IRTG 1792 Discussion Paper 2018-038

Tail-Risk Protection Trading Strategies

Natalie Packham *
Jochen Papenbrock *2
Peter Schwendner *3
Fabian Woebbeking *4

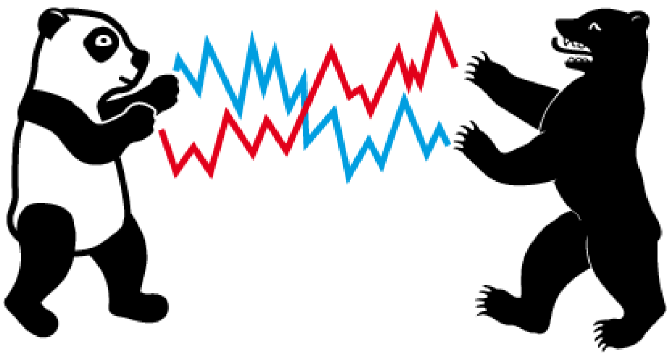

* Berlin School of Economics and Law, Germany
*2 Firamis and PPI AG, Germany
*3 Zurich University of Applied Sciences, Switzerland
*4 Goethe University Frankfurt, Germany

This research was supported by the Deutsche Forschungsgemeinschaft through the International Research Training Group 1792 "High Dimensional Nonstationary Time Series".

http://irtg1792.hu-berlin.de
ISSN 2568-5619

International Research Training Group 1792

TAIL-RISK PROTECTION TRADING STRATEGIES

N. Packham, J. Papenbrock, P. Schwendner and F. Woebbeking

AUGUST 15, 2016

Starting from well-known empirical stylised facts of financial time series, we develop dynamic portfolio protection trading strategies based on econometric methods. As a criterion for riskiness we consider the evolution of the value-at-risk spread from a GARCH model with normal innovations relative to a GARCH model with generalised innovations. These generalised innovations may for example follow a Student t, a generalised hyperbolic (GH), an alpha-stable or a Generalised Pareto (GPD) distribution. Our results indicate that the GPD distribution provides the strongest signals for avoiding tail risks. This is not surprising as the GPD distribution arises as a limit of tail behaviour in extreme value theory and therefore is especially suited to deal with tail risks. Out-of-sample backtests on 11 years of DAX futures data, indicate that the dynamic tail-risk protection strategy effectively reduces the tail risk while outperforming traditional portfolio protection strategies. The results are further validated by calculating the statistical significance of the results obtained using bootstrap methods. A number of robustness tests including application to other assets further underline the effectiveness of the strategy. Finally, by empirically testing for second order stochastic dominance, we find that risk averse investors would be willing to pay a positive premium to move from a static buy-and-hold investment in the DAX future to the tail-risk protection strategy.

Keywords: tail-risk protection, portfolio protection, extreme events, tail distributions.

JEL Classification: C15, G11, G17.

1. INTRODUCTION

Starting from well-known empirical stylised facts of financial time series, we develop dynamic portfolio protection trading strategies based on econometric methods. The principal idea is to investigate whether the information present in financial time series can be used to detect risk-build up and to what extent this can be put to use to protect against large downturns. The main motivation for such a tail-risk protection trading strategy is that, despite being unable to guarantee a hedge against arbitrary extreme events, it may still be effective, while avoiding the high cost associated with traditional portfolio protection.

Especially in times of low equity risk premia, institutional and private investors may be underinvested in equity due to investment constraints, downside tail risk, value-at-risk limits or behavioural effects like loss aversion. The literature identifies a particular aversion against tail risks among investors (Bollerslev and Todorov, 2011). This in turn leads to high prices of put options as reflected in the downward slope of the equity implied volatility curve as a function of the strike price (Kozhan *et al.*, 2013), and the upward slope of VIX futures prices as a function of expiry (Zhang *et al.*, 2010; Luo and Zhang, 2012; Eraker and Wu, 2014). Both market patterns are more pronounced than would be reasonable by the realised return distributions and may therefore curtail a successful long-term static hedge position against tail risks.

For banks, an increased focus on tail risk will be enforced by the new "Minimum capital requirements for market risk" (BIS, 2016) that resulted from the "Fundamental Review of the

Natalie Packham, Berlin School of Economics and Law, natalie.packham@hwr-berlin.de
Jochen Papenbrock, Firamis and PPI AG, jp@firamis.de
Peter Schwendner, Zurich University of Applied Sciences, peter.schwendner@zhaw.ch
Fabian Woebbeking, Goethe University Frankfurt, woebbeking@finance.uni-frankfurt.de
The authors would like to thank an anonymous referee for his or her valuable suggestions that helped to improve the paper in several respects.

Trading Book" by the Basel Committee on Banking Supervision after the global financial crisis. A quantitative impact study (BIS, 2015) suggests an increased capital requirement of about 40% averaged across 78 global banks. The stricter requirements will likely lead to a weaker passive risk absorption capacity for banks and encourage them to actively manage those risks.

More specifically, we develop dynamic trading strategies that aim at protecting against large downturns by taking into account the time-variation and dynamics of distributional parameters of financial time series. The principal idea is to exploit the information given by so-called "empirical stylised facts" of financial time series (Cont, 2001; McNeil *et al.*, 2005; Engle and Patton, 2001; Thurner *et al.*, 2012; Godin, 2015). First, accounting for the time-dependent dynamics of distributional parameters via a GARCH process allows to incorporate volatility clustering and autoregressive behaviour in volatility, both of which are well-documented stylised facts. Second, by fitting the GARCH innovations to flexible distribution families incorporating both normal and extreme behaviour allows to determine whether, in a given time period, extreme events are more likely to occur than suggested by normal innovations.

In our setup, the key input is the spread between value-at-risk (VaR) from a GARCH model with innovations following a generalised Pareto distribution (GPD), which is a distribution focussing on extreme risks, and a GARCH model with normally distributed innovations. We also considered Student t, skewed normal, skewed Student t, generalised hyperbolic (GH) and α-stable distributions as alternatives for the innovations, but found little explanatory power for extreme events. Because of the GARCH component, the magnitude of VaR, when viewed as a process over time, quickly adapts to changes in volatility. The distributional properties of the innovation process on the other hand provide information on skewness, excess kurtosis and in particular on the tail risk in the data. The resulting VaR spread can therefore be used to derive an expectation on the frequency of extreme events; in particular, changes in the VaR spread point to risk changes in the market and, as such, may be taken as signals anticipating the presence of tail risks. The principal idea of the VaR spread is similar to the approach followed by Rachev *et al.* (2010), who focus on α-stable innovations.

The choice of distributions is motivated as follows: The Student t, GH and α-stable distributions incorporate the normal distribution as a special case. The GH distribution is a flexible distribution family comprising light- and medium-tailed distributions, and has been successfully applied for modelling financial time series (e.g. McNeil *et al.*, 2005 and references therein). The Student t distribution is a special limiting case of a GH distribution, spanning the entire range of heavy-tailed behaviour. On the other hand, α-stable distributions are very heavy-tailed with the exception of the normal distribution, which is light-tailed. As such, all distribution families considered incorporate heavy-tailed behaviour, at least in a limiting sense. The GPD arises as the limit of considering only (extreme) outcomes beyond a threshold. Fitting innovations to a GPD therefore provides an entirely different approach by focussing on the extreme behaviour rather than the whole data.

Aside from examining daily returns, we include overnight returns in our analysis. This allows to capture a wider array of different signals in the data. We inspected intraday returns, but they turned out to be less useful for the analysis.

A trading strategy is generated from measuring the risk build-up in terms of the evolution of the VaR spreads of each data set. As already mentioned, the VaR spread with GPD innovations from daily returns and from overnight returns yields the most reliable results. This is not surprising, as the GPD arises as the limit distribution of the tail behaviour in extreme value theory; because of its focus on tail behaviour, it is particularly well suited for a tail-risk protection strategy.

In an extensive out-of-sample backtest covering more than eleven years of DAX future data we find that the tail-risk protection strategy outperforms the DAX future in terms of mean return, standard deviation, Sharpe ratio, worst drawdown and Calmar ratio. In addition, the results are compared with the performance of two traditional protection strategies, *protective put (PP)* and *delta-replicated put (DRP)*. PP is a strategy that eliminates downside risk for a specific asset and period of time at the cost of an option premium, whereas DRP replicates the delta of a hypothetical put with a dynamic trading position in the underlying without actually buying a physical option.

This allows for analysing three entirely different approaches to tail-risk protection: while the strategy developed here is a dynamic econometric approach based on historical data, PP and DRP are both strategies focussing on market data. PP is a semi-static strategy, whereas DRP is a dynamic strategy and its performance is model-dependent and subject to gap risk. For the time period under consideration, we find empirical evidence that the tail-risk protection strategy outperforms PP and DRP, both in terms of performance measures and when taking investors' risk preferences into account. In fact, empirical tests for second-order stochastic dominance reveal that all risk averse investors would prefer tail-risk protection over the benchmark, and there is also a strong indication for preference over PP and DRP. This can also be expressed as follows: a risk averse investor would be willing to pay a positive premium to move from a buy-and-hold investment in the DAX future to the tail-risk protection strategy. This premium is in the range of 40 basis points.

A number of statistical tests reinforce the results. Using bootstrapping techniques, we test the hypothesis that the tail-risk protection strategy does not outperform the DAX future with a random permutation of the trading signals, the PP and DRP strategies, respectively. Further robustness checks are conducted on different markets, asset classes as well as modelling choices.

Our results point to a great potential for tail-risk protection strategies based on time series data and econometric methods. Despite failing for totally surprising extreme losses, a property inherent to historical data, the strategy succeeds in detecting tail-risk build-up. This suggests an alternative to classic risk-protection strategies such as protective put, in particular as these may be too expensive in times of small risk premia to generate adequate or even positive returns.

It is worth mentioning that the strategy introduced here is quite simple yet effective. To achieve a high level of generality, we restrict the setup to capturing stylised facts that are commonly present in financial time series. In a concrete setting, more sophisticated strategies can be developed accounting for asset-class specific refinements (e.g. asymmetric GARCH variants for equity data).

The paper is structured as follows: Section 2 provides a brief review of the relevant literature. The econometric background is given in Section 3, and the tail-risk protection strategy described in Section 4. The empirical results including extensive out-of-sample backtesting and validation against the benchmark and the traditional protection strategies are given in Section 5, robustness checks are presented in Section 6. Section 7 concludes.

2. Background and Literature

The recent period of strong market movements has motivated a line of research focusing on irregular and rare events that occur away from the central region of the return distribution. These extreme or tail events come as a major threat to many investment portfolios and the economy as a whole. The trading year of 2008 alone illustrates the importance of tail-risk protection with several return events that should, under the assumption of independent, normally distributed returns, occur statistically only once in 90 years. Table I illustrates the necessity of incorporating more extreme events in modelling than a normal distribution would suggest. The table specifies *return periods* for different distributions. A return period is the time period in which one would statistically expect to see one occurrence of a particular event. The result shows that fitting a typical financial time series to a Student t-distribution drastically reduces return periods when compared to a normal distribution. A comparison with the number of actual occurrences in our dataset of DAX future returns ($\approx$ 22.75 years) indicates that even the Student-t distribution may not be an appropriate model for capturing extremes, advocating employing distributions from extreme value theory (EVT) instead.

TABLE I

RETURN PERIODS (IN YEARS) FOR DAILY LOSSES GREATER THAN $3-8$ STANDARD DEVIATIONS. THE STUDENT t-DISTRIBUTION WAS FITTED TO DAILY DAX FUTURE RETURNS IN THE PERIOD 5 JAN 1993 THROUGH 26 OCT 2015.

Std. dev.	Normal	Student t ($\nu = 10$)	observed occurrences per year
3	3	1	2
4	126	7	0.5
5	13954	35	0.1
6	4.05438×10^6	156	0.04
7	3.12533×10^9	585	0
8	6.0048×10^{12}	1925	0

The literature deals with the phenomenon of extreme events in different ways. First, it has long been known that observed equity premia are too high to be explained by classical equilibrium models – this is known as the equity premium puzzle introduced by Mehra and Prescott (1985). These models suggest that the risk premium is a monotone function of volatility. However, it has been observed empirically that this relationship may fail to hold. In other words, against the implications of classical theories, there are less volatile stocks that appear to be more profitable. (Baker *et al.*, 2011; Lemperiere *et al.*, 2014; Frazzini and Pedersen, 2014). Bollerslev and Todorov (2011) demonstrate that this failure of classical theories can be explained by extreme events, where the excess risk premium may be ascribed to jump risk compensation. In particular, controlling for the part of risk premium that is attributed to tail events brings the observed risk premium, on average, back in-line with the implications from the classical equilibrium models. Kelly and Jiang (2014) provide evidence that tail risk has large predictive power for aggregate stock market returns over time horizons of a month to several years, the rationale being again that investors are especially averse to tail risk and hence demand higher returns on tail risky portfolios.

Second, Bhansali and Davis (2010a,b) and Bhansali (2014) show that tail protection has a utility not only for defensive purposes but also allows for better long term asset allocation. The special focus on tail events is largely motivated by the breakdown of portfolio diversification in stress scenarios. This is owed to the empirical observation that asset correlation tends to increase during market downturns, especially amid stress and crisis scenarios (Longin and Solnik, 2001).

Third, a number of recent papers incorporate methods from extreme value theory (EVT) and combine them with GARCH volatility modelling in order to analyse the performance of classical risk measures such as value-at-risk and expected shortfall (McNeil and Frey, 2000; Loh and Stoyanov, 2014). There is rising evidence that practitioners already use comparable methods to dynamically hedge against tail risks (Strub, 2013; Madan, 2016).

3. GARCH MODEL WITH GENERALISED INNOVATIONS

The model approach for the underlying financial time series is three-fold: First, a GARCH process is fitted to filter the time-varying volatility. Second, the remaining GARCH innovations are fitted to various distributions. Third, value-at-risk (VaR) forecasts are calculated. The spread between the various VaR's relative to a VaR from normal innovations is used as input for the tail-risk protection strategy.

On a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\in\mathbb{Z}}, \mathbf{P})$ we assume an asset return process $(X_t)_{t\in\mathbb{Z}}$ that follows a GARCH(1,1) process (Bollerslev, 1986)

$$X_t = \sigma_t Z_t$$
$$\sigma_t^2 = \alpha_0 + \alpha_1 X_{t-1}^2 + \beta\sigma_{t-1}^2,$$

where the innovations Z_t, $t = 1, 2, \ldots$ are independent and identically distributed, and $\alpha_0 > 0$, $\alpha_1 \geq 0$ and $\beta \geq 0$. Imposing the condition $\alpha_1 + \beta < 1$ ensures that the process has an unconditional finite variance.

Assuming normally distributed innovations is often not justified, and it is not uncommon to presume t-distributed innovations instead (Bollerslev, 1987). In a general model setup one would choose various distributions for the innovations capturing different features of the data. We considered the skewed normal, (skewed) Student t distribution, generalised hyperbolic (GH) distribution, α-stable distribution and GPD distribution for innovations, of which the former three are generalisations of the normal distribution. It turned out that only the GPD innovations were useful for the tail-risk strategy, so we focus on this distribution here, but we point out that the other innovations distributions may be useful when considering strategies with other purposes. Therefore, a description of the GH and the α-stable distributions is given in the appendix in Appendix A. Briefly, the GH distribution arises as a special case of a normal distribution with (independent) stochastic variance. The t-distribution is a special GH distribution. α-stable distributions are characterised by random variables whose linear combinations remain in the same family.

3.1. *Generalised Pareto distribution*

The GPD arises as the limit of fitting extreme data exceeding a given threshold (the limit is taken with respect to the threshold). As such, it is the underlying distribution in the method of *threshold exceedances* to describe the tail of the data. Hence, contrary to the other distributions considered, most of the data is neglected, and only the extreme outcomes are included in the estimation.

The distribution function of the GPD is given by

$$G_{\xi,\beta}(x) = \begin{cases} 1 - (1 + \xi x/\beta)^{-1/\xi}, & \xi \neq 0, \\ 1 - \mathbf{e}^{-x/\beta}, & \xi = 0, \end{cases}$$

where $\beta > 0$, $x \geq 0$ when $\xi \geq 0$ and $0 \leq x \leq -\beta/\xi$ when $\xi < 0$. The parameters ξ and β are the *shape* and the *scale* parameters, respectively. The GPD is generalised in the sense that it comprises a number of special cases: If $\xi > 0$, then $G_{\beta,\xi}$ is a Pareto distribution with parameters $1/\xi$ and β/ξ; if $\xi = 0$, then $G_{\beta,\xi}$ is an exponential distribution; and if $\xi < 0$, then $G_{\beta,\xi}$ is a Pareto type II distribution. These three distribution essentially describe the tail behaviour via the so-called *max-domain* or *maximum domain of attraction* (e.g. McNeil *et al.*, 2005; Embrechts *et al.*, 2003): the Pareto distribution is heavy-tailed, the exponential distribution is light-tailed and the Pareto type II distribution is short-tailed.

The key ideas underlying the method of threshold exceedances are as follows: For a random variable X, resp. distribution function F, the *excess distribution* $F_u(x)$ over a threshold u is given by

$$F_u(x) = \mathbf{P}(X - u \leq x | X > u) = \frac{F(x + u) - F(u)}{1 - F(u)},$$

where $0 \leq x \leq x_F - u$ with x_F the right-endpoint of F. The following theorem justifies using the GPD as a proxy for the excess distribution:

Theorem 1 (Pickands-Balkema-de Haan). *There exists a positive function $\beta(u)$ such that*

$$\lim_{u \to x_F} \sup_{0 \le x < x_F - u} |F_u(x) - G_{\xi,\beta(u)}(x)| = 0,$$

if and only if F is in the MDA of H_ξ, $\xi \in \mathbb{R}$, where x_F is the right-endpoint of F.

In practice, there are various methods for determining an appropriate threshold to fit the data. We shall choose the thresholds to be the 95%-quantile of the loss distribution implying that 5% of the data are included in estimations.[1]

3.2. *Model calibration*

In a GARCH model with normally distributed innovations, the observations are conditionally normally distributed. The likelihood function based on n observations $X_1, \ldots, X_n$ is thus given by

$$L(\alpha_0, \alpha_1, \beta; \mathbf{X}) = \prod_{t=1}^{n} \frac{1}{\sigma_t} \mathrm{n}\left(\frac{X_t}{\sigma_t}\right),$$

where $\mathrm{n}(x)$ is the standard normal density and where $\sigma_t = \sqrt{\alpha_0 + \alpha_1 X_{t-1}^2 + \beta \sigma_{t-1}^2}$. Since the value of σ_0 is not observed one typically chooses as starting value the sample variance the observations.

To calibrate the GARCH models with non-normally distributed innovations, one can directly specify the likelihood function involving the respective probability density. However, it is custom to calibrate the GARCH process via *quasi-maximum likelihood estimation* (QMLE; sometimes called pseudo maximum likelihood estimation), e.g. McNeil and Frey (2000), which postulates a normal distribution assumption for the innovations, and to fit the residuals to the GH distribution via MLE in a second step. This two-step approach ensures that the time-varying volatility behaviour is fully captured by the GARCH parameters, while the conditional behaviour, in particular the heaviness of the tails, is captured by the parameters of the innovations' distribution.[2]

3.3. *Value-at-Risk*

Let X_t be a future, unknown asset return. The time-t *value-at-risk (VaR)* at confidence level $\alpha \in (0, 1)$ is given by

$$\mathrm{VaR}_{\alpha,t} = \inf\{l \in \mathbb{R} : \mathbf{P}(-X_t > l) \le 1 - \alpha\} = \inf\{l \in \mathbb{R} : \mathbf{P}(-X_t \le l) \ge \alpha\}.$$

In other words, $\mathrm{VaR}_{\alpha,t}$ is just the α-quantile of the distribution of the loss variable $-X_t$. If X_t admits a strictly positive density, the value-at-risk is characterised by

$$\mathbf{P}(X_t \le -\mathrm{VaR}_{\alpha,t}) = 1 - \alpha.$$

[1]This choice of threshold is somewhat arbitrary, in particular given the fact that the fitted distribution parameters may vary strongly with the choice of the threshold. However, the main element of the trading strategy will be to capture the time-variation of the parameters. Keeping the threshold and consequently the number of samples entering the GPD estimation constant implies that a major factor in the variation of the GPD samples over times lies in the occurrence of extreme events. This reasoning is backed by empirical tests showing that the results of the strategy are robust against quantile choices other than the 95%-quantile, see Section 6.3, as well as an inferior performance when employing data driven dynamic threshold determination such as Desmettre and Deege (2016).

[2]To better capture tail risks, one could, for the case of non-normally distributed innovations, consider GARCH variants that deal with asymmetries of positive and negative shocks such as QGARCH (Sentana, 1995).

Note that in this setting, VaR_α is defined on asset returns rather than asset values, which would be more appropriate for calculating a capital requirement.

To calculate $\text{VaR}_{\alpha,t}$ conditional on the information $\mathcal{F}_{t-1}$, we use the property that the conditional asset return distribution follows the distribution of the innovations. The volatility forecast for one time period from the GARCH model is given by

$$\hat{\sigma}_t^2 = \mathbb{E}(X_t^2|X_{t-1}, \sigma_{t-1}) = \alpha_0 + \alpha_1 X_{t-1}^2 + \beta\sigma_{t-1}^2,$$

The VaR at confidence level α for a one-period time horizon then given by

$$\text{VaR}_{\alpha,t} = -\hat{\sigma}_t F_{1-\alpha}, \tag{1}$$

where $F_{1-\alpha}$ denotes the $1-\alpha$-quantile of the innovations' distribution. This approach extends to portfolios via historical simulation.

4. TAIL-RISK PROTECTION STRATEGY: STEEPNESS OF VAR-SPREAD

From an investor's perspective a *tail event* could be defined as a return X_t that lies below value-at-risk (VaR), i.e. $X_t \leq \text{VaR}_{\alpha,t}$, with $\text{VaR}_{\alpha,t}$ the VaR assuming normally distributed innovations and where the confidence level α may for example be derived from the investor's risk preferences. The goal of a tail-risk protection trading strategy is to be invested in some risky asset or portfolio, but to avoid tail events, for example by hedging the exposure when a tail event is likely to occur. Of course, ex-ante it is impossible to rule out being invested when tail events occur, so a strategy needs to focus on patterns that signal a high likelihood of a tail event occurring. As the approach is based on historical data, one would hope for such a strategy to perform well when there is a build-up of risk in the market. Evidently, such a strategy cannot avoid tail events that occur "totally out of the blue" ("type I error"). Also, given the empirical stylised fact that return data feature little or no autocorrelation implies that such a strategy may signal de-investment when ex-post an investment would have been optimal ("type II error").

Given are time series consisting of financial returns, the corresponding volatility forecasts from GARCH models and the VaR figures from the GARCH-fitted models with different distribution assumptions on the innovations (normal and e.g. Student t, GH, alpha-stable and GPD). In the empirical examples below, there are two data sets: first, one-day VaR's based on the returns of closing prices of the previous 300 trading days are calculated, and second, overnight VaR's from the returns between daily opening and closing prices, again based on a history of 300 trading days, are calculated. The choice of 300 trading days is motivated by the trade-off of a long data set for statistical purposes versus a short data set to pick up extreme events related to the near past. Choosing a longer data horizon could lead to an accumulation of "old" extreme events in the tail and hence a failure to reflect current market conditions.

The *steepness-of-VaR-spread strategy* takes as input *VaR spreads* calculated from GARCH models with different innovations processes (e.g. normal distribution and GPD distribution),

$$s_t^F = \text{VaR}_{\alpha,t}^F - \text{VaR}_{\alpha,t}^{\text{N}},$$

where $\text{VaR}_{\alpha,t}^F$ denotes the value-at-risk from F-distributed innovations. If the VaR spread is small, then the GARCH model with normal innovations is an appropriate model for current market conditions, whereas for a large VaR spread the normal innovations are inappropriate.

Based on a rolling window of a given number of days, a steep VaR spread indicates a build-up of risk over the past trading days, whereas a flat VaR spread implies a constant risk setting.

The slope of the VaR spread therefore aims at detecting risk build-up. More specifically, define a linear model for the VaR spread over p days $t_1, \ldots, t_p$ by

$$s_t^F = \alpha + \beta_{t_p}(t - t_0) + \varepsilon_{t,t_0}^F, \quad t = t_1, \ldots, t_p.$$

Given time series data, $\hat{\beta}_{t_p}$ can be estimated by a least squares approach by regressing the VaR spread on the range $1, 2, \ldots, p$.

The implementation of the strategy depends on two parameters: Aside from p, the number of days determining the steepness of the VaR spread, one needs to define an investment threshold, so that an investor is invested whenever the steepness lies below the threshold and is hedged otherwise. The threshold may for example depend on the historical average steepness or on the historically determined percentage of investment days had the strategy been followed. In our setting, we define the threshold via a multiple q of the average of the positive part of the slope, that is, a risk build-up for the time period $(t-1, t]$ is signalled, if

$$\hat{\beta}_{t-1} \geq q \sum_{i=1}^{t-1} \max(\hat{\beta}_i, 0)/(t-1) =: q\,\overline{\beta}_{t-1}^{+}$$

Restricting attention to positive coefficients reflects that we are looking for indicators of increased risk. When VaR spreads are calculated on several data sets, such as daily returns on closing prices, overnight returns and intraday data, de-investment takes place if any of the signals is positive. In practice, de-investment could be realised through some hedge or overlay.

The two parameters of the strategy, the length of the look-back period p and the threshold multiplier q, are determined as the parameters that optimise the past cumulative alpha of the steepness-of-VaR-spread strategy, where cumulative alpha is defined as the difference between the to-day accumulated return of the tail-risk protection strategy and the to-day accumulated return of the static buy-and-hold strategy. Formally, letting $(r_t)_{t\geq 0}$ denote the daily log-returns of the underlying security, cumulative alpha at time T is defined as

$$\alpha_T(p, q) = \sum_{t=1}^{T} \left(r_t \cdot \mathbf{1}_{\{\hat{\beta}_{t-1} < q\,\overline{\beta}_{t-1}^{+}\}} - r_t \right),$$

and the optimisation problem at each point in time is $\max_{p,q} \alpha_T(q, p)$. Optimising cumulative alpha expresses that positive returns are to be generated from being de-invested when tail events occur, while avoiding de-investment when positive returns occur (preventing “type II errors”).

In order to do out-of-sample testing of the strategy, we split the data set into a “learning set”, initially comprising the older 50% of the time series and a “test set” consisting of the more recent 50% of the time series. The parameters are initially determined from the learning set, and then successively recalibrated as the data set grows through time. In the examples below, the parameters will be calibrated every ten trading days.

5. Empirical Results

5.1. *Data*

For the value-at-risk calculations we use daily DAX future prices (Bloomberg ticker: GX1 Index) from 21 October 1991 until 26 October 2015. Daily VaR’s are calculated based on daily returns and based on overnight returns. Additionally incorporating an intraday VaR based on DAX minute data had no effect on the performance of the trading strategy.

Calculations are performed on a rolling window of 300 trading days, implying that the first VaR is calculated for 5 January 1993. Fixing a time window, the respective DAX future returns

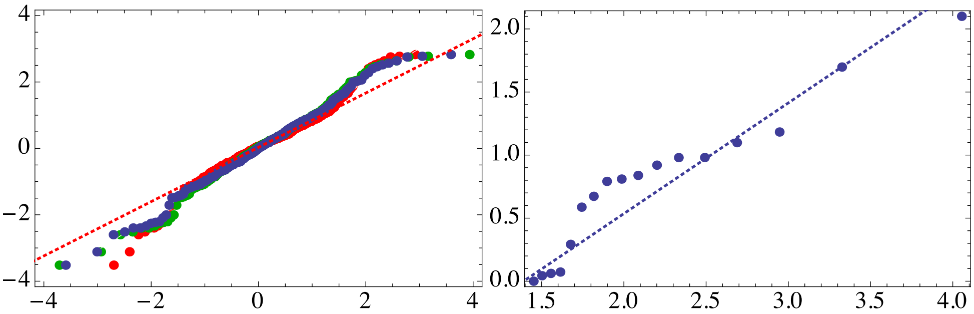

FIGURE 1.— QQ plots of GARCH residuals from 21 January 2014. Left: normal distribution (red), Student-t distribution (green), GH distribution (blue); right: GPD distribution (samples beyond 5%-threshold $u = 1.417$).

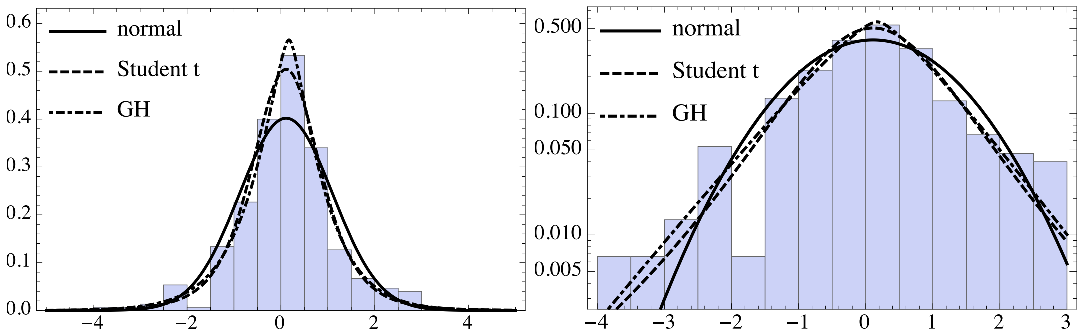

FIGURE 2.— Densities of GARCH residuals from 21 January 2014 together with relative frequencies of innovations. Left: normal scale; right: logarithmic scale.

$(r_t)_{0\le t\le 300}$ are fitted to a GARCH model via QMLE, yielding the volatility process $(\sigma_t)_{0\le t\le 300}$ and residuals $(\varepsilon_t)_{0\le t\le 300}$ with $\varepsilon_t = r_t/\sigma_t$. The residuals are then fitted to the following distributions: (skewed) normal, (skewed) Student-t, GH, α-stable and GPD.[3] Calibration of the GPD is based on the worst 5% of the outcomes. VaR's at $\alpha = 0.99$ are calculated using Equation (1). As an example, residual QQ-plots of one day are given in Figure 1 and the corresponding densities are given in Figure 2. The daily and overnight VaR's and VaR spreads are shown in Figure 9 in the appendix together with the daily volatility forecast.

5.2. *Trading strategies*

Aside from the application of the tail-risk protection strategy, we briefly describe two "traditional" strategies, protective put and delta-replicated put, in Sections 5.2.2 and 5.2.3, respectively. These will be applied in the subsequent sections to benchmark the results of the tail-risk protection strategy.

5.2.1. *Tail-risk protection strategy*

In our results it turns out that only the VaR spread based on GPD innovations produces a satisfactory tail-risk protection strategy. This is plausible, as the GPD focusses on the tail behaviour whereas the other distributions take into account the whole range of the innovations' outcomes. As a consequence, only the GPD spreads from the daily and overnight VaR are

[3]Alpha-stable innovations were calibrated using the GNU R STABLE package, see www.RobustAnalysis.com.

considered. The VaR spreads and a trading signal are determined on a daily basis as outlined in Section 4.

The earlier 50% of the data until 24 June 2004 are used as the initial "learning set" for the parameters p and q, determining the regression length and the threshold factor, respectively. These parameters are re-calibrated every ten trading days using a successively growing learning set. The parameters are chosen to optimise the cumulative alpha from the beginning of the time series to the current calibration date. The regression length p is subject to the constraint that $p \in [8, 15]$, where the lower bound is justified by information criteria such as the AIC, which point to an optimal lag length of 10, resp. 14, for the daily, resp. overnight, GPD VaR spreads in an autoregressive model. Likewise, the autocorrelation functions are significant for lags greater than 8. In our results, the regression length varies around 14 for most of the time, but hits the lower bound of 8 for some time periods. The threshold factor q can vary between 0 and 15, and varies between 1.25 and 2.65 for most of the time, but can take values of around 14 in time periods where the regression length is small. This indicates that a shorter risk build-up time horizon corresponds to a higher threshold criterion. The annual fraction of days that the tail-risk protection strategy is invested varies between 65% and 97%.

5.2.2. *Protective put*

The protective put strategy (PP) uses put options to insure a risky position against adverse market movements. If the put expires out of the money, the investor loses the option premium. However, if the option expires in the money, the gains offset losses from the underlying asset. In practice, the protective put strategy typically involves a sequence of options with different maturities that must be rolled over once their maturity is reached. Figlewski *et al.* (1993) analyse different trading trajectories for the option portfolio through time. Intuitively, the amount and level of protection determine the strategy's price. Counterparty risk aside, PP provides a guaranteed and model-independent protection mechanism (hedge).

The out-of-sample backtest consists of a strategy where an investor buys a three-month European put option on the DAX future with a strike price of 90% of the current futures price. Every three months at expiry of the put option, a new put option is bought. The volatility parameter for pricing the put option is approximated by using the VDAX-NEW implied volatility index. This underestimates the price of the put due to the volatility skew. Purchase prices and any option values are calculated using the Black formula for options on futures.

5.2.3. *Delta-replicated put*

The delta-replicated put strategy (DRP) is similar to PP, but with a dynamic delta replication instead of purchasing physical puts. We calculate the Black-Scholes delta of a hypothetical put option with a strike price of 90% of the current DAX future price each day. The delta of this put option is added to the existing DAX future; according to put-call parity, this results in holding the delta of a 90% call option. The maturity of the hypothetical option is always set to three months. We use the same model and parameters as for the protective put. An advantage of the DRP compared to the protective put is be that it does not suffer from high prices for put options as mentioned in the introduction. However, it is subject to model risk and to gap risk.

The strategy is comparable to the well-known *Constant Proportion Portfolio Insurance (CPPI)* (Black and Jones, 1987; Perold and Sharpe, 1988; Black and Perold, 1992), which aims at creating the asymmetric payout profile of a hypothetical perpetual call option. In contrast to CPPI, which relies on a fixed protection level, the strike of the hypothetical put option in the DRP strategy is always 90% of the current DAX futures price. This is necessary for our type of analysis, as we want to compare strategies that keep similar characteristics in time. A CPPI strategy would change its characteristics once the market increases substantially above the protection

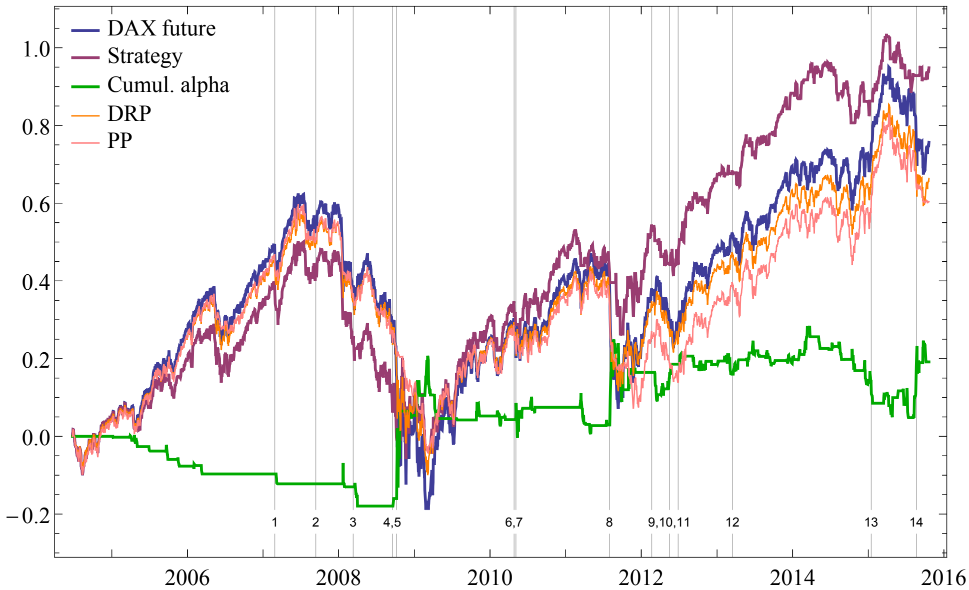

FIGURE 3.— Cumulative log returns of the DAX future (GX1), the tail-risk protection strategy, as well as DRP and protective put strategies; green line: cumulative alpha of the tail-risk protection strategy against the DAX future. Time horizon: out-of-sample test period 25 June 2004 through 21 October 2015. Vertical lines indicate prominent market events: 1 – Chinese stock market correction; 2 – Northern Rock bank run; 3 – Bear Stearns resolution; 4 – Lehman default; 5 – global contagion and interventions; 6 – EU debt crisis; 7 – Flash Crash; 8 – U.S. rating downgrade; 9 – Greek bailout; 10 – Bankia bailout; 11 – Spain,Cyprus support request, 12 – Cyprus bailout; 13 – EURCHF unpegged; 14 – global stock selloff.

level or decreases towards the protection level, essentially resulting in a complete unwind of the risk position.

Compared to PP, the DRP strategy classifies extreme events slightly differently; this is due to the rebalancing at a 90% strike level every day, as opposed to PP where the 90% strike is calculated every three months. Hence by construction (and similar to the tail-risk protection strategy), DRP will not provide tail-risk protection against slow downturns.

5.3. *Out-of-sample backtest*

Because the initial 50% of the data set are used solely for calibration purposes, out-of-sample testing is conducted for the time period 25 June 2004 through 21 October 2015. Here, on each business day a trading signal is generated using the data so far and applied to the subsequent return. The current set-up does not account for market frictions, such as bid-ask spreads, transaction and other associated costs (for institutional investors on liquid future markets these factors should have little impact; a detailed review is provided in Section 6.1).

The outcome of the backtest is shown in Figure 3, together with the evolution of the protective put and delta-replicated put strategies. For a strategy with perfect forecasting capabilities, one would expect a jump in cumulative alpha (green line) every time a tail event occurs. The following prominent extreme events are indicated by vertical lines:

1 February 2007: a large correction in Chinese stock markets spilled over to other major financial markets worldwide.
2 September 2007: Northern Rock Plc. experienced a bank run that forced the bank into public ownership.

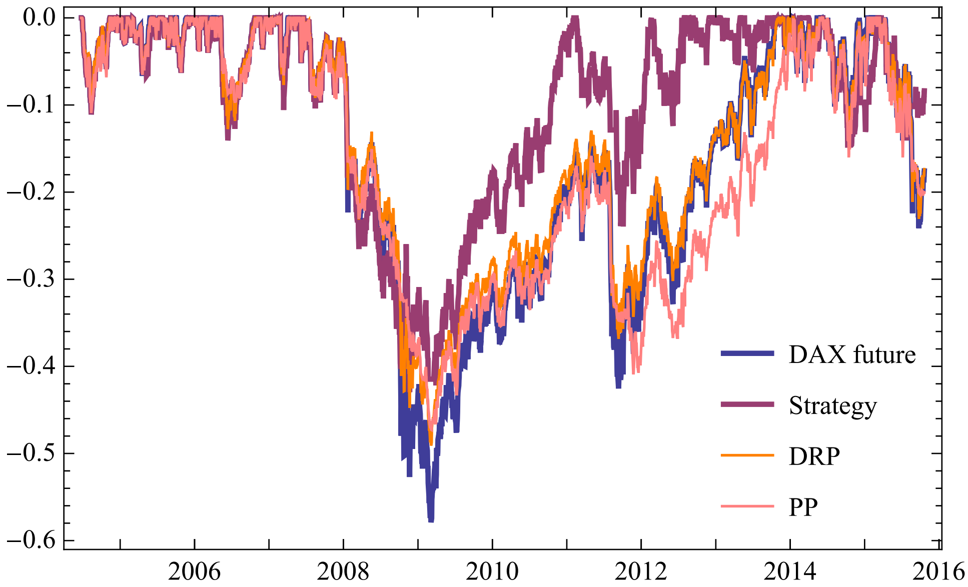

FIGURE 4.— Drawdown of the DAX future, tail-risk protection, strategy, DRP and protective put strategies, in the out-of-sample test time horizon.

3 March 2008: Bear Stearns Inc. received an emergency loan facility by the Federal Reserve Bank of New York. Subsequently, to mitigate a collapse, the bank was sold to JPMorgan.

4 September 2008: Lehman Brothers Holding Inc. filed for bankruptcy. This is often referred to as the peak of the 2008 financial crisis and resulted in heavy market movements.

5 October 2008: Massive global interventions and instantaneous rescue efforts to prevent a chain of defaults. Dominique Strauss-Kahn, the Managing Director of the IMF, said on October 11th: "Intensifying solvency concerns about a number of the largest US-based and European financial institutions have pushed the global financial system to the brink of systemic meltdown."

6 April 2010: the rating agency Standard & Poor's downgraded the sovereign debt of Greece, four days after the EU-IMF bail out was put in place.

7 May 2010: a drop of nearly 1,000 points in DJI with subsequent recovery of 900 points occurred; this is subsumed under the name "Flash Crash". A report by the SEC and CFTC mentions a so-called 'fat-finger-trade' and subsequent order execution of automated trading systems as the cause.

8 August 2011: the equity markets observed large drops in the US, EU and Asia after the rating downgrade of the U.S.

9 February 2012: Greece received its second financial aid package.

10 May 2012: Bankia, the 4th largest Spanish bank, negotiated bailout funding. The rating agency Standard & Poors responds with a downgrade to BB+.

11 June 2012: Spain and Cyprus request financial support from the euro area members and the IMF.

12 March 2013: Cyprus receives bailout funding in return for ageing on several conditions. These include the closure of Cyprus second largest bank and imposing a one-time bank deposit levy.

13 January 2015: The Swiss National Bank unpegged the Franc from the Euro. This had a large effect on the EURCHF exchange rate as well as spill over effects to other financial markets worldwide.

14 August 2015: in the wake of the Chinese stock market crash, many global stock indices

experienced large aftershocks.

Not all extreme events are identified by the tail-risk protection strategy, in particular as events that occur entirely “out of the blue” cannot be captured by a strategy based on historical data. Table II provides a comparison of the average performance of different strategies in a window of 20 trading days around the previously mentioned market events. The columns show relative performance of the tail-risk protection strategy versus the market, DRP and protective put. Positive numbers indicate out-performance, whereas negative numbers indicate miss-identification. In case the strategy failed to produce any signal during the event, the difference to the DAX futures benchmark (Str. − DAX) is zero. The data shows that, for example, the Northern Rock incidence and EURCHF peg removal were not captured by the strategy. This result is expected, as these events occurred “out of the blue”. On average, the strategy out-performs the DAX future as well as DRP and PP during the sample market events. The relatively small out-performance over PP shows that the insurance function of protective put works during downturn scenarios.

TABLE II

Mean performance in a 20 day window around prominent market events.

Event	Str. - DAX	Str. - DRP	Str. - PP
1	-0.001212	-0.001321	-0.001501
2	0.000000	0.000358	0.000502
3	-0.001369	-0.001739	-0.001225
4	0.000912	0.000113	-0.001066
5	0.005787	0.002692	-0.002980
6	-0.000555	-0.000624	-0.000509
7	0.000987	0.000633	0.000043
8	0.005996	0.004441	0.003797
9	0.000000	-0.000071	0.000073
10	0.003034	0.002457	0.002546
11	0.000816	0.001260	0.002180
12	-0.000940	-0.000863	-0.000899
13	-0.002786	-0.002296	-0.003375
14	0.006727	0.006389	0.004405
Mean	0.001243	0.000816	0.000142

Figure 4 shows the drawdown of the DAX future and the three strategies under consideration. Here, drawdown is the return relative to the peak to-date (high watermark) achieved by each strategy. Performance measures for the out-of-sample test period are reported in Table III. For ease of interpretation the daily returns are expressed per annum. The Sharpe ratio is calculated as the ratio of average return and standard deviation (ignoring the risk-free rate). The Calmar ratio is the average mean (p.a.) relative to the worst drawdown.

Both, the drawdown in Figure 4 and the performance measures in Table III establish that overall the tail-risk protection strategy outperforms the other strategies: for most of the time period the strategy's drawdown is superior to the DAX future, the overall worst drawdown is smaller than for all other strategies and the Sharpe and Calmar ratios are greater than the respective ratios for the other strategies.

In addition, we analyse performance for time horizons shorter than the entire backtest period of 11 years. Figure 5 presents box-and-whisker plots associated with the performance of investments in the DAX future, the tail-risk protection strategy, DRP and protective put over a time period of 252 trading days. For every possible starting point in the test set, the performance of a subsequent 252 trading day investment are calculated. The tail-risk protection's return distribution mostly affects the boundaries. The distributions of standard deviation display in particular the effectiveness of protective put, but also the reduction in risk from DRP

TABLE III

DAX FUTURE VERSUS PERFORMANCE OF TAIL-RISK PROTECTION, DRP AND PROTECTIVE PUT STRATEGIES

	DAX future	tail-risk protection	DRP	Protective Put
mean return (p.a.)	0.0656	0.0821	0.0575	0.0524
standard deviation	0.2169	0.1791	0.1739	0.1550
Sharpe ratio	0.3022	0.4587	0.3306	0.3378
worst drawdown	0.5791	0.4296	0.4910	0.4730
Calmar ratio	0.1132	0.1912	0.1171	0.1107

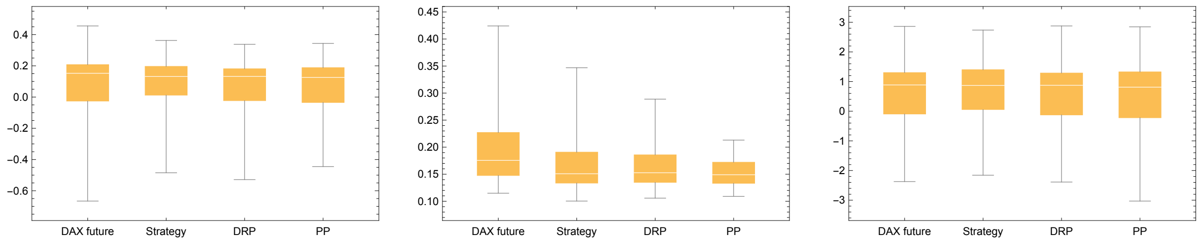

FIGURE 5.— Box-and-whisker plots of performance on a 252-day investment horizon; each box and each whisker represents 25% of the data; all values are annualised. Left: return; middle: standard deviation; right: Sharpe ratio.

and the tail-strategy. The risk adjusted out-performance of the tail-strategy is again verified in the Sharpe ratio analysis.

5.4. *Validation*

The outcome of the out-of-sample backtest of the previous section is very promising. In this section, we provide further evidence that this may not be due to sheer luck in the time period under consideration. To this end, we test the hypotheses that the tail-risk protection strategy has no predictive power over various benchmarks. Some of these tests require application of bootstrapping (Efron, 1979), which is widespread in the evaluation of trading strategies, see e.g. Brock *et al.* (1992); West (1996); White (2000); Sullivan *et al.* (1999). To acknowledge dependence in the underlying time series, we use the stationary bootstrap technique developed by Politis and Romano (1994). In the following, we largely adopt the notation of White (2000).

The observed returns of the benchmark are generated by an adapted process $(X_t)_{t\in\mathbb{Z}}$. It is assumed that $(X_t)_{t\in\mathbb{Z}}$ is a stationary strong α-mixing sequence with marginal distributions identical to a random variable X.[4] Trading signals are generated for times $R, \ldots, T$ where $T = R + n - 1$. The signals are estimators $\hat{p}_{R+1}, \ldots, \hat{p}_{T+1}$, with $\hat{p}_{R+k+1} \in \{0, 1\}$ based on $(X_1, \ldots, X_{R+k})$.

We wish to test the hypothesis that the strategy has no predictive power over a benchmark asset or benchmark strategy. Formally, the hypothesis is expressed as

$$\mathbb{E}[f^\star] \leq 0, \qquad \text{with } f^\star = Xp^\star - Y,$$

where $p^\star = \operatorname{plim} \hat{p}_T$ and where Y is the benchmark strategy. In the simplest case, $Y = X$. The

[4]A strong α-mixing sequence is a stochastic process consisting of dependent random variables that behave more like independent random variables the farther they are separated. Formally, for a stochastic process $(X_t)_{t\in\mathbb{Z}}$, let $\mathcal{B}_t$ be the Borel field generated by $(X_s)_{s\leq t}$ and $\mathcal{F}_t$ be the Borel field generated by $(X_s)_{s\geq t}$. Define $\alpha(s) := \sup\{|\mathbf{P}(A \cap B) - \mathbf{P}(A)\mathbf{P}(B)| : A \in \mathcal{B}_t, B \in \mathcal{F}_{t+s}, t \in \mathbb{Z}\}$. The process $(X_t)_{t\in\mathbb{Z}}$ is *strong α-mixing* if $\alpha(s) \to 0$ as $s \to \infty$ (Rosenblatt, 1956, 1971).

TABLE IV

p-VALUES CORRESPONDING TO THE HYPOTHESIS THAT THE TAIL-RISK PROTECTION STRATEGY HAS NO PREDICTIVE POWER OVER THE BENCHMARK. THE BOTTOM THREE STRATEGIES ARE SIMULATED USING STATIONARY BOOTSTRAPPING WITH PARAMETER 0.1. THE SAMPLE LENGTH IS $n = 2883$, AND THE NUMBER OF SIMULATIONS IS 500.

Benchmark	benchmark strategy Y_{t+1}	p-value
Permutation	$p_{R+\pi_t} X_{t+1}$	0.14
DAX future	X_{t+1}	0.26
DRP	X_{t+1}^{DRP}	0.15
Protective Put	X_{t+1}^{PP}	0.14

test statistic is

$$\overline{f} = \frac{1}{n} \sum_{t=R}^{T} \hat{f}_{t+1},$$

with $\hat{f}_{t+1} = X_{t+1} p_{t+1} - Y_{t+1}$.

Any tests are based on the assumption that the following central limit theorem holds (conditions are given in West (1996), see also White (2000)):

$$\sqrt{n}(\overline{f} - \mathbb{E}(f^{\star})) \xrightarrow{\mathcal{L}} \mathrm{N}(0, \sigma^2),$$

expressing that the scaled and normalised estimator converges in distribution to a normal distribution.

First, we test the hypotheses that the strategy's return is not superior to various benchmark returns. One benchmark is a random permutation of the same number of long and neutral trading positions as the strategy itself. In other words, letting $\pi = \{1, \ldots, n\}$ be a random permutation drawn from the distribution that assigns equal probability to all permutations, set

$$Y_{t+1} = p_{R+\pi_t} X_{t+1}.$$

The permutation test gives an indication of whether a simple random strategy could have achieved a similar average return.

Other benchmarks are the DAX future return, the return of the DRP strategy and the return of the protective put strategy. Here, we use the stationary bootstrap by Politis and Romano (1994) to account for dependence in the time series. In a conventional bootstrap sample of length n, the samples are drawn independently from the return distribution. In the stationary bootstrap, the sample is made up of sequences of returns of random length, determined by a geometric distribution. In our case, the parameter chosen for the geometric distribution is 0.1 implying an expected sequence length of 10, which is consistent with the autocorrelation present in the data, cf. Section 5.2. The resulting p-values are given in Table IV.

The results can be interpreted as follows: the null hypothesis cannot be rejected for the DAX futures benchmark. For the random permutation, PP and DRP, the result is inconclusive, but gives some evidence that the average return achieved by the tail-risk protection strategy is not purely random. The result regarding the DAX futures benchmark is inline with the fact that the strategy's objective is not formulated as "beating the benchmark's return".

Rather, the strategy's objective is to reduce risk by avoiding tail events. Therefore, we test the hypothesis that the strategy's Sharpe ratio, resp. worst drawdown, does not outperform the DAX future's Sharpe ratio, resp. worst drawdown. For comparison purposes, we also give the p-values when performing the same tests, but replacing the tail-risk strategy with the DRP, resp. the protective put strategies. The results are given in Table V.

TABLE V

p-VALUES CORRESPONDING TO THE HYPOTHESIS THAT THE VARIOUS STRATEGIES DO NOT OUTPERFORM THE DAX FUTURE IN TERMS OF SHARPE RATIO AND IN TERMS OF WORST DRAWDOWN. ALL RESULTS ARE OBTAINED USING STATIONARY BOOTSTRAPPING WITH PARAMETER 0.1. THE SAMPLE LENGTH IS $n = 2883$, AND THE NUMBER OF SIMULATIONS IS 500.

	p-values	
Strategy	Sharpe ratio	worst drawdown
Tail-risk protection	0.16	0.03
DRP	0.19	0.00
Protective Put	0.32	0.04

Here we find clear evidence that the tail-risk protection strategy reduces the worst drawdown, whereas the results about the Sharpe ratio are more inconclusive. The results only change marginally when including transaction costs.

Summing up, the tests indicate that the tail-risk protection strategy outperforms the DAX future in terms of worst drawdown. There is also some slightly weaker evidence that the return performance is not purely random and that the strategy's Sharpe ratio is better than the DAX futures' Sharpe ratio.

5.5. *Investor's utility*

A systematic approach for capital investment calls for taking into consideration investors' risk preferences. First, what kind of utility function is required to prefer the tail-risk protection strategy over alternative strategies? Second, how can the tail-risk protection strategy be extended to capture individual risk and return preferences?

Before going into the details regarding the first question, we briefly consider the second question. In the spirit of classical portfolio theory, the tail-risk protection strategy maximises returns while restricting the set of investment strategies according to some risk criterion, in this case by seeking to avoid tail events. To incorporate individual risk preferences, the restrictions would need to be formulated in an investor-specific way. For example, risk preferences could differ with respect to the degree of tolerance of accepting loss events, taking into account that a stronger dislike of losses trades off with a higher number of false de-investments. A full treatment is beyond the scope of the paper, in particular taking into account the path dependency of the strategy, cf. (Cox and Leland, 2000).

Turning to the first question raised above, one could ask to what extent do investors prefer this strategy over the given traditional alternatives. Preferences are captured by utility functions that are twice continuously differentiable monotone and concave, that is, utility functions u with $u \in C^2$ with $u' > 0$ and $u'' < 0$. A strong criterion is *uniform preference*, capturing that a payoff or strategy is preferred over another by *all* investors. This is identical to the notion of second order stochastic dominance (SSD). More specifically, SSD is a partial order on the set of integrable random variables; for two random variables X and Y with distribution functions F and G, respectively, X second order stochastically dominates Y if and only if (e.g. Föllmer and Schied, 2002),

$$\int_{-\infty}^{k} F(x)dx \leq \int_{-\infty}^{k} G(x)dx \quad \text{for all } k.$$

Kaur *et al.* (1994) designed a test to empirically establish if an SSD-relationship holds between two distributions. We apply this test to the outcomes of investments into the DAX future, the tail-risk protection strategy, DRP and PP. The investment horizon under consideration is fixed to 252 trading days.[5] The following hypotheses are tested:

$$H_0 : \int_{-\infty}^{k} F(x)dx \leq \int_{-\infty}^{k} G(x)dx \quad \text{for some } k \in [a,b],$$

versus

$$H_1 : \int_{-\infty}^{k} F(x)dx > \int_{-\infty}^{k} G(x)dx \quad \text{for all } k \in [a,b].$$

We set a to be the smallest common value of X and Y and b the largest common value, respectively.

TABLE VI

SECOND ORDER STOCHASTIC DOMINANCE TEST OF THE TAIL-RISK PROTECTION STRATEGY AGAINST ALTERNATIVE STRATEGIES AND THE DAX FUTURE. SHOWN IS THE PERCENTAGE OF THE INTERVAL $[a,b]$ WHERE THE TAIL-RISK PROTECTION STRATEGY SIGNIFICANTLY DOMINATES THE RESPECTIVE ALTERNATIVE STRATEGY.

	10%*	5%**	1%***
PP	82.93%	81.54%	76.46%
DRP	100.00%	100.00%	96.99%
DAX	100.00%	100.00%	100.00%

Table VI shows the percentage of $k \in [a,b]$ where H_0 is rejected in favour of H_1 with 10%, 5% and 1% significance. According to this test, the DAX future and DRP are clearly dominated at all significance levels except the 1% level for DRP. This strongly suggests that any risk-averse investor would prefer the tail-risk protection strategy over a static buy-and-hold investment or DRP. In fact, a static buy-and-hold investment in the DAX future is still dominated at the 5% significance level when deducting an annual fee of 40 basis points.

The result is weaker when considering the tail-risk protection strategy against protective put. For the upper 82.93% of the interval $[a,b]$, PP is second order stochastically dominated by an investment in the strategy. That the dominating part spans the right side of the return distribution is not surprising as a guaranteed protection from negative tail events can only be achieved via PP.

6. ROBUSTNESS

In this section, we provide further robustness checks by incorporating trading costs, testing the outcome of the strategy when considering different thresholds for the GPD innovations and considering other assets.

6.1. *Trading costs*

To analyse the trading costs of the tail-risk protection strategy we consider realised bid/ask-prices for the DAX future (GX1). In addition, we analyse trading volume and trading range at the specific trade dates to infer a more refined view of the market situation, in particular given that there may be difficulties placing trades at specific points in time. This is a particular concern because the strategy aims at identifying negative tail events that, by definition, occur commonly in times of market distress. An extreme market environment could even inhibit a change of position, even if a valid signal is generated.

[5] Robustness checks with varying investment horizons did not produce notably different results.

TABLE VII

TRADE DATE ANALYSIS OF THE TAIL-RISK PROTECTION STRATEGY. B/A-SPREADS ARE RELATIVE TO THE ASK PRICE, THAT IS, $(A_t - B_t)/A_t$. TRADE RANGE DESCRIBES THE RELATIVE DIFFERENCE OF THE INTRA-DAY HIGH AND LOW PRICE $(H_t - L_t)/L_t$. STANDARD DEVIATION IN PARENTHESES.

Trading days	Strategy	All
Time horizon (years)	11.29	
Number of trades	160	
Percentage of days invested	0.79	
B/A-spread paid	0.019939	
per year average	0.001766	
Average B/A-spread	0.000249	0.000238
	(0.000218)	(0.000215)
Average trade range	0.020598	0.018316
	(0.014077)	(0.013565)
Average daily volume	140617.48	134201.63
	(51873.22)	(51491.10)

Table VII shows the analysis on the trade dates of the tail-risk protection strategy. Within the 11.29 years test set, the strategy produced an average of 14 trades per year, resulting in a long delta one exposure in 78.93% of all trading days. On every trade, half the bid ask spread is paid such that opening and closing a position costs the full bid/ask-spread. This creates average annual trading costs of 17.66 basis points, which is well below the premium of 40 basis points calculated in Section 5.5 that risk averse investors would be willing to pay.

The lower half of Table VII provides information on the market environment on the average strategy-trading day compared to the full test set period. More specifically, we consider the relative spread between the bid price B_t and the ask price A_t. Similarly, the percentage range between the highest price H_t and the lowest price L_t is analysed. The average trade volume is used as an indicator for market depth. On average, all three measures show increased figures for the tail risk strategy. However, all numbers are within one standard deviation of the overall average. Hence, we do not find evidence that the strategy structurally demands trading during unusual or illiquid market times.

6.2. *Other markets*

To further validate the above results, the tail-risk protection strategy is applied to further assets. This includes the S&P 500 stock index future (Bloomberg: SP1 Index) as well as the West Texas Intermediate crude oil future (Bloomberg: CL1 Comdty).

It appears that the strategy performs well for the S&P 500 equity index future, see Figure 6. Several major market events are captured (e.g. the Lehman default in 2008), whereas "out of the blue events" remain undetected. WTI, on the other hand, is not only an entirely different asset class, but also features a market with decline of most of the test period. This is due to the oil storage costs that are reflected in a futures term structure in contango: the storage costs lead to a negative drift for a buy-and-hold investor in oil futures (Erb and Campbell (2016)). The continuous futures time series is generated in Bloomberg using the ratio rollover method (Went (2010)) and thus significantly deviates from spot oil prices. In addition to the long-term negative drift, the year 2008 was characterised by a drastic shift in inflation expectations that led to an oil rally until July 2008 and a long liquidation phase until February 2009. It can be observed that the strategy is not designed to outperform in extended but stable periods of market decline. In other words, tail events are defined relative to the current market environment; if this is in a steady decline, then these negative events will not be classified as tail events.

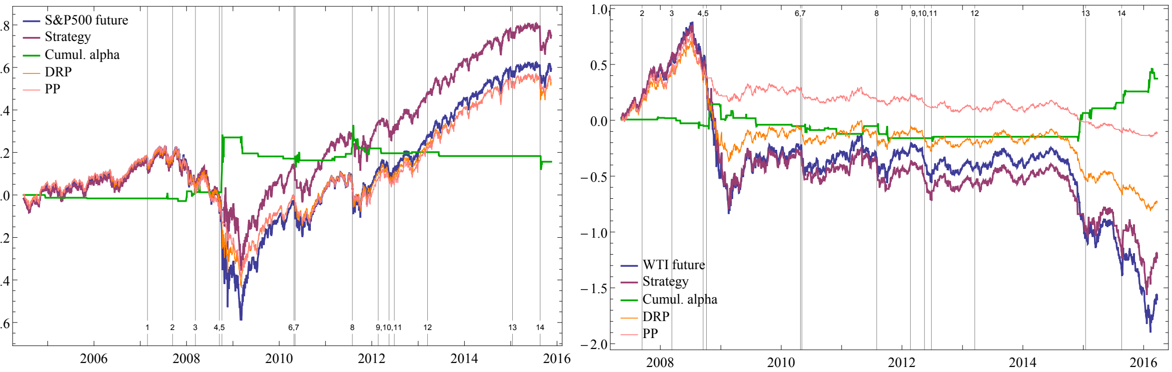

FIGURE 6.— Cumulative log returns of the S&P 500 future (left) and WTI future (right) against the tail-risk protection strategy, as well as DRP and protective put strategies, respectively. Time horizon: out-of-sample test period 25 June 2004 through 29 March 2016. Vertical lines indicate prominent market events, see Section 5.3 for a detailed description.

6.3. *GPD thresholds*

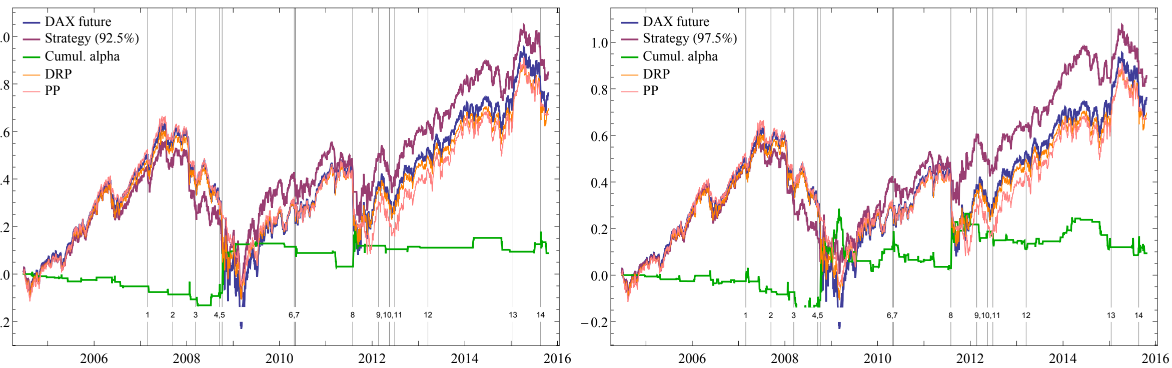

FIGURE 7.— Cumulative log returns of the DAX future against the tail-risk protection strategy with a GPD threshold of 92.5% (left) and 97.5% (right), as well as DRP and protective put strategies, respectively. Time horizon: out-of-sample test period 25 June 2004 through 21 October 2015. Vertical lines indicate prominent market events, see Section 5.3 for a detailed description.

One possible caveat for the performance of the tail-risk protection strategy is the somewhat arbitrary threshold choice for fitting the GPD distribution. In the following, we consider the tail-risk protection strategy from Section 5 while, ceteris paribus, fitting the GPD with thresholds of 92.5% and 97.5%. Figure 7 indicates that the strategy is robust to such model changes. As outlined earlier, because the strategy is sensitive to risk build-up, the arrival of extreme events will affect the tail entering the GPD estimation. It is hence plausible that a constant threshold is more effective than a dynamic threshold. A lower threshold, that is, including more samples in the tail, will decrease the weight associated with each sample, and therefore potentially weaken the risk build-up detection.

6.4. *Alternative implementations of the tail-risk protection strategy*

The tail-risk protection strategy delivers "long" and "flat" signals with respect to a risk exposure in a futures market. The actual implementation of a trading strategy could be done in several ways. The purest form would be to directly implement the signals using long and flat positions in the corresponding future. In financial industry product language, this could be labelled a “rule-based smart beta strategy”. A client could consider such a product as an alternative to a long-only product to reduce tail risk.

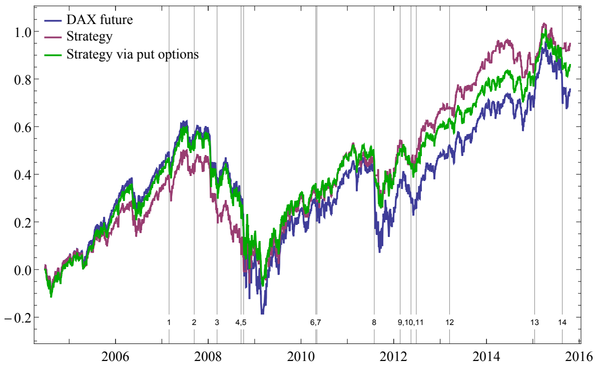

FIGURE 8.— Cumulative log returns of the DAX future against the tail-risk protection strategy that was implemented via futures and put options. Time horizon: out-of-sample test period 25 June 2004 through 21 October 2015. Vertical lines indicate prominent market events, see Section 5.3 for a detailed description.

A different implementation would be to protect an already existing long equity position consisting e.g. of ETFs or an index-like basket of single stocks with an overlay of futures. Whenever the tail-risk protection strategy signals to go “flat”, a short position in the future would be implemented to protect the given equity position. Such an overlay could be part of a risk management strategy of an institutional investor like a pension fund or a hedge fund. In such an environment, also other financial instruments beyond short positions in futures could be considered to implement the tail risk signals. For example, the signals to unwind the market position could be implemented via long put options. This requires an additional decision in terms of the maturity and strike price of the puts to be selected. In practice, a listed European put option offered by an exchange like the EUREX or the purchase of a put warrant from an issuing bank would be viable instruments. Figure 8 shows a performance comparison of the tail-risk protection strategy using futures and three-month puts on futures at a strike price of 90% of the underlying price at the purchase date. The comparison shows a performance difference for the “hybrid” implementation using the puts relative to the “pure” implementation using the futures, but no systematic advantage or disadvantage across the whole time window. As the devised tail-risk strategy does not provide a “target time horizon” for each tail risk signal, the choice of the put maturity date cannot be directly derived from the signals. In addition to that, a European put before maturity shows a weaker delta than −100%. Thus, the resulting delta position immediately after the purchase of the put is lower than 100%, but not zero like in the base implementation using the futures directly. This explains why the implementation using the put follows more the underlying market than the implementation using the futures.

Finally, banks that play an active role in market making and risk warehousing run substantial volatility risk (“vega”) positions (Taleb (1997)). In many cases of “hedged” structured products desks, the resulting portfolios show negative convexity risk in volatility, Jeffery (2004). Due to the illiquidity of the structured product positions and the corresponding vega hedges, a short-term reaction to sudden market movements on the option market is often not possible. Therefore, tail risk protection using futures could also be relevant for the trading desks of those banks.

7. CONCLUSION

We develop a dynamic trading strategy that aims at protecting against large market downturns by taking into account the time-variation and dynamics of distributional parameters of financial time series. First, accounting for the time-dependent dynamics of distributional parameters via a GARCH process allows to incorporate volatility clustering and autoregressive behaviour in volatility. Second, by fitting the GARCH residuals (innovations) to flexible distribution families incorporating both normal and extreme behaviour allows to determine whether, in a given time period, extreme events are more likely to occur than suggested by e.g. normal innovations.

The key input is the spread between a value-at-risk (VaR) from a GARCH model with GPD innovations and a VaR from a GARCH model with normally distributed innovations. The GARCH process filters the time-varying volatility behaviour. Because of the GARCH component, the magnitude of VaR, when viewed as a process over time, adapts quickly to changes in volatility. The distributional properties of the innovation process on the other hand provide, amongst other things, information on the tail risk in the data. The resulting VaR spread can therefore be used to derive an expectation on the frequency of extreme events, and as such, may be taken to produce signals about the presence of tail risks.

We show that the general idea of applying tail risk measures and EVT to identify extreme events yields promising results in tail-risk protection and risk management. An extensive out-of-sample backtest of the strategy indicates that the tail-risk protection strategy is able to protect against those extreme tail-risks that are preceded by a risk build-up. Those risks that occur "out of the blue" cannot be captured by a strategy that is built on historical data. We find that risk averse investors would be willing to pay a positive premium in the range of 40 basis points in order to move from a static buy-and-hold investment in the DAX future to the tail-risk protection strategy. Further statistical tests involving permutation tests and bootstrapping provide evidence that the out-of-sample performance of the tail-risk protection strategy is not due to sheer luck. Application to further assets (S&P 500 future and WTI Oil future) illustrate the strategy's different behaviour when comparing true tail-risk events and longer-term liquidation periods.

The model leaves room for several extensions: First, to achieve a high level of generality, the strategy proposed here focuses only on the most prominent empirical stylised facts that are observed in most financial time series. One could enhance the strategy by incorporating more specific knowledge about the asset class at hand. Second, apart from the current cumulative alpha optimisation, other calibration strategies could be considered. One example is the optimisation based on quantiles or conditional tail measures. Alternatively, one could also consider an economic loss function (e.g. in the spirit of Fleming *et al.* (2001)) and explicitly take into account investor risk preferences.

REFERENCES

Baker, M., B. Bradley, and J. Wurgler. Benchmarks as limits to arbitrage: Understanding the low-volatility anomaly. *Financial Analysts Journal*, 67(1):40–54, 2011.

Barndorff-Nielsen, O. Exponentially decreasing distributions for the logarithm of particle size. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 353(1674):401–419, 1977.

Bhansali, V. and J. Davis. Offensive risk management: Can tail risk hedging be profitable? *Available at SSRN 1573760*, 2010.

Bhansali, V. and J. M. Davis. Offensive risk management ii: The case for active tail hedging. *The Journal of Portfolio Management*, 37(1):78–91, 2010.

Bhansali, V. *Tail risk hedging: Creating Robust Portfolios for Volatile Markets.* McGraw-Hill, 2014.

BIS. Fundamental review of the trading book - interim impact analysis. Basel Committee on Banking Supervision, Bank for International Settlements, 2015.

BIS. Minimum capital requirements for market risk. Basel Committee on Banking Supervision, Bank for International Settlements, 2016.

Black, F. and R. W. Jones. Simplifying portfolio insurance. *The Journal of Portfolio Management*, 14(1):48–51, 1987.

Black, F. and A. Perold. Theory of constant proportion portfolio insurance. *Journal of Economic Dynamics and Control*, 16(3):403–426, 1992.

Bollerslev, T. and V. Todorov. Tails, fears, and risk premia. *The Journal of Finance*, 66(6):2165–2211, 2011.

Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

Bollerslev, T. A conditionally heteroskedastic time series model for speculative prices and rates of return. *The Review of Economics and Statistics*, pages 542–547, 1987.

Brock, W., J. Lakonishok, and B. LeBaron. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance*, pages 1731–1764, 1992.

Cont, R. and P. Tankov. *Financial Modelling with Jump Processes*. Chapman & Hall/CRC, 2004.

Cont, R. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.

Cox, J. C. and H. E. Leland. On dynamic investment strategies. *Journal of Economic Dynamics and Control*, 24(11):1859–1880, 2000.

Desmettre, S. and M. Deege. Modeling redemption risks of mutual funds using Extreme Value Theory. *Journal of Risk*, August 2016.

Eberlein, E. and E. A. v. Hammerstein. Generalized hyperbolic and inverse Gaussian distributions: limiting cases and approximation of processes. In *Seminar on Stochastic Analysis, Random Fields and Applications IV*, pages 221–264. Springer, 2004.

Eberlein, E. and U. Keller. Hyperbolic distributions in finance. *Bernoulli*, 1(3):281–299, 1995.

Eberlein, E. and K. Prause. The generalized hyperbolic model: financial derivatives and risk measures. In *Mathematical Finance-Bachelier Congress*, pages 245–267, 2000.

Eberlein, E. Generalized hyperbolic models. In Cont, R., editor, *Encyclopedia of Quantitative Finance*, pages 833–836. John Wiley & Sons Ltd., 2010.

Efron, B. Bootstrap methods: another look at the jackknife. *Annals of Statistics*, pages 1–26, 1979.

Embrechts, P., C. Klüppelberg, and T. Mikosch. *Modelling Extremal Events for Insurance and Finance*. Springer, 2003. 4th printing.

Engle, R. F. and A. J. Patton. What good is a volatility model. *Quantitative Finance*, 1(2):237–245, 2001.

Eraker, B. and Y. Wu. Explaining the negative returns to VIX futures and ETNs: An equilibrium approach. *Available at SSRN 2340070*, 2014.

Erb, C. and R. Campbell. Conquering misperceptions about commodity futures investing. *Financial Analysts Journal*, 72(4):26–35, 2016.

Feller, W. *An Introduction to Probability Theory and Its Applications*, volume 2. John Wiley & Sons, New York, 2nd edition, 1971.

Figlewski, S., N. Chidambaran, and S. Kaplan. Evaluating the performance of the protective put strategy. *Financial Analysts Journal*, 49(4):46–56, 1993.

Fleming, J., C. Kirby, and B. Ostdiek. The economic value of volatility timing. *The Journal of Finance*, 56(1):329–352, 2001.

Föllmer, H. and A. Schied. *Stochastic Finance. An Introduction in Discrete Time*. de Gruyter, 2002.

Frazzini, A. and L. H. Pedersen. Betting against beta. *Journal of Financial Economics*, 111(1):1–25, 2014.

Godin, F. Minimizing CVaR in global dynamic hedging with transaction costs. *Quantitative Finance*, forthcoming, 2015.

Jeffery, C. Reverse cliquets: end of the road? *RISK*, February 2004.

Kaur, A., B. P. Rao, and H. Singh. Testing for second-order stochastic dominance of two distributions. *Econometric theory*, 10(05):849–866, 1994.

Kelly, B. and H. Jiang. Tail risk and asset prices. *Review of Financial Studies*, 27(10):2841–2871, 2014.

Kozhan, R., A. Neuberger, and P. Schneider. The skew risk premium in the equity index market. *Review of Financial Studies*, 26(9):2174–2203, 2013.

Lemperiere, Y., C. Deremble, T.-T. Nguyen, P. A. Seager, M. Potters, and J.-P. Bouchaud. Risk premia: Asymmetric tail risks and excess returns. Working paper, 2014.

Loh, L. and S. Stoyanov. Tail risk of equity market indices: An extreme value theory approach. Technical report, EDHEC Risk Institute, February 2014.

Longin, F. and B. Solnik. Extreme correlation of international equity markets. *The Journal of Finance*, 56(2):649–676, 2001.

Luo, X. and J. E. Zhang. The term structure of vix. *Journal of Futures Markets*, 32(12):1092–1123, 2012.

Madan, D. B. Adaptive hedging. Working paper, Robert H. Smith School of Business, University of Maryland, June 2016.

McNeil, A. and R. Frey. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7(3):271–300, 2000.

McNeil, A., R. Frey, and P. Embrechts. *Quantitative Risk Management*. Princeton University Press, Princeton, NJ, 2005.

Mehra, R. and E. C. Prescott. The equity premium: A puzzle. *Journal of Monetary Economics*, 15(2):145–161, 1985.

Perold, A. and W. Sharpe. Dynamic strategies for asset allocation. *Financial Analysts Journal*, 44(1):16–27, 1988.

Politis, D. N. and J. P. Romano. The stationary bootstrap. *Journal of the American Statistical Association*, pages 1303–1313, 1994.

Rachev, S. and S. Mittnik. *Stable Paretian Models in Finance*. John Wiley & Sons, 2000.

Rachev, S., B. Racheva-Iotova, and S. Stoyanov. Capturing fat tails. *Risk Magazine*, pages 72–77, May 2010.

Rosenblatt, M. A central limit theorem and a strong mixing condition. *Proceedings of the National Academy of Sciences of the United States of America*, 42(1):43–47, 1956.

Rosenblatt, M. *Markov processes, Structure and Asymptotic Behavior*. Springer, 1971.

Sentana, E. Quadratic arch models. *The Review of Economic Studies*, 62(4):639–661, 1995.

Strub, I. S. Tail hedging strategies. Available at SSRN 2261831, 2013.

Sullivan, R., A. Timmermann, and H. White. Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance*, pages 1647–1691, 1999.

Taleb, N., editor. *Dynamic Hedging: Managing Vanilla and Exotic Options*. Wiley, 1997.

Thurner, S., J. D. Farmer, and J. Geanakoplos. Leverage causes fat tails and clustered volatility. *Quantitative Finance*, 12(5):695–707, 2012.

Went, P. Creating continuous futures price series. Working paper, 2010.

West, K. D. Asymptotic inference about predictive ability. *Econometrica*, pages 1067–1084, 1996.

White, H. A reality check for data snooping. *Econometrica*, pages 1097–1126, 2000.

Zhang, J. E., J. Shu, and M. Brenner. The new market for volatility trading. *Journal of Futures Markets*, 30(9):809–833, 2010.

APPENDIX A: FURTHER GENERALISED INNOVATIONS

A.1. *Generalised hyperbolic distribution*

The GH distribution arises as a special case of the normal mean-variance mixture (NMVM) distribution. Random variables following an NMVM distribution can be represented as

$$X \overset{\mathcal{L}}{=} m(W) + \sqrt{W}\sigma Z, \tag{2}$$

where $Z \sim \mathrm{N}(0,1)$, that is, Z is standard normally distributed, $W \geq 0$ is a random variable independent of Z, the so-called *mixing variable*, $\sigma \in \mathbb{R}_+$ and $m : [0,\infty) \to \mathbb{R}$ is a measurable function, cf. e.g. Section 3.2.2 of (McNeil *et al.*, 2005). The GH distribution arises as the special case where W follows a generalised inverse Gaussian (GIG) distribution, and the Student t distribution arises when W is inverse gamma (IG) distributed. Representation (2) illustrates that both the GH and the t distribution are generalisations of the normal distribution as they essentially consist of a normal random variable with a stochastic standard deviation.

The GH distribution was introduced by (Barndorff-Nielsen, 1977) and has been shown to achieve an almost perfect statistical fit to stock market return data, e.g. Eberlein and Prause (2000); Eberlein and Keller (1995).

The GH density is given by, see e.g. (Eberlein, 2010),

$$f(\lambda,\alpha,\beta,\delta,\mu)(x) = a \cdot (\delta^2 + (x-\mu)^2)^{(\lambda-\frac{1}{2})/2} \cdot K_{\lambda-\frac{1}{2}}(\alpha\sqrt{\delta^2+(x-\mu)^2})\exp(\beta(x-\mu)),$$

with

$$a := a(\lambda,\alpha,\beta,\delta,\mu) = \frac{(\alpha^2-\beta^2)^{\lambda/2}}{\sqrt{2\pi}\alpha^{\lambda-\frac{1}{2}}\delta^\lambda K_\lambda(\delta\sqrt{\alpha^2-\beta^2})},$$

and where K_λ denotes the modified Bessel function of the third kind with index λ. The parameters are interpreted as follows: $\alpha > 0$ is the shape parameter, β, with $0 \leq |\beta| < \alpha$ is the skewness ($\beta = 0$ is the symmetric case), $\mu \in \mathbb{R}$ is the location parameter, $\delta > 0$ is a scaling parameter and $\lambda \in \mathbb{R}$ characterises the subclass, which essentially determines the weight in the tails.

Special cases of the GH distribution arise when $\lambda = 1$, in which case one gets the class of hyperbolic distributions, and when $\lambda = -1/2$, which yields the class of normal inverse Gaussian (NIG) distributions. Further, for $\lambda > 0$ and as $\delta \to 0$ one obtains a variance-gamma (VG) distribution. In the case $\lambda < 0$ and $\delta^2 = -2\lambda$, one obtains a skewed Student-t distribution with δ^2 degrees of freedom, (Eberlein and Hammerstein, 2004). The special case of the normal distribution thus arises when $\beta = 0$ and as $\delta \to \infty$.

A.2. *α-stable distribution*

A different generalisation of the normal distribution is the family of α-stable distributions. A (non-degenerate) random variable X is stable if for each $n > 0$ there exist constants c_n and γ_n such that $X_1 + \cdots + X_n \overset{\mathcal{L}}{=} c_n X + \gamma_n$. It turns out that $c_n = n^{1/\alpha}$ with $0 < \alpha \leq 2$. In particular, the normal distribution is the only 2-stable distribution, that is, a stable distribution with $\alpha = 2$. The normal distribution is the only α-distribution with a finite variance, so that other types of α-stable distributions are very heavy tailed. The development of the theory surrounding α-stable distributions goes back to Paul Lévy's work in the 1920s and 1930s, see e.g. (Feller, 1971; Cont and Tankov, 2004). A comprehensive overview of α-stable distributions in finance is given in (Rachev and Mittnik, 2000).

In the non-normal case, that is, when $\alpha < 2$, there is no closed-form for the α-stable distribution or density. The characteristic function, however, admits a closed form:

$$\int_{-\infty}^{\infty} \mathbf{e}^{itx}\,\mathrm{d}F(x) = \begin{cases} \exp\{-c^\alpha\,|t|^\alpha(1 - i\beta\,\mathrm{sign}(t)\tan(\pi\alpha/2)) + i\delta t\}, & \text{if } \alpha \neq 1, \\ \exp\{-c\,|t|(1 + i\beta\,2/\pi\,\mathrm{sign}(t)\ln|t|) + i\delta t\}, & \text{if } \alpha = 1, \end{cases}$$

where F is the corresponding α-stable distribution function, $\alpha \in (0,2]$ is the index of stability, $\beta \in [-1,1]$ determines the skewness, $\delta \in \mathbb{R}$ determines the location and $c > 0$ is the scale parameter.

APPENDIX B: VAR'S AND VAR SPREADS

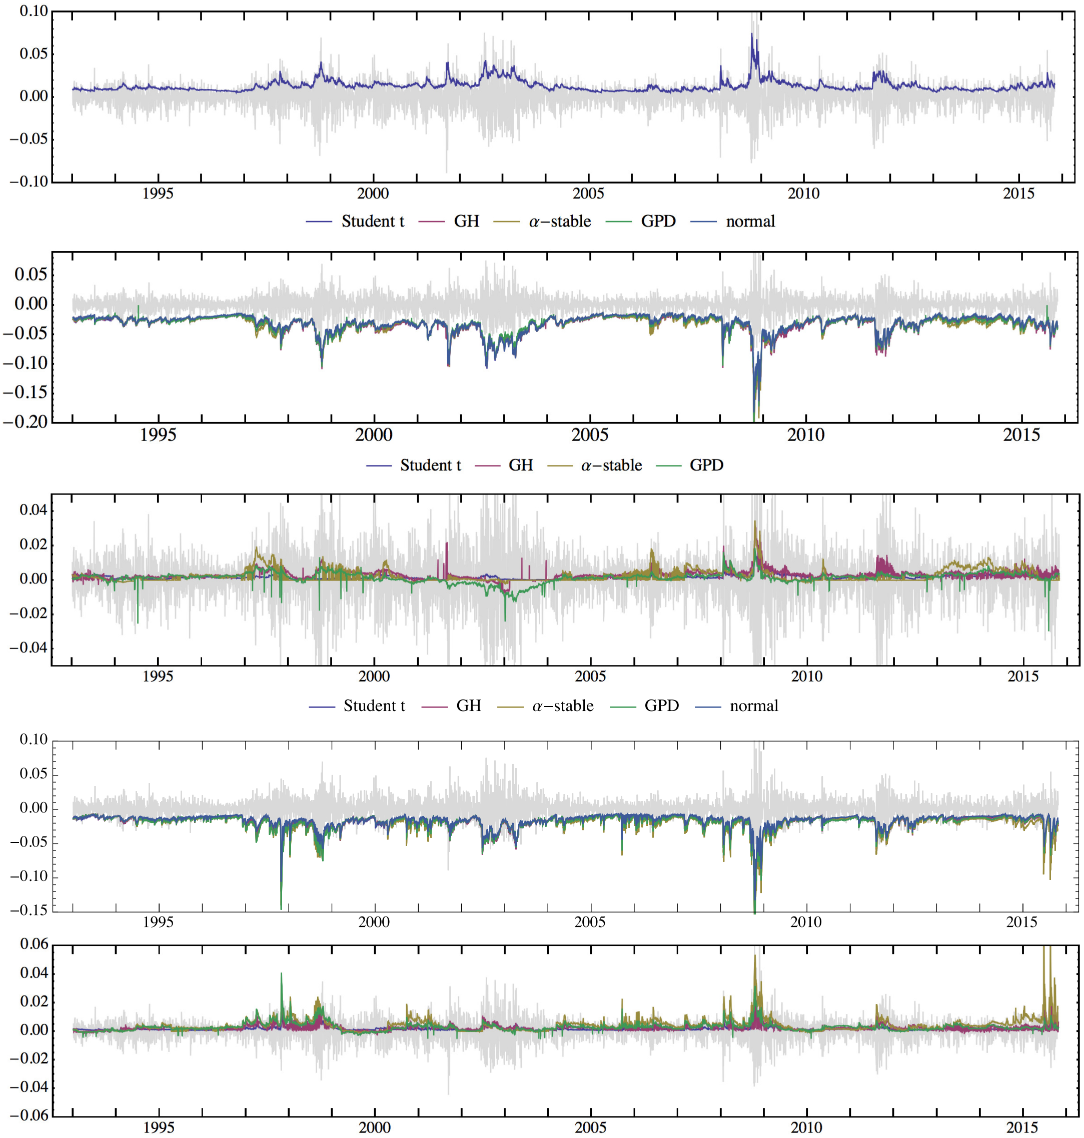

FIGURE 9.— From top: volatility forecast from GARCH model; daily value-at-risk at a confidence level of $\alpha = 0.99$; daily value-at-risk spread relative to VaR with normally distributed residuals; overnight VaR; overnight VaR spread. Right axis corresponds to DAX future returns.

IRTG 1792 Discussion Paper Series 2018

For a complete list of Discussion Papers published, please visit irtg1792.hu-berlin.de.

001 "Data Driven Value-at-Risk Forecasting using a SVR-GARCH-KDE Hybrid" by Marius Lux, Wolfgang Karl Härdle and Stefan Lessmann, January 2018.
002 "Nonparametric Variable Selection and Its Application to Additive Models" by Zheng-Hui Feng, Lu Lin, Ruo-Qing Zhu asnd Li-Xing Zhu, January 2018.
003 "Systemic Risk in Global Volatility Spillover Networks: Evidence from Option-implied Volatility Indices " by Zihui Yang and Yinggang Zhou, January 2018.
004 "Pricing Cryptocurrency options: the case of CRIX and Bitcoin" by Cathy YH Chen, Wolfgang Karl Härdle, Ai Jun Hou and Weining Wang, January 2018.
005 "Testing for bubbles in cryptocurrencies with time-varying volatility" by Christian M. Hafner, January 2018.
006 "A Note on Cryptocurrencies and Currency Competition" by Anna Almosova, January 2018.
007 "Knowing me, knowing you: inventor mobility and the formation of technology-oriented alliances" by Stefan Wagner and Martin C. Goossen, February 2018.
008 "A Monetary Model of Blockchain" by Anna Almosova, February 2018.
009 "Deregulated day-ahead electricity markets in Southeast Europe: Price forecasting and comparative structural analysis" by Antanina Hryshchuk, Stefan Lessmann, February 2018.
010 "How Sensitive are Tail-related Risk Measures in a Contamination Neighbourhood?" by Wolfgang Karl Härdle, Chengxiu Ling, February 2018.
011 "How to Measure a Performance of a Collaborative Research Centre" by Alona Zharova, Janine Tellinger-Rice, Wolfgang Karl Härdle, February 2018.
012 "Targeting customers for profit: An ensemble learning framework to support marketing decision making" by Stefan Lessmann, Kristof Coussement, Koen W. De Bock, Johannes Haupt, February 2018.
013 "Improving Crime Count Forecasts Using Twitter and Taxi Data" by Lara Vomfell, Wolfgang Karl Härdle, Stefan Lessmann, February 2018.
014 "Price Discovery on Bitcoin Markets" by Paolo Pagnottoni, Dirk G. Baur, Thomas Dimpfl, March 2018.
015 "Bitcoin is not the New Gold - A Comparison of Volatility, Correlation, and Portfolio Performance" by Tony Klein, Hien Pham Thu, Thomas Walther, March 2018.
016 "Time-varying Limit Order Book Networks" by Wolfgang Karl Härdle, Shi Chen, Chong Liang, Melanie Schienle, April 2018.
017 "Regularization Approach for NetworkModeling of German EnergyMarket" by Shi Chen, Wolfgang Karl Härdle, Brenda López Cabrera, May 2018.
018 "Adaptive Nonparametric Clustering" by Kirill Efimov, Larisa Adamyan, Vladimir Spokoiny, May 2018.
019 "Lasso, knockoff and Gaussian covariates: a comparison" by Laurie Davies, May 2018.

IRTG 1792, Spandauer Straße 1, D-10178 Berlin
http://irtg1792.hu-berlin.de

This research was supported by the Deutsche Forschungsgemeinschaft through the IRTG 1792.

IRTG 1792 Discussion Paper Series 2018

For a complete list of Discussion Papers published, please visit irtg1792.hu-berlin.de.

020 "A Regime Shift Model with Nonparametric Switching Mechanism" by Haiqiang Chen, Yingxing Li, Ming Lin and Yanli Zhu, May 2018.
021 "LASSO-Driven Inference in Time and Space" by Victor Chernozhukov, Wolfgang K. Härdle, Chen Huang, Weining Wang, June 2018.
022 " Learning from Errors: The case of monetary and fiscal policy regimes" by Andreas Tryphonides, June 2018.
023 "Textual Sentiment, Option Characteristics, and Stock Return Predictability" by Cathy Yi-Hsuan Chen, Matthias R. Fengler, Wolfgang Karl Härdle, Yanchu Liu, June 2018.
024 "Bootstrap Confidence Sets For Spectral Projectors Of Sample Covariance" by A. Naumov, V. Spokoiny, V. Ulyanov, June 2018.
025 "Construction of Non-asymptotic Confidence Sets in 2 -Wasserstein Space" by Johannes Ebert, Vladimir Spokoiny, Alexandra Suvorikova, June 2018.
026 "Large ball probabilities, Gaussian comparison and anti-concentration" by Friedrich Götze, Alexey Naumov, Vladimir Spokoiny, Vladimir Ulyanov, June 2018.
027 "Bayesian inference for spectral projectors of covariance matrix" by Igor Silin, Vladimir Spokoiny, June 2018.
028 "Toolbox: Gaussian comparison on Eucledian balls" by Andzhey Koziuk, Vladimir Spokoiny, June 2018.
029 "Pointwise adaptation via stagewise aggregation of local estimates for multiclass classification" by Nikita Puchkin, Vladimir Spokoiny, June 2018.
030 "Gaussian Process Forecast with multidimensional distributional entries" by Francois Bachoc, Alexandra Suvorikova, Jean-Michel Loubes, Vladimir Spokoiny, June 2018.
031 "Instrumental variables regression" by Andzhey Koziuk, Vladimir Spokoiny, June 2018.
032 "Understanding Latent Group Structure of Cryptocurrencies Market: A Dynamic Network Perspective" by Li Guo, Yubo Tao and Wolfgang Karl Härdle, July 2018.
033 "Optimal contracts under competition when uncertainty from adverse selection and moral hazard are present" by Natalie Packham, August 2018.
034 "A factor-model approach for correlation scenarios and correlation stress-testing" by Natalie Packham and Fabian Woebbeking, August 2018.
035 "Correlation Under Stress In Normal Variance Mixture Models" by Michael Kalkbrener and Natalie Packham, August 2018.
036 "Model risk of contingent claims" by Nils Detering and Natalie Packham, August 2018.
037 "Default probabilities and default correlations under stress" by Natalie Packham, Michael Kalkbrener and Ludger Overbeck, August 2018.
038 "Tail-Risk Protection Trading Strategies" by Natalie Packham, Jochen Papenbrock, Peter Schwendner and Fabian Woebbeking, August 2018.

IRTG 1792, Spandauer Straße 1, D-10178 Berlin
http://irtg1792.hu-berlin.de

This research was supported by the Deutsche Forschungsgemeinschaft through the IRTG 1792.