Speculative markets: bubbles or balloons?

Christian Hafner, Andrew Harvey (ach34@cam.ac.uk) and Linqi Wang

CORE, University of Cambridge and QMUL

Quant Strats, London, October 15th, 2025

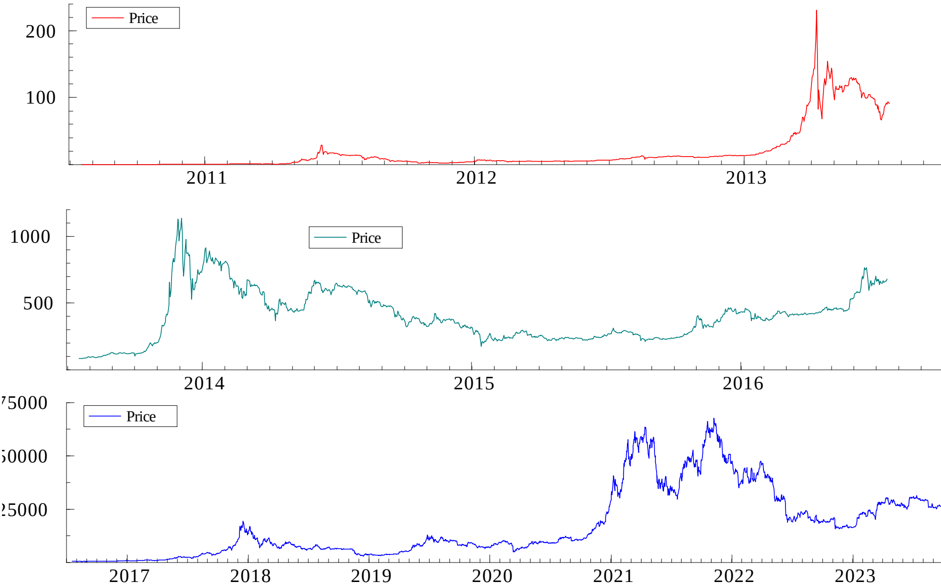
Price
200
100
2011
2012
2013
Price
1000
500
2014
2015
2016
Price
75000
50000
25000
2017
2018
2019
2020
2021
2022
2023


LPrice
10
8
6
4
2
0
-2
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023


LBTC
11.0
10.5
10.0
9.5
9.0
8.5
8.0
2018
2019
2020
2021
2022
2023
2024
20

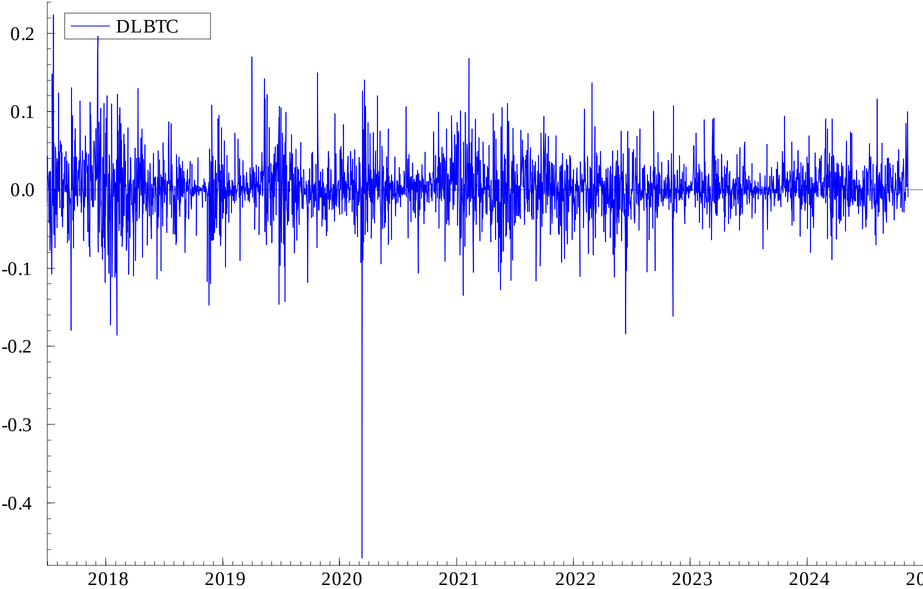
DLBTC
0.2
0.1
0.0
-0.1
-0.2
-0.3
-0.4
2018
2019
2020
2021
2022
2023
2024

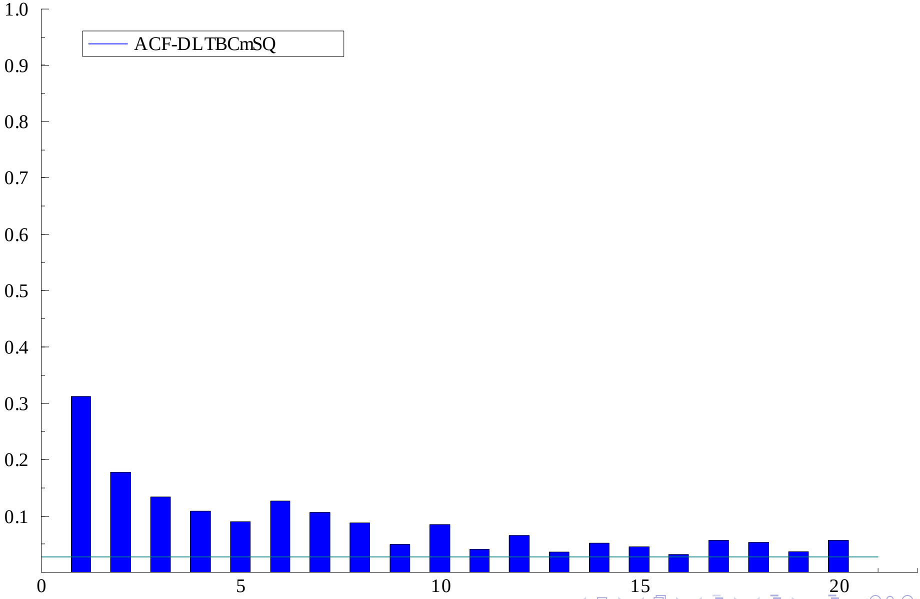
ACF-DLTBCmSQ
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
5
10
15
20

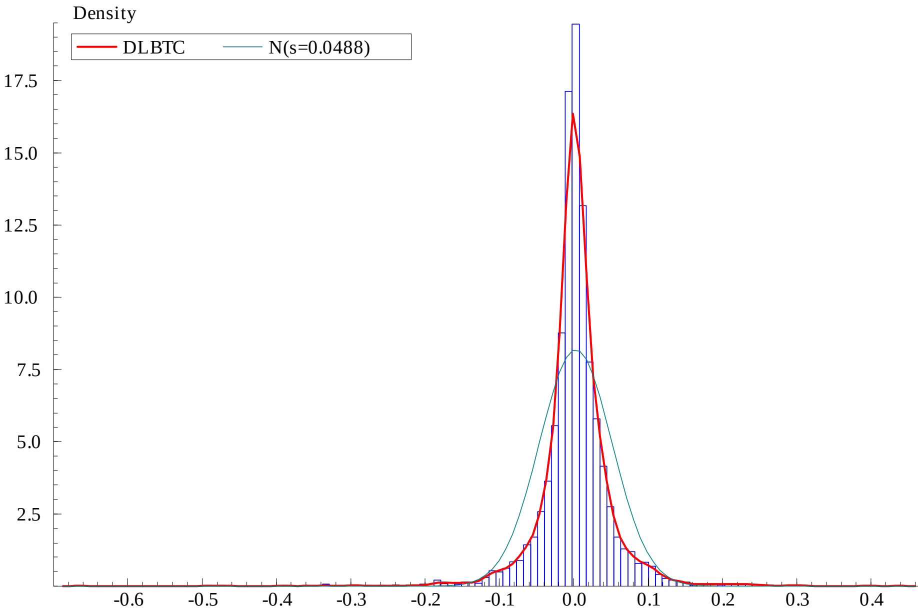
Density
DLBTC
N(s=0.0488)
17.5
15.0
12.5
10.0
7.5
5.0
2.5
-0.6
-0.5
-0.4
-0.3
-0.2
-0.1
0.0
0.1
0.2
0.3
0.4

Speculative markets: bubbles and balloons

Market with imperfect information, so returns are not a MD.
Heteroscedasticity and non-normality.
A classic bubble will normally show a steady increase until it bursts. A balloon does the opposite (unless a pin is stuck in it). FOMO.
Repeated bubbles and balloons. Aim is to track movements for nowcasting and forecasting rather than test for bubbles or balloons.
**
Modeling Prices from Speculative Markets: Bursting Bubbles or Deflating Balloons?
Christian Hafner, Andrew Harvey. Linqi Wang. 2025 Cambridge Working Papers in Economics, CWPE 2523

Introduction to Score driven models

A unified and comprehensive theory for a class of nonlinear time series models in which the dynamics of a changing parameter, such as location or scale, is driven by the score of the conditional distribution.

Extensions to multivariate time series. Correlation or association may change over time. Time-varying copulas.

Harvey, A.C. **Dynamic models for volatility and heavy tails**. CUP 2013

Creal et al (2011, JBES, 2013, JAE).

Harvey, A (2022) Score-driven time series models. *Annual Review of Statistics and Its Application*, 9, 321-42. doi: 10.1146/annurev-statistics-040120-021023

Statistical theory has developed in the last 10 years. Asymptotic theory - Blasques et al (2023, JE). BBKL.

Invertibility.

Patton, A. and Y. Simsek (2025) Generalized Autoregressive Score Trees and Forests.

Outline of talk

1) Location models.
2) Heteroscedasticity. The case for EGARCH.
3) Dynamic location and scale.
4) Explosive models: bubbles and balloons
5) Dynamic tail index

Exponential Generalized Beta of the Second Kind (EGB2)

If x is distributed as $GB2(\alpha, \nu, \xi, \varsigma)$ and $y = \ln x$, the PDF of the $EGB2$ variate y is

$$f(y; \mu, \nu, \xi, \varsigma) = \frac{\nu \exp\{\xi(y-\mu)\nu\}}{B(\xi, \varsigma)(1 + \exp\{(y-\mu)\nu\})^{\xi+\varsigma}}.$$

What was the logarithm of scale in GB2 now becomes location in EGB2.
Thin tailed. All moments exist.
When $\xi = \varsigma$, the distribution is symmetric; for $\xi = \varsigma = 1$ it is a *logistic* distribution and when $\xi = \varsigma \to \infty$ it tends to a *normal* distribution.
For $\xi = \varsigma = 0$, the distribution is *double exponential or Laplace.*
Figure below is Realised Variance (RV) for FTSE and its logarithm.

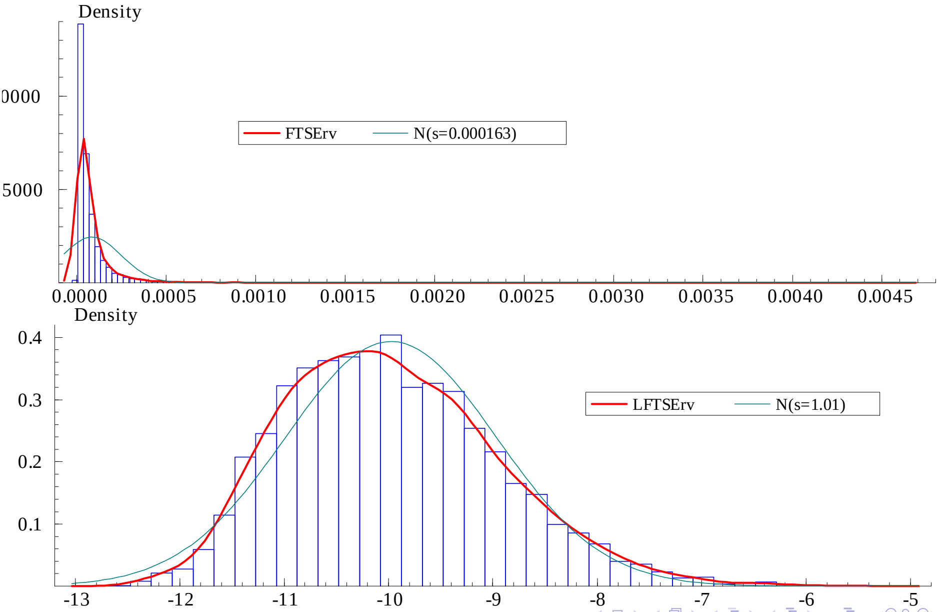
Density
0000
5000
FTSErv
N(s=0.000163)
0.0000
0.0005
0.0010
0.0015
0.0020
0.0025
0.0030
0.0035
0.0040
0.0045
Density
0.4
0.3
0.2
0.1
LFTSErv
N(s=1.01)
-13
-12
-11
-10
-9
-8
-7
-6
-5

Generalized t-distribution

Student-t, normal and GED are special cases of generalized-t. Two shape parameters - tail index (df for Student t) and a second parameter that is 2 for Student t.

Absolute value of gen-t is GB2.

The flexibility of Gen-t goes a long way towards meeting the objection that parametric models are too restrictive and hence vulnerable to misspecification -see McDonald and Newey (Econometrica, 1987).

Harvey and Lange (2017, JTSA). Volatility Modeling with a Generalized t-distribution.

Dynamic location: Generalized t-distribution

The homoscedastic location-t model is

$$y_t = \mu + \mu_{t|t-1} + \varepsilon_t \exp \lambda, \quad t = 1, \ldots, T, \tag{1}$$

where $SD = \sigma = \exp \lambda$ and

$$\mu_{t+1|t} = \phi \mu_{t|t-1} + \kappa u_t^{\mu}, \qquad |\phi| < 1, \quad \kappa > 0 \tag{2}$$

When y_t has a generalized–t distribution with conditional mean $\mu + \mu_{t|t-1}$ and scale $\exp \lambda$, the conditional score is

$$\frac{\partial \ln f}{\partial \mu_{t|t-1}} = e^{-\lambda} \frac{\eta + 1}{\eta} (1 - b_t) \left| \varepsilon_t \right|^{\upsilon - 1} sgn(\varepsilon_t),$$

where $\varepsilon_t = (y_t - \mu - \mu_{t|t-1}) \exp(-\lambda)$ and

$$b_t = \frac{\varepsilon_t^2 / \nu}{1 + \varepsilon_t^2 / \nu}, \qquad 0 \leq b_t \leq 1,$$

Dividing by $I_{\mu\mu}$ but ignoring constants gives standardized score

$$u_t^{\mu} = e^{\lambda} (1 - b_t) \left| \varepsilon_t \right|^{\upsilon - 1} sgn(\varepsilon_t). \tag{3}$$

Dynamic location model: EGB2

The score function with respect to location in EGB2 gives a gentle form of Winsorizing.
For a logistic distribution, $\xi = \varsigma = 1$,

$$u_t = \frac{3\sigma}{h} \frac{e^{(y_t - \mu_{t|t-1})h/\sigma} - 1}{e^{(y_t - \mu_{t|t-1})h/\sigma} + 1} \qquad (4)$$

with $h = 1.81$.
The less restrictive invertibility requirements for the EGB2 score make it an attractive alternative to the true score for a gen-t in a QSD model.

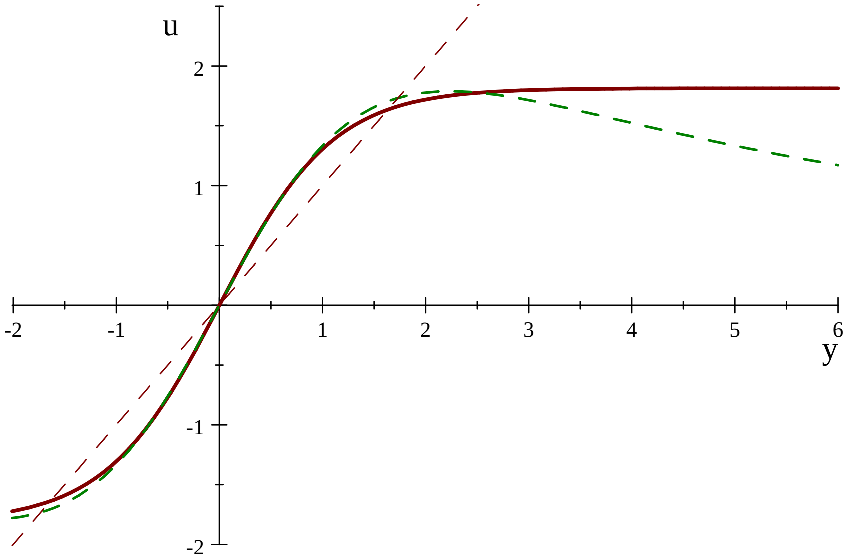

Figure: Standardized score functions for t_7 and logistic with $\sigma = 1$, together with the (linear) score for a normal distribution.

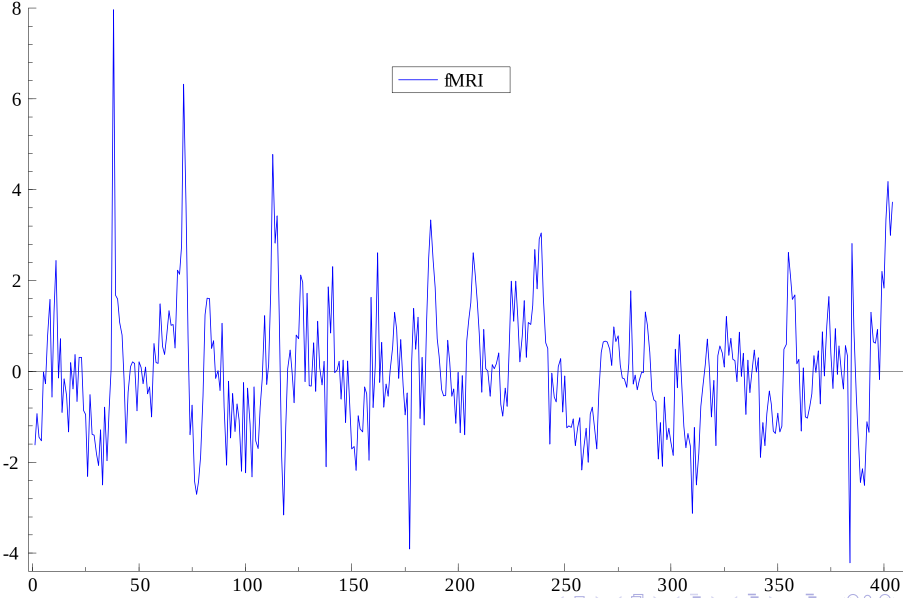
fMRI
8
6
4
2
0
-2
-4
0
50
100
150
200
250
300
350
400

Example: fMRI

MRI scans.

QSD with conditional Student-t and a logistic score gives a better fit than SD-t or SD-EGB2.

Degrees of freedom is 5.65.

$\phi = 0.775, \quad \kappa = 0.673$

GARCH-t and Beta-t-EGARCH

Stock returns are non-normal
Assume that z_t has a Student t_ν-distribution, where ν denotes degrees of freedom - GARCH-t model.
The t-distribution is employed in the predictive distribution of returns and used as the basis for maximum likelihood (ML) estimation of the parameters, but it is not acknowledged in the design of the equation for the conditional variance.
The specification of the $\sigma^2_{t|t-1}$ *as a linear combination of squared observations is taken for granted*, but the consequences are that $\sigma^2_{t|t-1}$ responds too much to extreme observations and the effect is slow to dissipate.

Exponential DCS Volatility Models

$$y_t = \varepsilon_t \exp \lambda_{t|t-1}, \quad t = 1, ..., T,$$

where the serially independent, zero mean variable ε_t has a $t_\nu-$distribution with degrees of freedom, $\nu > 0$, and the filtering equation for the log of scale is score standardized by dividing by information quantity is

$$\lambda_{t+1|t} = \delta + \phi\lambda_{t|t-1} + \kappa u_t.$$

The conditional score is standardized by dividing by information quantity to give

$$u_t = \frac{(\nu+3)}{2\nu}\left[\frac{(\nu+1)y_t^2}{\exp(\lambda_{t|t-1}) + y_t^2} - 1\right], \quad -1 \leq u_t \leq \nu, \quad \nu > 0$$

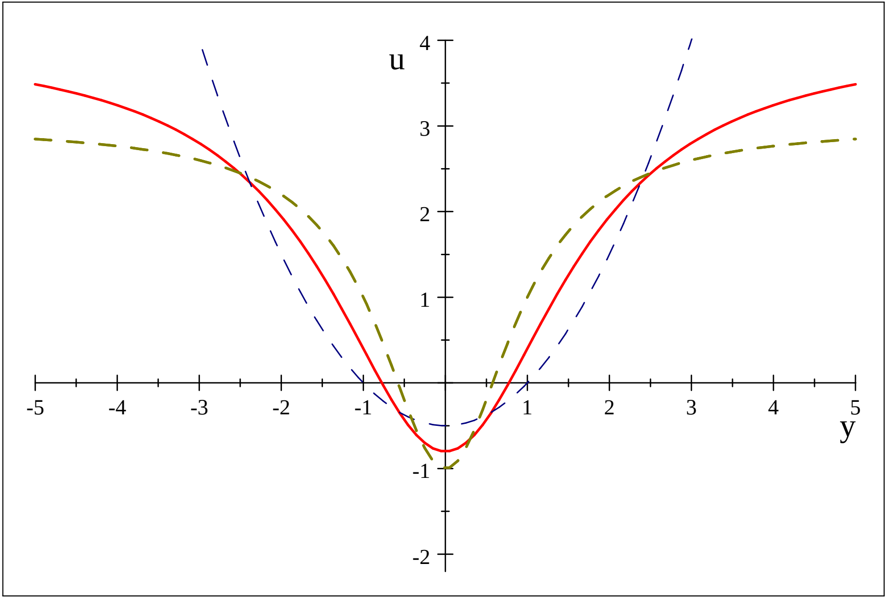

Standardized scores for t-distributions with unit variance: $\nu = 3$ is thick brown; $\nu = 5$ is thin red and Gaussian is dashed blue

Beta-t-EGARCH

The conditional score may be expressed as

$$\frac{\partial \ln f}{\partial \lambda_{t|t-1}} = (\nu + 1)b_t - 1, \qquad 0 < \nu < \infty,$$

where

$$b_t = \frac{\varepsilon_t^2/\nu}{1 + \varepsilon_t^2/\nu}, \qquad 0 \leq b_t \leq 1.$$

is distributed as $Beta(1/2, \nu/2)$. The $u_t's$ are IID and the model is stationary.

Moments can be calculated. Harvey (2013, ch 4).

Invertibility not a problem.

Leverage and Two Components

The leverage effect in volatility models suggests that volatility rises when the asset price falls. This asymmetry can be captured by an EGARCH model by modifying the dynamic equation to

$$\lambda_{t+1|t} = \omega(1-\phi) + \phi\lambda_{t|t-1} + \kappa u_t + \kappa^* sgn(-(y_t - \mu))(u_t + 1),$$

where κ^* is a new parameter which, because the negative of the sign of the return is taken, will normally be positive.
When two component volatility models are employed, it is often found that the leverage effect is confined to the short-term component.
In this case, the evolution of the long-run component will be less susceptible to the influence of strongly negative returns.

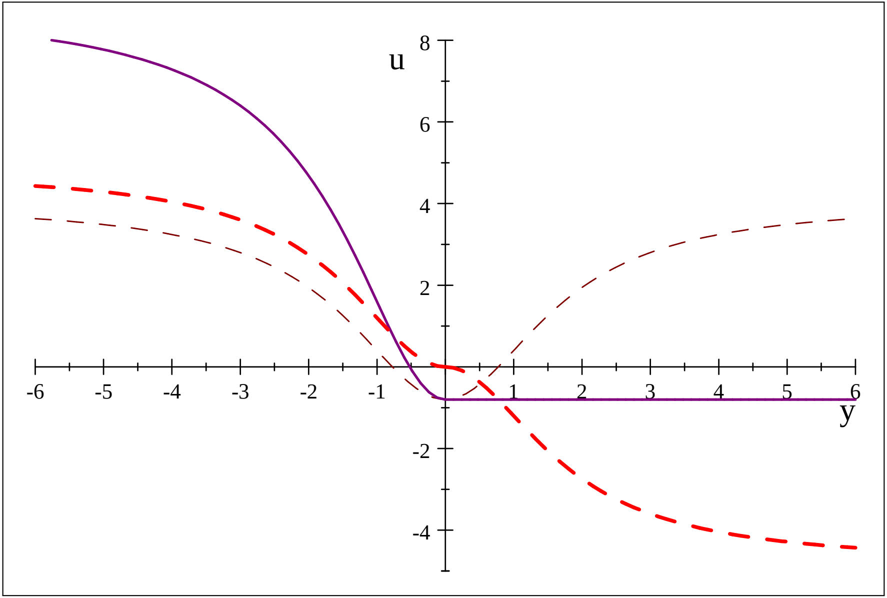

Figure: Impact curves for t_5 with unit variance. Thin dashes is score; purple line has $\kappa = \kappa^* = 1$ and thick red dashes has $\kappa = 0$ and $\kappa^* = 1$.

Dynamic location and scale

A SD or QSD model for dynamic location with an additive EGARCH error may be set up for a generalized–t distribution:

$$\begin{aligned} y_t &= \mu + \mu_{t|t-1} + \varepsilon_t \exp \lambda_{t|t-1} \\ \mu_{t+1|t} &= \phi_\mu \mu_{t|t-1} + \kappa_\mu u_{\mu t}, \\ \lambda_{t+1|t} &= \delta + \phi_\lambda \lambda_{t|t-1} + \kappa_\lambda u_{\lambda t} \end{aligned}$$

where $u_{\mu t}$ is the conditional score divided by the information quantity at time t is

$$A1: \qquad u_t^\mu = e^{\lambda_{t|t-1}} \frac{\eta + 1}{\eta} (1 - b_t) \left| \varepsilon_t \right|^{\upsilon - 1} sgn(\varepsilon_t),$$

where $\varepsilon_t = (y_t - \mu - \mu_{t|t-1}) \exp(-\lambda_{t|t-1})$ and

$$b_t = \frac{\varepsilon_t^2 / \nu}{1 + \varepsilon_t^2 / \nu}, \qquad 0 \leq b_t \leq 1.$$

Dynamic location and scale

Or quasi-score

$$A1/QSD: \qquad u_t^{\mu} = \frac{3\exp\lambda_{t|t-1}}{h}\frac{e^{h\varepsilon_t}-1}{e^{h\varepsilon_t}+1}$$

where $h = \sqrt{2\psi'(1)} = \pi\sqrt{1/3} = 1.814$. [$\exp\lambda_{t|t-1}$ redefined as $\sigma_{t|t-1}$ rather than scale]
The presence of $e^{\lambda_{t|t-1}}$ as the leading term in u_t^{μ} means that the score is no longer IID.
Dividing by the square root of the information gives

$$A2/QSD: \qquad u_t^{\mu} = \sqrt{3}\frac{e^{h\varepsilon_t}-1}{e^{h\varepsilon_t}+1}. \qquad -1.66 \le u_t^{\mu} \le 1.66.$$

Dynamic location and scale

Location and scale may be combined multiplicatively to give

$$y_t = \mu + \mu_{t|t-1} \exp \lambda_{t|t-1} + \varepsilon_t \exp \lambda_{t|t-1},$$

The heteroscedasticity corrected observations are implicitly modeled as

$$(y_t - \mu) \exp(-\lambda_{t|t-1}) = \mu_{t|t-1} + \varepsilon_t,$$

where ε_t has unit standard deviation.
The information quantity for μ does not depend on λ, suggesting that we use the raw score. For QSD

$$M/QSD: \qquad u_t^{\mu} = h \frac{e^{h\varepsilon_t} - 1}{e^{h\varepsilon_t} + 1} = 1.81 \frac{e^{1.81\varepsilon_t} - 1}{e^{1.81\varepsilon_t} + 1}.$$

where

$$\varepsilon_t = (y_t - \mu) \exp(-\lambda_{t|t-1}) - \mu_{t|t-1}.$$

Dynamic location and scale

The score for $\lambda_{t|t-1}$ in the M model is

$$u_t^{\lambda} = (\eta + 1)\, b_t - 1 + \frac{\eta + 1}{\eta}(1 - b_t)\, |\varepsilon_t|^{\upsilon - 1}\, \mu_{t|t-1} sgn(\varepsilon_t)$$

Note that u_t^{μ} and u_t^{λ} are uncorrelated if $\mu = 0$. The score for scale Winsorizes.

Bitcoin

Y_t = logarithm of price. $y_t = \Delta Y_t$. Model for returns, y_t.
1) Two component gen-t EGARCH with leverage works well. υ around 1.2.
2) M and M/QSD better (higher $\ln L$) than SD and QSD. Furthermore the μ is not significant for M and M/QSD so BIC shows them to be even better.
3) Although leverage effects are significant, the BIC is smaller when they are dropped.

Explosive models

A run of positive shocks to returns means $\mu_{t|t-1}$ increases and Y_t increases more rapidly. This feeds speculation leading to bigger effects on location than can be produced by the standard model. A dynamic equation for $\ln \mu_{t|t-1}$ might mimic this behaviour by enhancing the effect of a high level of returns. Sharp downward movements may similarly enhanced.
Nonlinear link function

$$\begin{aligned} \mu_{t|t-1} &= \alpha_1[\exp(\beta_{t|t-1}) - 1] - \alpha_2[\exp(-\beta_{t|t-1}) - 1], \\ &= \alpha_2 - \alpha_1 + \alpha_1 \exp(\beta_{t|t-1}) - \alpha_2 \exp(-\beta_{t|t-1}), \qquad \alpha_1, \alpha_2 \geq 0, \end{aligned}$$

where

$$\beta_{t+1|t} = \phi^\beta \beta_{t|t-1} + \kappa\, u_t^\beta, \qquad \kappa > 0,\ 0 \leq \phi^\beta \leq 1$$

with u_t^β as in A2 or M model.
When $\beta_{t|t-1} = 0$, $\exp(\beta_{t|t-1}) = \exp(-\beta_{t|t-1}) = 1$ which is the reason for subtracting one.

Explosive models

When $\alpha_2 = 0$, the response is as shown by the blue dashed line in figure. Large positive movements in $\beta_{t|t-1}$ are amplified, leading to a faster rise in Y_t. (Balloon) On the other hand, when $\alpha_1 = 0$, the opposite happens. (Bubble)
When $\alpha_1 = \alpha_2 = \alpha$, the model is

$$y_t = \mu + 2\alpha \sinh(\beta_{t|t-1}) + \varepsilon_t \exp\left(\lambda_{t|t-1}\right)$$

where $\sinh(x) = (e^x - e^{-x})/2$. This is the solid line in Figure.
Large movements of $\beta_{t|t-1}$ in either direction are amplified, leading to pronounced rises and falls in Y_t.

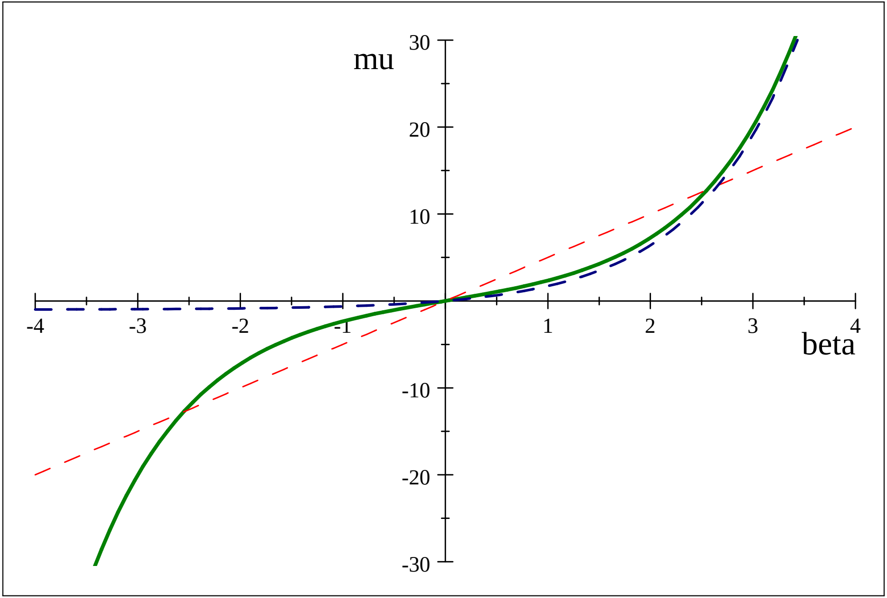

Exponential, sinh and linear response

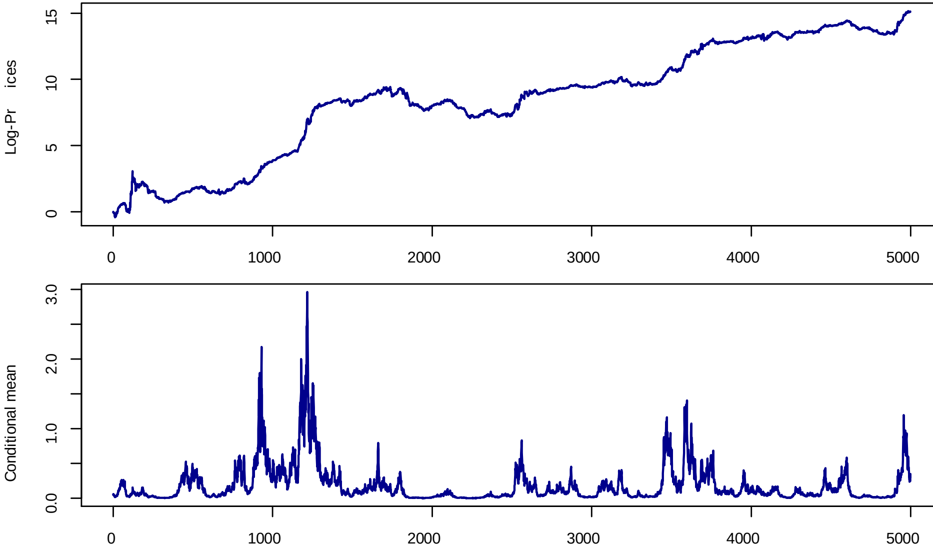

Figure: Simulated balloon

Explosive models

The score with respect to $\beta_{t|t-1}$ is

$$\frac{\partial \ln f_t}{\partial \beta_{t|t-1}} = u_t^{\mu} \left\{ \alpha_1 \exp\left(\beta_{t|t-1}\right) + \alpha_2 \exp\left(-\beta_{t|t-1}\right) \right\}$$

Dividing by the square root of the information quantity yields

$$u_t^{\beta} = u_t^{\mu},$$

where u_t^{μ} is the score or quasi-score.

LINEX*

A modification is to combine the exponential function with a linear response; cf LINEX loss function. The result is

$$\mu_{t|t-1} = \alpha_2 - \alpha_1 + \alpha_3 \beta_{t|t-1} + \alpha_1 \exp(\beta_{t|t-1}) - \alpha_2 \exp(-\beta_{t|t-1})$$

The blue dashed line in the previous figure is now the solid line produced with $\alpha_1 = 1$, $\alpha_2 = 0$ and $\alpha_3 = 2$.

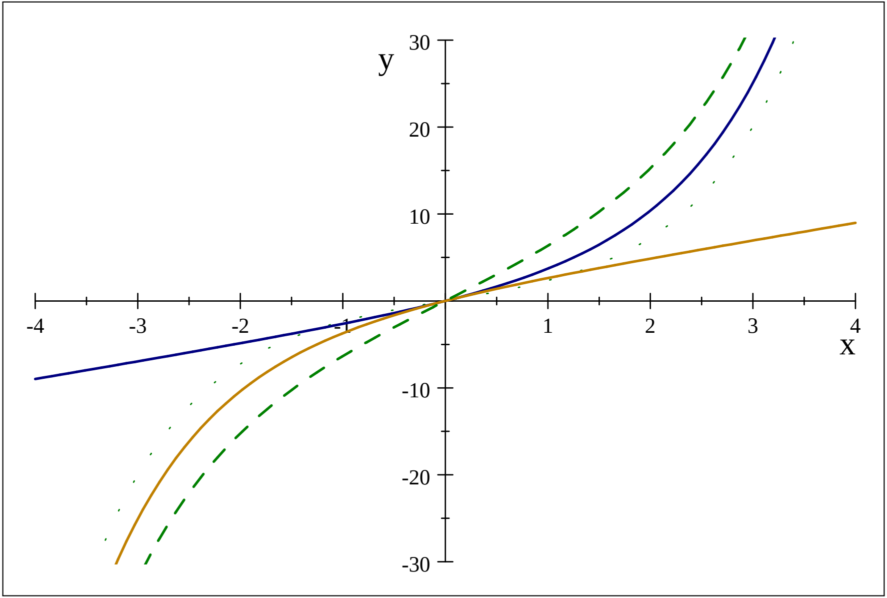

Exponential combined with linear response

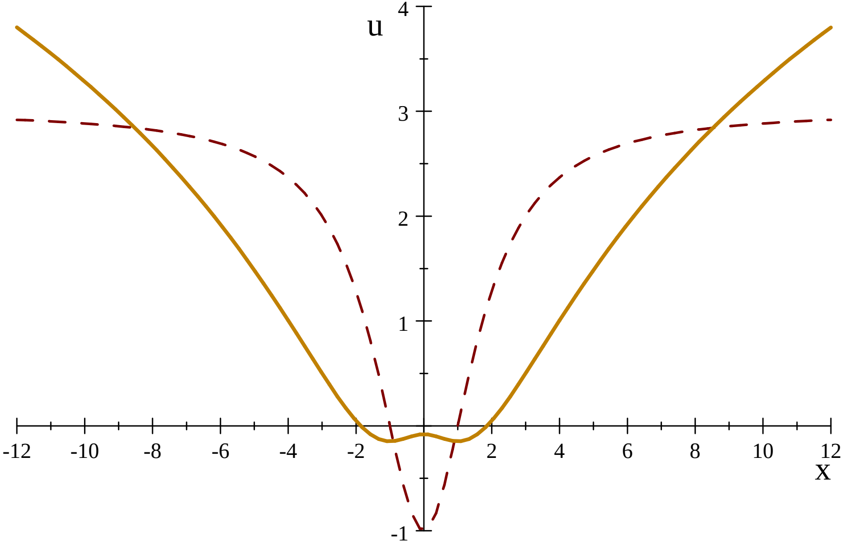

Figure: Score for $u_t = -\ln \nu$ with the score for scale plotted for a standardized t-variate ($\nu = 3$).


tail index (d.o.f.) 2010-07-19 / 2024-11-12
5.0
4.5
4.0
3.5
3.0

Conclusions

1) Location - QSD works well for generalized-t. Simple and removes concern about invertibility.
2) EGARCH - usual conclusions, ie gen-t, two components with their own leverage.
In terms of LL and AIC, models with leverage are better but with BIC, without leverage is better.
3) Location and dispersion -additive and multiplicative forms. Multiplicative best.
4) Explosive models. Bitcoin exhibits repeated bubbles and/or balloons. Balloons seem to dominate.
5) Change in tail index. From 3 to 5.

On State-Space Models in Time's Domain

Within the bounds where data shifts unseen,
A hidden state evolves through time's embrace.
Its measured form, though partial and routine,
Reflects the past, yet leaves a masked trace.
The transition guides how states will flow,
A process steeped in noise and chance's call.
While observation hints at what we know,
It veils the truth behind a measured wall.
Through Kalman's gaze, we seek to filter clear,
To estimate the path from noise unkind.
With each new step, prediction draws it near,
A latent truth within the data mined.
Thus, state-space frames what time may yet reveal,
A structured dance where past and now congeal.