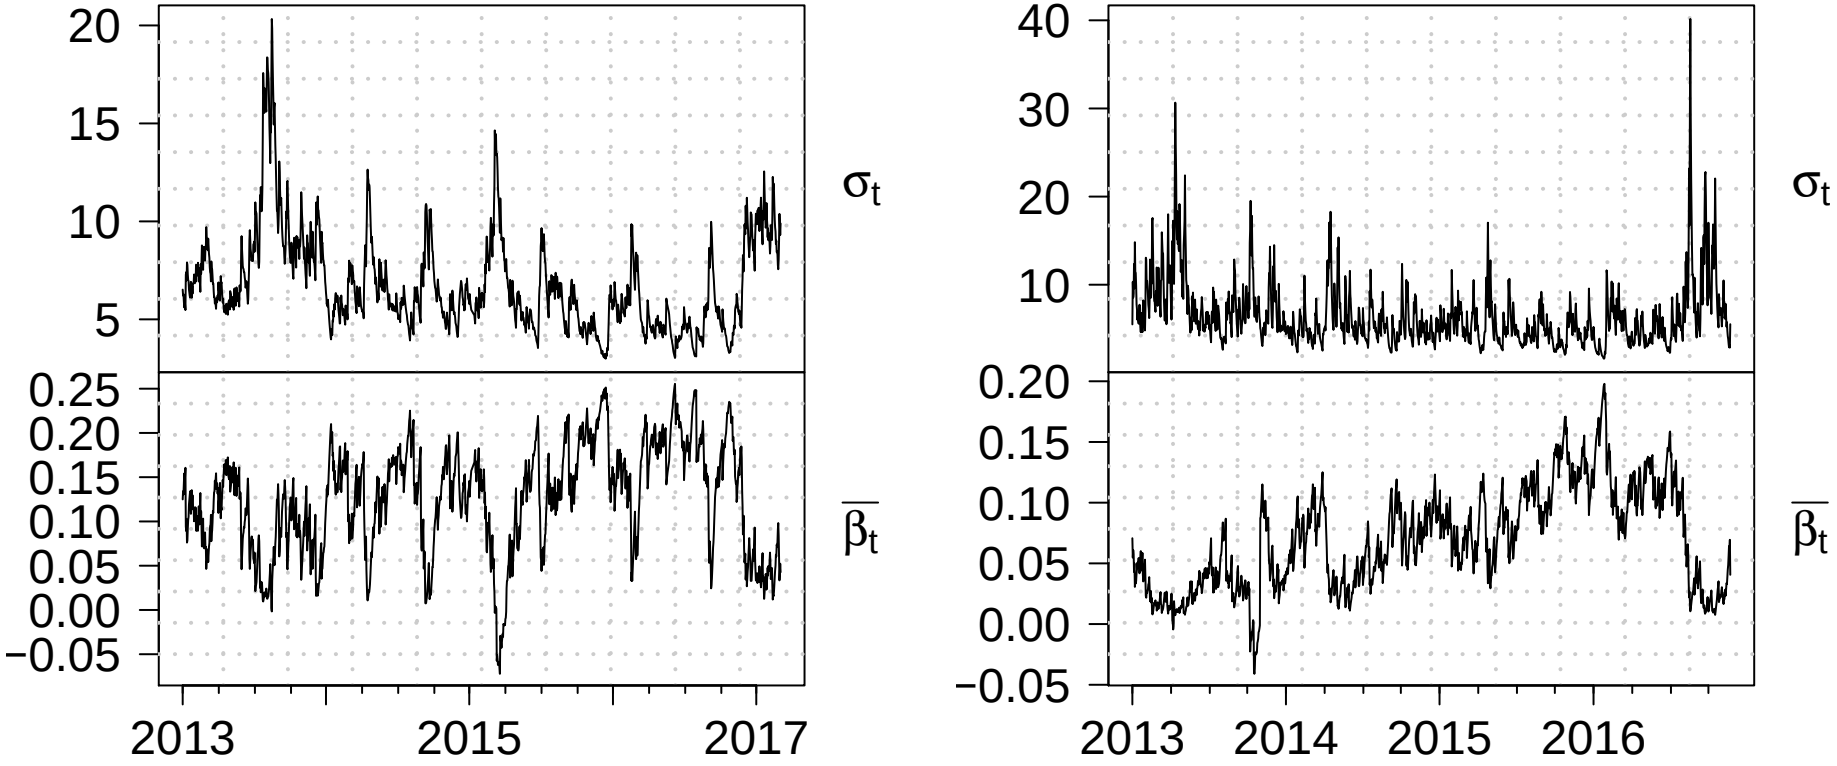

(a) Litecoin (b) Ripple

Figure: Filtered volatility (σ_t) and skewness parameter ($\bar{\beta}_t$) for Litecoin and Ripple over the full sample.

Out of sample results

Quantitative risk management applications

- In Catania et al. (2018a) we study the benefits of a robust filter for the volatility process of crypto–currencies time–series, we report a forecasting analysis comparing predictions delivered by the robust Score–Driven GHSKT filter and the GARCH filter.
- The quantity of interest for this analysis is the H–step ahead Value–at–Risk (VaR) level of the crypto–currency time–series at confidence level $\alpha \in (0, 1)$.
- Starting from our data set of 289 time–series, we select those with a sample size of at least 1.200 observations leaving us with 49 series. The in–sample size of the first forecast is set to $T = 600$, such that we have at least 600 predictions in order to compare the two models. Models' parameters are estimated using the first 600 observations and subsequently getting updated each time a new observation becomes available using an extending window.

Quantitative risk management applications

- For both models, multi step ahead predictions ($H > 1$) of the VaR level are computed by Monte Carlo simulation using 10.000 draws from the predictive distribution.
- In order to compare the Score–Driven and GARCH GHSKT models we compute the quantile loss associated to each VaR prediction delivered by the two models.
- Given a prediction $VaR^{\alpha}_{T+H|T}$, and the realization of the crypto–currency log–return, y_{T+H}, the quantile loss is defined as:

$$QL^{\alpha}_{T+H|T} = (\alpha - \mathbf{1}_{\{y_{T+H} < VaR^{\alpha}_{T+H|T}\}})(y_{T+H} - VaR^{\alpha}_{T+H|T}), \tag{9}$$

  where $\mathbf{1}_{\{y_{T+H} < VaR^{\alpha}_{T+H|T}\}} = 1$ if $y_{T+H} < VaR^{\alpha}_{T+H|T}$ and zero otherwise, see González-Rivera et al. (2004).
- Quantile loss values are subsequently averaged over the forecasting horizons, the model with the lower averaged QL value is preferred.

Quantitative risk management applications

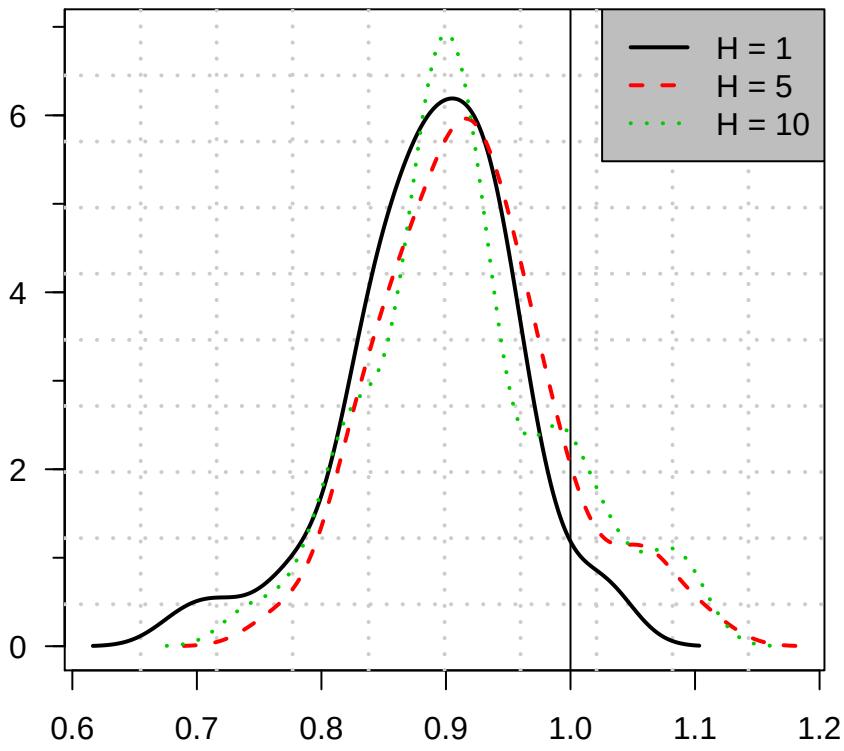

(a) $\alpha = 1\%$

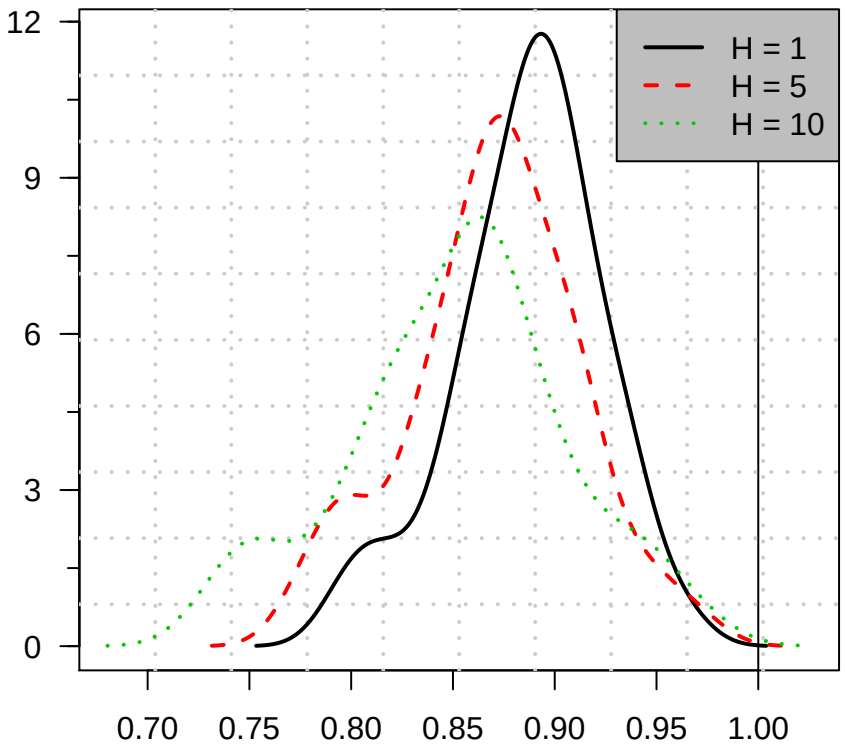

(b) $\alpha = 5\%$

Kernel density estimated on the ratio between the average quantile loss of the Score–Driven and GARCH GHSKT models. Left figure reports results for the $\alpha = 1\%$ VaR confidence level, right figure reports results for $\alpha = 5\%$. Values lower than one indicate outperformance of the predictions based on the Score–Driven volatility filter versus the GARCH alternative. The vertical bar at one indicates equal performance between the two specifications.

Conditional Mean Prediction

- In Catania et al. (2018b) we study the predictability of cryptocurrencies time series
- We compare several alternative univariate and multivariate models in point and density forecasting of four of the most capitalized series: Bitcoin, Litecoin, Ripple and Ethereum
- We apply a set of crypto-predictors and rely on Dynamic Model Averaging to combine a large set of univariate Dynamic Linear Models and several multivariate Vector Autoregressive models with different forms of time variation
- We find statistical significant improvements in point forecasting when using combinations of univariate models and in density forecasting when relying on selection of multivariate models.

The forecast design

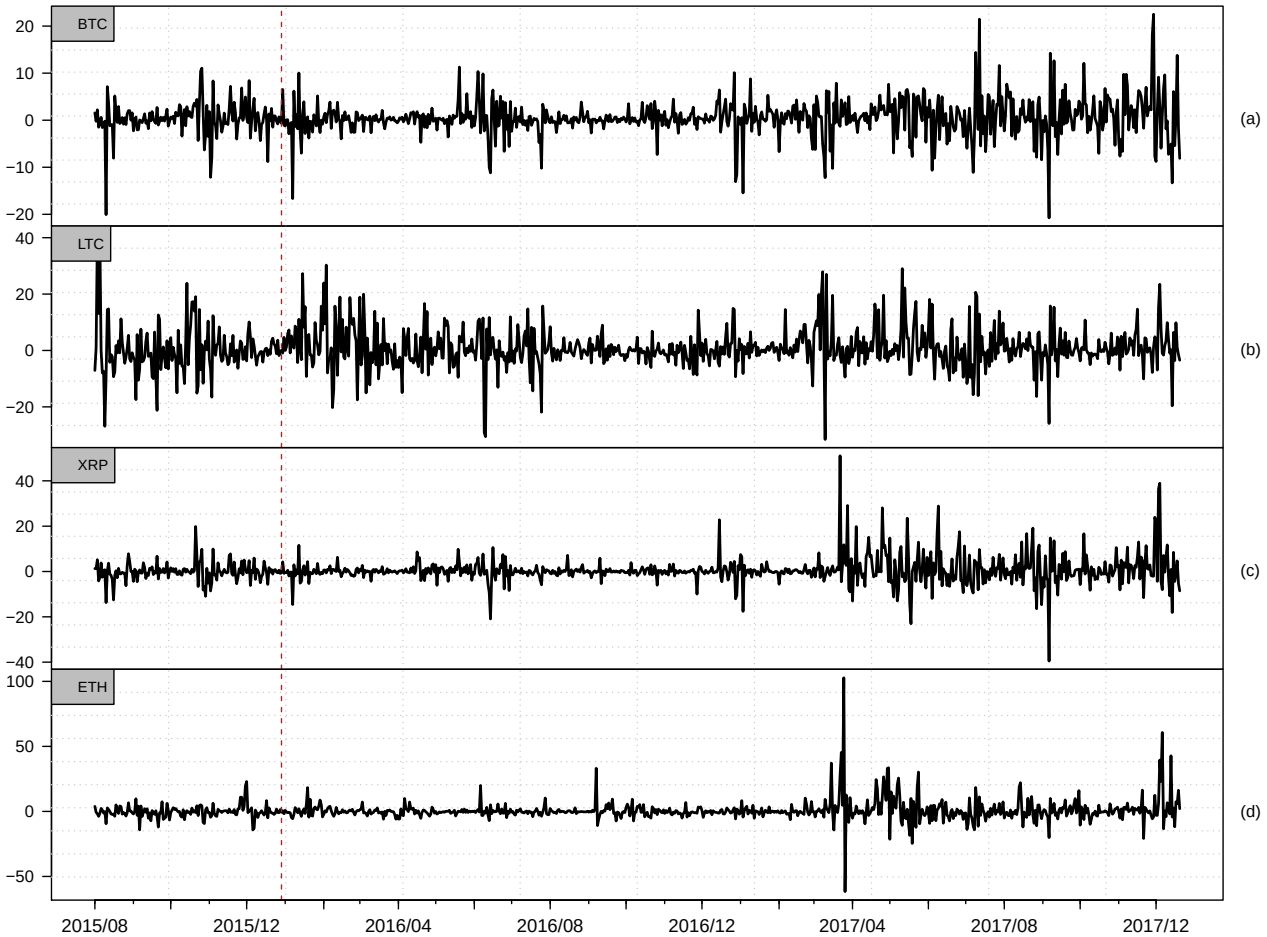

Figure: The plots show the four crypto–currencies daily percentage log returns considered in this study: Bitcoin (BTC) in panel (a); Litecoin (LTC) in panel (b); Ripple (XRP) in panel (c); and Ethereum (ETH) in panel (d). The dashed horizontal red line indicates the beginning of the out-of-sample period on the 6th June, 2016. The full sample spans from 9th August, 2015 to 28th December, 2017 for a total of 873 observations.

Cryptocurrencies and Crypto–predictors

Data Overview		
Abbreviation	Full name	Transformation
Cryptocurrencies time series		
BTC	Bitcoin	First difference of Log
ETH	Ethereum	First difference of Log
XRP	Ripple	First difference of Log
LTC	Litecoin	First difference of Log
Additional crypto–explicative time series		
BTC_HL	Bitcoin High minus Bitcoin Low	Log
ETH_HL	Ethereum High minus Ethereum Low	Log
XRP_HL	Ripple High minus Ripple Low	Log
LTC_HL	Litecoin High minus Litecoin Low	Log
Additional financial and macro time series		
CDS_5y	Europe credit default swap index 5 years	First difference of Log
ES_600	Stoxx Europe 600 - Price Index	First difference of Log
GLD	Gold Bullion LBM	First difference of Log
NK_225	Nikkei 225 Stock Average - Price Index	First difference of Log
SP_500	S&P 500 Composite - Price Index	First difference of Log
SV	Silver Handy & Harman Base Price	First difference of Log
BD_1m	1-Month US Treasury Constant Maturity Rate	First difference
BD_10y	10-Year US Treasury Constant Maturity Rate	First difference
VIX	VIX closing price	Log

Univariate models

All univariate models have been estimated using the R package **eDMA** of Catania and Nonejad (2018)

Abbreviation	*Full Description*
AR(1)	Autoregressive model of order one, benchmark model.
KS	Kitchen Sink specification.
KS–noregr	Kitchen Sink with no covariates.
DMA	DMA across all models and forgetting factor combinations. See Dangl and Halling (2012).
DMS	Dynamic Model Selection (DMS).
DMA–noregr	DMA with only the lagged values of the series as covariates.
DMS–noregr	DMS with only the lagged values of the series as covariates.

Table: Univariate models considered in the forecasting exercise. The first column is the model's abbreviation. The second column provides a brief description of each individual model.

Univariate forecasting results: MSE

h	1	2	3	4	5	6	7
				Bitcoin			
AR1	42.49	42.28	41.54	41.62	41.55	41.42	41.12
KS	1.52	12.25	1.84	**0.96**	1.06	1.06	1.12
KS-noreg	2.07	7.03	1.61	1.02	1.01	1.03	1.01
DMA-noreg	**0.96**	**0.98**	0.99	1.00	1.00	0.99	1.00
DMS-noreg	**0.97**	1.00	0.99	1.03	1.02	1.01	1.02
DMA	**0.97**	**0.97**	1.01	1.04	1.02	1.02	1.13
DMS	1.01	1.02	1.06	1.06	1.02	1.05	1.15
				Litecoin			
AR1	134.27	132.88	133.05	133.43	133.60	133.25	131.71
KS	1.02	1.17	7.64	1.01	1.17	1.11	1.09
KS-noreg	**0.96**	1.03	1.88	1.00	1.01	1.02	1.00
DMA-noreg	0.99	1.03	1.04	1.02	1.02	1.05	1.05
DMS-noreg	1.01	1.04	1.06	1.04	1.02	1.04	1.04
DMA	0.98	1.03	1.09	1.11	1.03	1.06	1.15
DMS	1.00	1.07	1.11	1.11	1.04	1.09	1.22

Table: Mean squared error (MSE), computed over the forecast horizon. Results are reported relative to the benchmark specification (AR1) for which the absolute score is reported. Models' description is reported in Table 1.

Univariate forecasting results: MSE

h	1	2	3	4	5	6	7
				Ripple			
AR1	224.02	221.31	222.02	221.13	218.93	219.62	219.45
KS	1.11	1.24	1.27	1.10	1.76	1.21	2.01
KS-noreg	1.03	1.04	1.02	1.01	1.08	1.00	1.02
DMA-noreg	0.99	1.03	1.03	1.05	1.06	1.00	1.05
DMS-noreg	1.02	1.03	1.04	1.07	1.08	1.03	1.08
DMA	1.20	1.03	1.22	1.25	1.10	1.18	1.17
DMS	1.27	1.05	1.22	1.26	1.11	1.21	1.21
				Ethereum			
AR1	180.57	174.99	175.61	175.56	175.79	175.90	174.08
KS	1.05	12.72	1.09	1.01	1.02	1.67	1.09
KS-noreg	1.01	3.40	1.02	1.00	1.00	1.01	1.00
DMA-noreg	**0.96**	1.00	0.98	**0.97**	**0.98**	0.98	1.00
DMS-noreg	0.98	1.01	1.01	0.99	0.99	1.00	1.01
DMA	0.97	1.01	1.03	1.01	1.04	1.04	1.04
DMS	1.02	1.04	1.08	1.05	1.05	1.09	1.04

Table: Values in **bold**, indicate rejection of the null hypothesis of Equal Predictive Ability between each model and the benchmark according to the Diebold–Mariano test at the 5% confidence level. Gray cells indicate those models that belong to the Superior Set of Models

Multivariate models

- *Constant parameter VARs*: We have four-variate VARs with three lags selected using BIC. We apply two version of them: a frequentist VAR estimated using OLS and a Bayesian VAR (BVAR) as in Koop and Korobilis (2013).

- *Time–varying parameter specifications*: Time-varying parameter vector autoregression model (TVP-VAR) as described in Koop and Korobilis (2013).

- The evidence is striking and results are almost opposite to the univariate case in terms of forecast metrics: no model provides economic gains when point forecasting cryptocurrencies; several models provide large gains when density forecasting.

- Focusing on MSE: simpler constant-parameter VAR and BVAR specifications are very imprecise at short horizons with losses up to 20%, but they perform more similar to the AR(1) at longer horizons with mild improvements for BVAR.

- Focusing on predictive log score: Time–varying parameter specifications provide statistically superior forecasts relative to the benchmark at almost all horizons.

References I

Torben G Andersen, Tim Bollerslev, Francis X Diebold, and Paul Labys. The distribution of realized exchange rate volatility. *Journal of the American statistical association*, 96:42–55, 2001.

Torben G Andersen, Tim Bollerslev, Peter F Christoffersen, and Francis X Diebold. Volatility and correlation forecasting. *Handbook of economic forecasting*, 1:777–878, 2006.

M. Bernardi and L. Catania. Portfolio Optimisation Under Flexible Dynamic Dependence Modelling. *ArXiv e-prints*, 2016.

Bloomberg. Nasdaq Plans to Introduce Bitcoin Futures. 2017.

Tim Bollerslev. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 31:307–327, 1986.

L. Catania, S. Grassi, and F. Ravazzolo. Predicting the volatility of cryptocurrency time–series. 2018a.

Leopoldo Catania and Stefano Grassi. Modelling crypto-currencies financial time-series. 2017.

Leopoldo Catania and Nima Nonejad. Dynamic model averaging for practitioners in economics and finance: The edma package. *Journal of Statistical Software, Articles*, 84(11):1–39, 2018. ISSN 1548-7660. doi: 10.18637/jss.v084.i11. URL `https://www.jstatsoft.org/v084/i11`.

Leopoldo Catania, Stefano Grassi, and Francesco Ravazzolo. Forecasting cryptocurrencies financial time series. 2018b.

Jeffrey Chu, Saralees Nadarajah, and Stephen Chan. Statistical Analysis of the Exchange Rate of Bitcoin. *PloS one*, 10:1–27, 2015.

Fulvio Corsi. A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*, 7:174–196, 2009.

David R Cox. Statistical analysis of time series: Some recent developments [with discussion and reply]. *Scandinavian Journal of Statistics*, 8(2):93–115, 1981.

Burton Crane. *The Sophisticated Investor: A Guide to Stock-market Profits*. Simon and Schuster, 1959.

Drew Creal, Siem Jan Koopman, and Andre' Lucas. Generalized Autoregressive Score Models with Applications. *Journal of Applied Econometrics*, 28:777–795, 2013.

Thomas Dangl and Michael Halling. Predictive Regressions with Time–Varying Coefficients. *Journal of Financial Economics*, 106: 157–181, 2012.

References II

Chicago Mercantile Exchange. CME Group Announces Launch of Bitcoin Futures. 2017.

Forbes. Emerging Applications For Blockchain. 2017.

Gloria González-Rivera, Tae-Hwy Lee, and Santosh Mishra. Forecasting volatility: A reality check based on option pricing, utility function, value-at-risk, and predictive likelihood. *International Journal of forecasting*, 20(4):629–645, 2004.

Bruce E Hansen. Autoregressive Conditional Density Estimation. *International Economic Review*, 35:705–730, 1994.

Andrew C Harvey. *Dynamic Models for Volatility and Heavy Tails: with Applications to Financial and Economic Time Series*, volume 52. Cambridge University Press, 2013.

A Hencic and C Gourieroux. Noncausal Autoregressive Model in Application to Bitcoin/USD Exchange Rate. *Proceedings of the 7th Financial Risks International Forum*, pages 1–25, 2014.

Gary Koop and Dimitris Korobilis. Large Time-Varying Parameter VARs. *Journal of Econometrics*, 177:185–198, 2013.

André Lucas and Xin Zhang. Score-driven Exponentially Weighted Moving Averages and Value-at-Risk Forecasting. *International Journal of Forecasting*, 32:293–302, 2016.

André Lucas, Bernd Schwaab, and Xin Zhang. Conditional euro area sovereign default risk. *Journal of Business & Economic Statistics*, 32(2):271–284, 2014. doi: 10.1080/07350015.2013.873540. URL `http://dx.doi.org/10.1080/07350015.2013.873540`.

André Lucas, Bernd Schwaab, and Xin Zhang. Modeling financial sector joint tail risk in the euro area. *Journal of Applied Econometrics*, 32(1):171–191, 2017. ISSN 1099-1255. doi: 10.1002/jae.2518. URL `http://dx.doi.org/10.1002/jae.2518`. jae.2518.

S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. *Available at: https://Bitcoin.org/Bitcoin.pdf*, 2009.

Wilfredo Palma. *Long-memory Time Series: Theory and Methods*, volume 662. John Wiley & Sons, 2007.

S Sapuric and A Kokkinaki. Bitcoin is Volatile! Isn't That Right? *Business Information Systems Workshops*, pages 255–265, 2014.