

Term structure of risk, the role of Known and Unknown Risks and Non-stationary Distributions

Riccardo Colacito* Robert F. Engle†

Abstract

In this paper we document the presence of a term structure of risk and we propose how to measure it using alternative models to forecast volatility and the Value at Risk at different horizons. We then quantify the benefits of an investor that is aware of the existence of a term structure of risk in the context of an asset allocation exercise.

First Draft: August 14, 2006. *This draft:* May 29, 2007.

*The University of North Carolina at Chapel Hill, Kenan Flagler School of Business.

†New York University, Stern School of Business.

1 Introduction

Many people believe that there are grave risks facing our financial markets. These include the massive budget deficits, the balance payments deficits, the high cost of energy and many other raw materials, the uncertainty over FED policy, war in Iraq that is going badly, global warming and the extraordinary amount of US Debt that is held by the Chinese government. In addition, there is a concern that the vast global derivatives market, the number of unregulated hedge funds, the merging of financial markets across national borders and the explosive growth of private equity funds, make the financial system more unstable and susceptible to meltdown. These concerns are not new but have been serious topics of discussion for several years.

The extraordinary fact however is that the volatility of financial markets today is about as low as it has ever been. This has been true for most of the years 2004-2006. This is the situation in the US equity market but it is also true in global equity markets. The volatility has fallen to very low levels in most equity markets around the globe as shown in figure 1. It is also reflected in options prices as can easily be seen by looking at a time series plot of the volatility index, VIX.

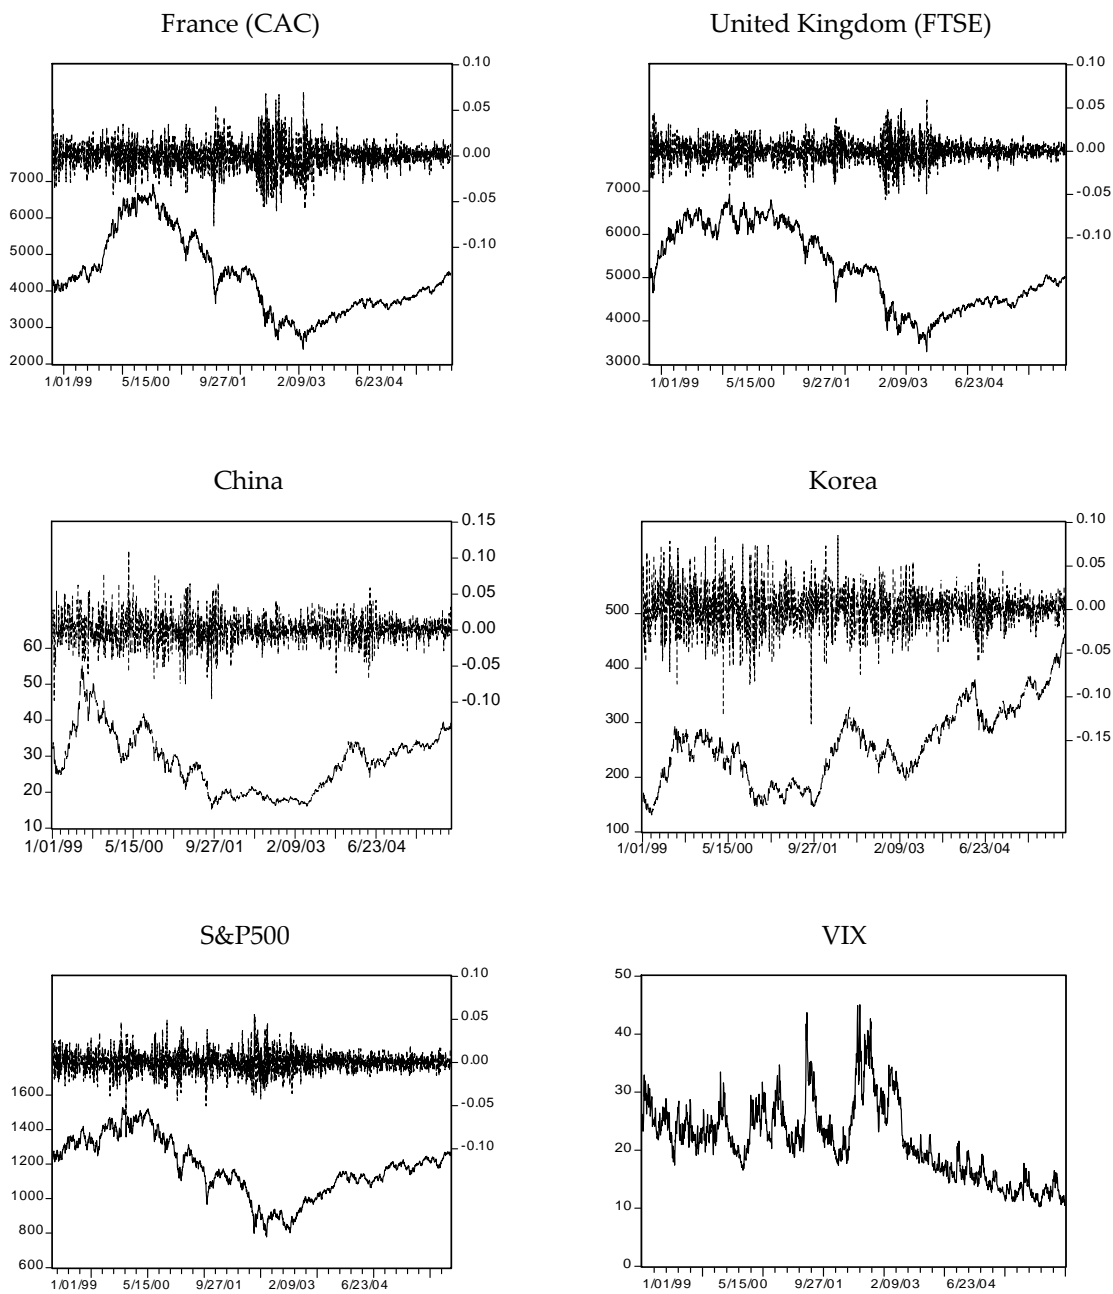


FIG. 1 - Volatilities of global equity markets. The first five plots report the returns (top line) and the levels (bottom line) of equity markets in France, United Kingdom, China, Korea and United States. The sixth plot shows the volatility index in the US.

These observations present a puzzle. Are financial markets ignoring these risks or are the risks not so serious? In this paper we present another resolution of the puzzle. Most of these risks are potential problems for the future. They are not risks in the short run, only in the long run. There may be a term structure of risk that faces financial markets in general and individual investors in particular. This concept must be carefully defined and examined empirically. Finally we must consider the implications for asset pricing and portfolio construction.

2 Measurement

In this paper we associate long run risks with the probability and magnitude of losses of a passive portfolio over a long horizon. Measuring this in nominal terms is only appropriate if the changes in price level or purchasing power of risk free rates are minor adjustments. The analysis could be carried out in any of these frameworks. We choose nominal returns to focus on the dynamics of the financial markets rather than the nominal economy as a whole.

In contrast Bansal and Yaron (2004) introduce long run risks by postulating a slowly varying factor in real consumption that induces variation in expected returns. The long run risk is thus the risk of a low consumption state which corresponds to a low

return state. Without further elaboration, the prediction of this risk in the distant future would not be changing over time as current information would have little ability to predict these events. Conceptually, a model more similar to ours would introduce the long run risks into the variance of consumption, rather than its level.

To quantify these long run risks, we follow Timmermann and Guidolin (2006). We consider the long run variance and long run value at risk, LRVaR. These measures are widely used in financial planning, but can be given a new interpretation with long horizons. Unknown and unforecastable risks appear in the historical data as surprising returns and are therefore a part of predicted variance and VaR. Non-stationary risks can sometimes be corrected for and therefore be used to improve risk assessment and decision making.

2.1 Volatility forecasts at various horizons

The task of forecasting volatility is one that can only be accomplished after a model has been specified. But what is the reasonable set of assumptions that one can make about the underlying economic model? It is common to assume that returns follow a stationary process, with the understanding that this is a statistical convenience and not an economic model. With stationary returns, long run risk is constant. This can

be shown in a simple example that allows us to introduce some of the notation that we will be using in the rest of the analysis. Let r_t be a mean zero random variable measuring the return on a financial asset and assume that it follows a GARCH(1,1) process:

$$\begin{aligned} r_t &= \sqrt{h_t} \varepsilon_t, \quad \varepsilon_t \sim N(0, 1) \\ h_t &= (1 - \alpha - \beta) \omega + \beta h_{t-1} + \alpha r_{t-1}^2 \end{aligned} \tag{1}$$

Taking the unconditional expectation of squared returns, we obtain

$$E[r_t^2] = E[h_t] = \omega$$

that is our constant estimate of long run risk. Long run risk is the time average of short run risk and the unconditional term structure of volatility risk is proportional to \sqrt{T} .

Nevertheless long run risk can change over time or at least there is no *a priori* reason to restrict our statistical model from this possibility. As a matter of fact, unknown and unknowable events can occur and if *ex post* we say that there is a shift in the distribution, *ex ante* we must assess the probabilities. The important question is whether the historical risks can be used to assess the future risks and this is a

question of the stationarity of the distribution. If the distribution is stationary, then unknown and unknowable risks are already sensibly incorporated in the forecasts of future risk. But if the distribution is changing, then these changes must be modeled.

An example of a model that allows for time varying long run risk is the spline GARCH of Engle and Rangel (2005), in which economic or exogenous variables such as recessions, inflation and macro volatility increase the long run variance. This is a multiplicative model in which the conditional variance is assumed to be the product of a long run and of a short run components and where both terms can be time varying. In particular, mean-zero returns follow the process:

$$\begin{aligned} r_t &= \sqrt{\tau_t g_t} \varepsilon_t, \quad \varepsilon_t \sim N(0, 1) \\ g_t &= (1 - \alpha - \beta) + \alpha \left(\frac{r_{t-1}^2}{\tau_{t-1}} \right) + \beta g_{t-1} \end{aligned} \tag{2}$$

where τ_t is a function of time and exogenous variables. By taking unconditional expectations of squared returns

$$E[r_t^2] = \tau_t E[g_t \varepsilon_t^2] = \tau_t E[g_t] = \tau_t$$

it is clear that τ_t can be interpreted as the long run forecast of variance. We will also refer to this component as the low frequency variance or sometimes the unconditional




FIG. 2 - Long vs. short run volatility of the S&P500. The thick line is short run volatility and the smooth thin line is the long run volatility.

variance when it is a function of deterministic or exogenous variables. One possibility is that the long run variance τ_t is an exponential quadratic spline of time:

$$\log(\tau_t) = \omega_0 + \omega_1 + \omega_2 + \sum_{k=1}^K \theta_k [\max(t - t_k, 0)]^2$$

Figure 2 reports the measure of short run and long run volatility for the S&P500 forecasted by the spline GARCH model. The figure shows how there may be periods in which the short run risk (thick line) is high, while the long run risk is low and

viceversa. The picture also shows how the current date volatility appears to be at a record low level, while long run volatility is higher. This is the case not only for the US, but also for a large number of countries, as it is documented in figure 3.

2.2 The term structure of value at risk

The Value at Risk T periods ahead from the current date is the α quantile of the conditional distribution of returns at time $t + T$. In formulas:

$$Pr_t(r_{t+T} \leq -VaR_{t+T}^\alpha) = \alpha$$

As a benchmark we will consider the case of *i.i.d.* mean zero returns: $r_t \sim N(0, h)$, $\forall t$. In this situation the value at risk is simply proportional to the square root of time:

$$VaR_{t+T}^\alpha = \sqrt{hT} \Phi^{-1}(\alpha) \tag{3}$$

It is often convenient to standardize the measure reported in (3) by \sqrt{T} , in which




FIG. 3 - Measures of volatility. In each subplot, the thick line represents the conditional volatility, the thin line is the unconditional volatility and the squares are the annualized realized volatility. Each panel is for a different countries. From the top left to the bottom right: India, Argentina, Japan, Brazil, South Africa and Poland.

case *i.i.d.* returns are equivalent to a constant term structure of risk.

When returns are not *i.i.d.*, the term structure of VaR can slope upward or downward. An interesting case to consider is the one in which returns follow a TARCH(1,1) process:

$$\begin{aligned} r_t &= \mu + \sqrt{h_t} \varepsilon_t \\ h_t &= \omega + \beta h_{t-1} + \alpha r_{t-1}^2 + \gamma r_{t-1}^2 I_{(r_{t-1} < 0)} \end{aligned} \tag{4}$$

The law of motion of the conditional variance is such that following periods of negative returns there is an expectation for a relatively higher variance in the future. Although one-period returns are symmetrically distributed at each point in time, multi-period returns are not as illustrated in figure 4. The probability that is attached to the extreme negative events that may occur many periods in the future has potentially important consequences that should be taken into account in the context of any asset allocation exercise.

Table 1 reports the estimate of the parameters of model (4) for a long dataset




FIG. 4 - Asymmetric volatility: binomial tree. At each node returns have a symmetric distribution, but following periods of positive returns the volatility lowers, while after periods of negative returns the volatility increases. This implies an asymmetric distribution of multi-period returns.

of daily observations on the S&P500, ranging from 1950 to 2006. As expected the asymmetric volatility parameter is positive and significant at a 95% level of confidence. Negative shocks have 3 times the effect of positive shocks in forecasting future variances.

Diebold, Hickman, Inoue, and Schuermann (1998) show that the common practice of converting one day volatilities to T-day estimates by scaling by \sqrt{T} is inappropriate and produces overestimates of the variability of long-horizon volatility. Our




FIG. 5 - Value at Risk of a TARCH Gaussian simulation. The left panel reports the VaR for T ranging from 0 to 100 of a TARCH(1,1) simulated from table 1 as a thick line, while the dotted line is the benchmark case of *i.i.d.* returns. The right panel reports $\frac{VaR_{t+h}^\alpha}{\sqrt{T}}$.

work relates to theirs in that we quantify the impact of a TARCH(1,1) volatility process on the estimate of the VaR at a given future point. We address this task by simulating one million excess returns following process (4) and calibrating its parameters according to table 1. We then let T vary between 1 and 100 and construct the corresponding 1% Value at Risk. Figure 5 reports the results of this simulation. In comparing the VaR when returns follow an asymmetric process (thick line) with the VaR obtained under the assumption that returns are *i.i.d.* (dashed line) it is apparent that according to this risk measure, both long run risk and short run risk according to (4) exceed the risk for *i.i.d.* shocks. Particularly important however is the fact that this difference increases with horizon. That is, the term structure of risk can slope upward all the time.

TABLE 1
TARCH(1,1): ESTIMATION OF PARAMETERS.

	ω	α	β	γ
Estimate	8.36×10^{-7}	0.035	0.918	0.074
Std. Err.	6.40×10^{-7}	0.003	0.003	0.002

Notes - The sample period is 1950-2006.

3 Implications for asset allocation

It has recently emerged that volatility timing and traditional market timing are fundamentally related, as well documented by Christoffersen and Diebold (2006). Fleming, Kirby, and Ostdiek (2001) and Fleming, Kirby, and Ostdiek (2003) study a one-day horizon asset allocation problem and document the economic value of various conditional volatility estimators and realized volatilities. Engle and Colacito (2006) pointed out that correct volatility and correlation timing is typically worth 50-60 basis points when the investment horizon is one day. However most portfolio managers have investment horizons longer than a day, even though they ultimately end up doing a static asset allocation exercise. It seems reasonable to think that an investor, aware that returns follow the TARCH process that we discussed in the earlier sections, would take the presence of a downside risk into account when choosing portfolio weights in this context. In this section we give a quantitative answer to the question of how much can an investor expect from optimally adjusting

portfolio weights when the variance is asymmetric.

We shall focus on the simplest case in which the investor can only allocate her wealth among a risky and a risk-less asset. We denote with w_{t+T} the share of the portfolio that is invested in the risky asset between time t and $t+T$ and with r_{t+T} the logarithm of the continuously compounded return on the risky asset in excess of the risk free rate r^f between t and $t+T$. We assume that the agent wants to maximize terminal wealth according to an exponential utility function:

$$\max_{w_{t+T}} E_t - \exp \{ -b (w_{t+T} r_{t+T}) \} \exp \{ -b r^f \} \quad (5)$$

where b is a preference parameter that reflects the absolute risk aversion. The risk free rate is constant at a daily frequency. If log-returns are conditionally distributed as normals, an investor seeking to maximize her utility according to (5) could simply solve a mean-variance exercise:

$$\max_{w_{t+T}} E_t [w_{t+T} r_{t+T}] - \frac{b}{2} E_t [w_{t+T}^2 r_{t+T}^2] \quad (6)$$

However if returns are not lognormally distributed, the equivalency of problems (5) and (6) does not hold anymore. We now develop an approximate procedure to choose portfolios according to (5) when the returns follow an asymmetric GARCH

model. The first step is to approximate the utility function accounting for higher moments. The result, derived in the appendix, is:

$$\max_{w_{t+T}} - \exp \{-bw_{t+T}\mu_{t+T}\} \left[1 + \frac{b^2}{2}w_{t+T}^2h_{t+T} - \frac{b^3}{6}w_{t+T}^3s_{t+T} \right] \quad (7)$$

where μ_{t+T} , h_{t+T} and s_{t+T} denote the conditional expectations of mean, variance and third centered moment, respectively. This utility function formalizes the idea according to which investors like positive first and third moments and dislike second moments. Alternatively, agents are now concerned about the lower tail of the distribution that is depicted in figure 4. The solution can be calculated numerically based on the forecast first, second and third moments. This optimization is simple, but does not produce a closed form solution. In the experiment described below, the mean is constant. To forecast the third central moment we use a recursion developed in the appendix. Essentially it computes $E_t [(r_{t+1} + r_{t+2} + \dots + r_{t+k})^3]$ in terms of $E_t [h_{t+k}^{3/2}]$. Then approximating this by Taylor series, the third moment can be forecast and used to optimize portfolios at each point in time. Clearly, the more negative the third moment, the less exposure to the risky asset will be chosen by this investor.

In order to quantify the benefit of knowing that returns follow an asymmetric volatility process, we simulate daily returns according to model (4) and then we compare

two investors with the same objective function (5): one makes forecasts of the distribution of returns based on the TARCH(1,1) process reported in (4), while the other one believes that returns are distributed according to the GARCH(1,1) process in (1).¹ For the results to be comparable, we will assume that the two models agree on the unconditional forecasts of mean returns and variance.²

The metric that we adopt to quantify these benefits is based on the criterion function 5. For a given risk free rate \tilde{r}^f , an agent that refrains from investing in the risky asset would obtain an average utility $U(\tilde{r}^f) = -\exp\{-b\tilde{r}^f\}$. By allocating a non-zero share of her portfolio in the risky asset at the actual risk-free rate r^f , she could instead expect a utility $U(r^f) = -E[\exp\{-b(w_{t+T}r_{t+T})\} \exp\{-br^f\}]$. The risk-free rate \tilde{r}^f that would make her indifferent between the two strategies can be easily shown³ to be equal to:

$$\tilde{r}^f = r^f + \frac{-\log E \exp\{-bw_t r_t\}}{b} \quad (8)$$

Hence our evaluation strategy that consists in obtaining sequences of optimal portfolio weights $\{w_{1,t}\}_{t=1}^T$ and $\{w_{2,t}\}_{t=1}^T$ based on forecasting from (4) and (1) respectively and then comparing the ‘certainty equivalent’ returns \tilde{r}_1^f and \tilde{r}_2^f based on the sample

¹More precisely, we use these models to describe the process of returns in excess of their mean.

²The appendix also reports the details on how to compute multi-period forecasts of third centered moments.

³We document this in the Appendix.

counterparts of the terms involving an expectation. It is natural to expect that \tilde{r}_2^f will typically be greater than \tilde{r}_1^f : we want to quantify this benefit.

Figure 6 reports the percentage annualized gains when the investment horizons are 20 days (left panel) and 1 year (right panel). A number like 1 on the vertical axis means that an investor that is informed of the asymmetry in the volatility process and optimally adjusts portfolio weights would need 100 basis points in excess of what an investor that ignores the asymmetry would need in order to refrain from investing in the risky asset. We plot these gains for increasing values of the coefficient of absolute risk aversion, b . The average gain can be as high as 220 basis points and it is decreasing with b just because the amount of wealth that is invested in the risky asset is decreasing with risk aversion. In moving from a 20 day to a year long exercise, there is still a sizeable gain to be made. The decrease of the average benefit has presumably to be attributed to the difficulty of accurately forecasting the distribution of multi-period returns as the horizon gets longer and longer.

Although this represent the outcome of the simplest example of portfolio allocation, the results reported in this section are suggestive of the fact that there is potentially a considerable gain that can be obtained by appropriately timing volatility over horizons that are longer than a day. Along these lines it is not hard to imagine that multivariate asset allocation experiments would yield even larger gains, that must




FIG. 6 - Annualized percentage gain from volatility timing when the investment horizon is 20 days (left panel) and 252 days (right panel). The vertical axis reports the extra return that an investor that is aware of the asymmetry of the volatility process could obtain. The horizontal axis reports the preference parameter b . The thick line is the average gain, while the dashed lines represent the 95% confidence interval.

be taken into account as the planning horizon increases.

4 Concluding remarks

In this paper we have documented the presence of a term structure of risk and provided tools that can be used by academics and practitioners to actively manage portfolios in the presence of downward risk. The implications in the context of a simple asset allocation exercise are suggestive of the fact that taking into account time varying asymmetries in the multi-period distributions of asset returns can po-

tentially result in significant financial gains. This provides a useful starting point for the exploration of the benefits that can be obtained in the context of large scale systems.

References

- Bansal, R. and A. Yaron (2004). Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance* 59, 1481–1509.
- Christoffersen, P. and F. Diebold (2006). Financial asset returns, direction-of-change forecasting, and volatility dynamics. *Management Science* 52, 1273–1288.
- Diebold, F., A. Hickman, A. Inoue, and T. Schuermann (1998). Scale models. *Risk Magazine* 11, 104–107.
- Engle, R. and R. Colacito (2006). Testing and valuing dynamic correlation for asset allocation. *Journal of Business and Economic Statistics* 24(2).
- Engle, R. and G. Rangel (2005). The spline garch model for unconditional volatility and its global macroeconomic causes. *working paper*.
- Fleming, J., C. Kirby, and B. Ostdiek (2001). The economic value of volatility timing. *Journal of Finance* 56, 329–352.
- Fleming, J., C. Kirby, and B. Ostdiek (2003). The economic value of volatility timing using realized volatility. *Journal of Financial Economics* 67, 473–509.
- Timmermann, A. and M. Guidolin (2006). Term structure of risk under alternative economic specifications. *Forthcoming in Journal of Econometrics*.

A Approximation of the utility function

Following the notation of section 3, the utility function is:

$$U_t = E_t - \exp\{-bW_{t+T}\}$$

where $W_{t+T} = w_{t+T}r_{t+T}$. A third order Taylor expansion around $w_{t+T}\mu$ delivers:

$$\begin{aligned} U_t &\approx -E_t \exp\{-bw_{t+T}\mu\} \left[1 - bw_{t+T}(r_{t+T} - \mu) + \frac{b^2}{2}w_{t+T}^2(r_{t+T} - \mu)^2 - \frac{b^3}{6}w_{t+T}^3(r_{t+T} - \mu)^3 \right] \\ &= -\exp\{-bw_{t+T}\mu\} \left[1 + \frac{b^2}{2}w_{t+T}^2h_{t+T} - \frac{b^3}{6}w_{t+T}^3s_{t+T} \right] \end{aligned}$$

where s_{t+T} denotes the conditional third centered moment of the distribution of r_{t+T} . This is the analytical form we worked with in section 3.

B Multi-period forecasts of second and third moments

Given the following process for the logarithm of excess returns:

$$\begin{aligned} r_t &= \sqrt{h_t} \varepsilon_t \\ h_t &= \omega + \alpha r_{t-1}^2 + \beta h_{t-1} + \gamma I_{r_{t-1} < 0} r_{t-1}^2 \end{aligned}$$

the conditional forecast of the variance of multi-period returns can be computed as:

$$\begin{aligned} E_t \left[\left(\sum_{j=1}^T r_{t+j} \right)^2 \right] &= \sum_{j=1}^T E_t [r_{t+j}^2] \\ &= \sum_{j=1}^T E_t h_{t+j} \\ &= h_{t+1} + \sum_{j=2}^T \left[\omega \sum_{i=0}^{j-2} \left(\alpha + \beta + \frac{\gamma}{2} \right)^i + \left(\alpha + \beta + \frac{\gamma}{2} \right)^{j-1} h_{t+1} \right] \end{aligned}$$

We shall denote the third centered conditional moment as:

$$s_{t+j} = E_t \left[\left(\sum_{i=1}^j r_{t+i} \right)^3 \right]$$

The one period ahead third moment is equal to zero:

$$s_{t+1} = E_t \left[h_{t+1}^{3/2} \varepsilon_{t+1}^3 \right] = 0$$

The conditional third moment of two periods continuously compounded returns is:

$$\begin{aligned} s_{t+2} &= E_t \left[(r_{t+1} + r_{t+2})^3 \right] \\ &= 3E_t \left[r_{t+1} r_{t+2}^2 \right] \\ &= 3E_t \left[\sqrt{h_{t+1}} \varepsilon_{t+1} \left(\omega + \alpha h_{t+1} \varepsilon_{t+1}^2 + \beta h_{t+1} + \gamma I_{\varepsilon_{t+1} < 0} h_{t+1} \varepsilon_{t+1}^2 \right) \right] \\ &= -\frac{12}{5} \gamma h_{t+1}^{3/2} \end{aligned}$$

Similarly:

$$\begin{aligned} s_{t+3} &= s_{t+2} + 3\sqrt{h_{t+1}} E_t \left[\varepsilon_{t+1} h_{t+3} \right] + 3E_t \left[\sqrt{h_{t+2}} \varepsilon_{t+2} h_{t+3} \right] \\ &= s_{t+2} + 3\sqrt{h_{t+1}} E_t \left[\varepsilon_{t+1} \left(\omega + \alpha h_{t+2} + \beta h_{t+2} + \frac{\gamma}{2} h_{t+2} \right) \right] - \frac{12}{5} \gamma E_t \left[h_{t+2}^{3/2} \right] \\ &= s_{t+2} + 3 \left(\alpha + \beta + \frac{\gamma}{2} \right) \sqrt{h_{t+1}} E_t \left[\varepsilon_{t+1} h_{t+2} \right] - \frac{12}{5} \gamma E_t \left[h_{t+2}^{3/2} \right] \\ &= s_{t+2} + 3 \left(\alpha + \beta + \frac{\gamma}{2} \right) (s_{t+2} - s_{t+1}) - \frac{12}{5} \gamma E_t \left[h_{t+2}^{3/2} \right] \end{aligned}$$

and

$$s_{t+j} = s_{t+j-1} + \left(\alpha + \beta + \frac{\gamma}{2} \right) (s_{t+j-1} - s_{t+j-2}) - \frac{12}{5} \gamma E_t \left[h_{t+j-1}^{3/2} \right], \quad \forall j \geq 4$$

where $E_t \left[h_{t+j}^{3/2} \right]$ can be approximated to a first order as:

$$\begin{aligned} E_t \left[h_{t+j}^{3/2} \right] &\approx \left(\bar{h}^{3/2} - \frac{3}{2} \bar{h}^{-3/2} \right) + \frac{3}{2} \bar{h}^{-1/2} E_t h_{t+j} \\ &= k_0 + k_1 E_t h_{t+j} \end{aligned}$$

C Computation of the certainty equivalent risk-free rate

Given the utility function discussed in this paper, a sequence of portfolio weights $\{w_t\}_{t=1}^T$ and the actual risk free rate r^f deliver the following expected utility:

$$U(r^f) = -E[\exp\{-bw_t r_t\}] \exp\{-br^f\} \quad (9)$$

An agent that allocates all of her wealth in the risk-less asset at the rate \tilde{r}^f obtains:

$$U(\tilde{r}^f) = -\exp\{-b\tilde{r}^f\} \quad (10)$$

The rate \tilde{r}^f that makes the investor indifferent between (9) and (10) is computed as:

$$\begin{aligned} U(r^f) &= U(\tilde{r}^f) \\ -\log E[\exp\{-bw_t r_t\}] \exp\{-br^f\} &= -\log \exp\{-b\tilde{r}^f\} \\ -\log E[\exp\{-bw_t r_t\}] + br^f &= b\tilde{r}^f \end{aligned}$$

from which it follows:

$$\tilde{r}^f = r^f + \frac{-\log E[\exp\{-bw_t r_t\}]}{b}$$