



SCHOOL OF ECONOMICS AND MANAGEMENT
FACULTY OF SOCIAL SCIENCES
AARHUS UNIVERSITY



CREATES
Center for Research in Econometric
Analysis of Time Series

CREATES Research Paper 2010-38

Predictable return distributions

Thomas Q. Pedersen

School of Economics and Management
Aarhus University
Bartholins Allé 10, Building 1322, DK-8000 Aarhus C
Denmark

Predictable return distributions*

Thomas Q. Pedersen[†]

July 2010

Abstract

This paper provides detailed insights into predictability of the entire stock and bond return distribution through the use of quantile regression. This allows us to examine specific parts of the return distribution such as the tails or the center, and for a sufficiently fine grid of quantiles we can trace out the entire distribution. A univariate quantile regression model is used to examine stock and bond return distributions individually, while a multivariate model is used to capture their joint distribution. An empirical analysis on US data shows that certain parts of the return distributions are predictable as a function of economic state variables. The results are, however, very different for stocks and bonds. The state variables primarily predict only location shifts in the stock return distribution, while they also predict changes in higher-order moments in the bond return distribution. Out-of-sample analyses show that the relative accuracy of the state variables in predicting future returns varies across the distribution. A portfolio study shows that an investor with power utility can obtain economic gains by applying the empirical return distribution in portfolio decisions instead of imposing an assumption of lognormally distributed returns.

JEL Classification: C21, C31, G11, G12, G17

Keywords: Return predictability, return distribution, quantile regression, multivariate model, out-of-sample forecast, portfolio choice

*This research is supported by *CREATES* (Center for Research in Econometric Analysis of Time Series), funded by the Danish National Research Foundation. Detailed comments from Tom Engsted are gratefully acknowledged.

[†]*CREATES*, School of Economics and Management, Aarhus University, Building 1322, DK-8000 Aarhus C, Denmark. E-mail: tqpedersen@creates.au.dk.

1 Introduction

Return predictability has been one of the most debated and analyzed topics within the financial literature during the last 20 years. However, despite still being at the forefront of the current research agenda, surprisingly little is known about predictability of other parts of the return distribution than the conditional mean and variance. In many areas of financial economics knowledge is required of either the entire return distribution or other parts of the distribution than the conditional mean. In asset pricing higher-order moments such as skewness and kurtosis have proven useful to explain variation in stock returns (see e.g. Harvey and Siddique, 2000, and Dittmar, 2002). In risk management, focus is usually on the lower tails of the return distribution. In portfolio management under standard preferences such as constant relative risk aversion, investors generally require an estimate of the entire distribution of future returns. Hence, understanding return predictability in more detail has great economic importance in many areas of financial economics.

This paper goes beyond predictability of the conditional mean and variance and examines predictability of the entire stock and bond return distribution as a function of economic state variables. The idea is to use quantile regression as originally introduced by Koenker and Basset (1978), which enables us to analyze predictability of different parts of the distribution as captured by specific quantiles. For example, by choosing the 0.05-, 0.50-, and 0.95-quantile, respectively, we can examine predictability of the lower tail, center, and upper tail of the distribution. Choosing a sufficiently fine grid of quantiles, we can analyze predictability of the entire distribution. The issue of predictable return distributions is not confined to individual assets. It is also highly relevant for joint distributions, which are used extensively in, for example, portfolio decisions and risk management. Only in the special case of independence between assets, can we use a combination of estimates from univariate quantile regressions. In general, we need an approach that can take the dependence structure into account when providing an estimate of the joint distribution. I use the multivariate quantile regression model introduced by Chakraborty (2003) to analyze predictability of the joint stock and bond return distribution as a function of economic state variables. The multivariate model allows us to specify quantiles in a manner similar to the univariate model, but with the difference that we now need to specify quantile combinations representing different parts of the *joint* distribution. For example, by choosing the (0.05,0.05)-quantile, we can examine predictability of the joint lower tail of the distribution, while the (0.05,0.95)-quantile represents the joint outcome of the return on the first asset falling in the lower part of the distribution and the return on the second asset falling in the upper part of the distribution. By choosing a sufficiently fine grid of quantile combinations, we can examine predictability of the entire joint distribution.

Predictability is evaluated both in- and out-of-sample. It is very likely that empirical forecasting models are, to some extent, almost always misspecified in the sense that the functional form of the model may be incorrect and/or that all the relevant state variables are not included compared to the true data-generating process. The relevant question

in the present context is, thus, not if the economic state variables correctly predict the different parts of the return distribution, but if any of the state variables perform better than others, and if the relative accuracy of competing forecasts is different across the distribution. It is not unlikely that a state variable is a relatively good forecaster of, for example, the lower tail of the distribution, while at the same time it is a relatively bad forecaster of the upper tail. Likewise, it is also possible that the relative accuracy of competing forecasts is very different across the tails and the center of the distribution.

Besides enabling us to evaluate predictability of the joint return distribution, the multivariate quantile model has an additional interesting feature. In the univariate quantile model the proportion of observations falling below the α -quantile is α by construction. In the multivariate quantile model the corresponding proportion is unknown from the outset except in the special case of independence between assets. By comparing the actual proportion to the proportion given independence, we can obtain detailed insights into the dependence structure between assets and map out in which part of the joint distribution, the dependence is present. In general, the multivariate quantile regression model gives a very flexible approach to modeling time-varying joint distributions that do not depend on any specific parametric assumptions.

As a final contribution, this paper contains an out-of-sample portfolio study to evaluate if it bears economic significance to model the empirical return distribution instead of imposing a distributional assumption, which is often done in portfolio decisions. The investor is assumed to have power utility, which generally implies that we need to obtain an estimate of the entire return distribution, except if returns are lognormally distributed. The combination of power utility and lognormality implies that we only need estimates of the first two moments of the return distribution. I compare the certainty equivalent return (CER) based on knowledge of the entire return distribution obtained from quantile regression with the CER based on the assumption that returns are lognormally distributed.

Based on monthly US data from 1941:5 to 2008:12, the main findings of this paper are as follows. First, for both stocks and bonds it is possible to predict certain parts of the return distribution, also even if the conditional mean is unpredictable. However, there are large differences between stocks and bonds. Conditional on the set of state variables used in this paper, the slope coefficients for stock returns are very similar across the distribution, which is also verified by formal testing. This implies that the state variables, generally, only predict location shifts in the stock return distribution. In contrast, the slope coefficients for bond returns are very different across the distribution. Furthermore, this difference shows itself in quite different ways, which implies that the state variables predict very different changes in the bond return distribution, including changes in volatility and skewness. Second, the joint distribution between stock and bond returns is also predictable. Based on the multivariate model, I furthermore show that stocks and bonds are dependent, which implies that a simple combination of quantile forecasts from the univariate setup will not suffice when attempting to capture the joint distribution. This is also verified in an out-of-sample analysis. Third, the out-of-sample analysis shows that the state variables' relative predictive ability varies a lot over the return distribu-

tion. This holds for both individual and joint distributions, and suggests that in order to obtain the best possible forecast of a future return distribution, a combination of different state variables might be preferable compared to single state variables. Fourth, the portfolio study shows that it can be associated with economic gains for an investor with power utility to model the empirical return distribution in portfolio decisions compared to imposing a lognormal distributional assumption about returns. The results are, however, very different for stocks and bonds. For stocks, where the state variables basically only capture location shifts, the CER is highest under the lognormal assumption. For bonds, where the state variables capture time-variation also in higher-order moments, the CER is highest under the empirical return distribution.

This paper is closely related to a recent paper by Cenesizoglu and Timmermann (2008). They were the first to propose the use of quantile regression to examine predictability of the entire return distribution. In an empirical application they apply the approach to stock returns and show that while the conditional mean is often unpredictable, predictability of different parts of the distribution, such as the tails and shoulders, shows itself more clearly. They also examine the economic significance of predictability of stock return quantiles through an asset allocation exercise. The present paper is motivated by the idea introduced by Cenesizoglu and Timmermann (2008) and extends their analysis in a number of ways. Regarding methodology, this paper proposes to model joint distributions through the use of multivariate quantile regression. Regarding the empirical analysis, this paper examines predictability of both stock and bond return distributions, including their joint distribution. Furthermore, it evaluates the relative accuracy of the state variables in forecasting the different parts of the distribution out-of-sample. Finally, it compares the optimal portfolio choice based on the empirical return distribution captured through quantile regression with the portfolio choice obtained by imposing a distributional assumption on returns.

Besides the paper by Cenesizoglu and Timmermann (2008), the present paper is related to many different strings of the financial literature. Naturally, the paper is closely related to the literature on predictability of the mean or volatility of stock and bond returns, especially the papers examining predictability as a function of different state variables; see e.g. Ilmanen (1995), Kirby (1997), Marquering and Verbeek (2004), Welch and Goyal (2008), Paye (2009), Lettau and Ludvigson (2010), and Viceira (2010). By choosing a very fine grid of quantiles when performing quantile regression, we can obtain an estimate of the entire future return distribution. Hence, the approach facilitates forecasts of the entire distribution instead of only specific moments, and in that sense the paper is also related to the literature on density forecasting; see e.g. Tay and Wallis (2002) and Corradi and Swansson (2006). As previously mentioned the multivariate quantile model used in this paper provides a very flexible way to model joint distributions. The need to conduct multivariate modeling has spurred the use of, for example, copulas (see e.g. Patton, 2004) and regime-switching models using mixtures of multivariate normal distributions (see e.g. Ang and Bekaert, 2002, and Guidolin and Timmermann, 2007). This paper complements the literature on multivariate modeling by introducing a non-parametric alternative to the existing methods. The multivariate quantile model is used

to model the joint distribution between stock and bond returns, and through the model’s feature of mapping out in which part of the joint distribution the dependence is present, the paper supplements the rapidly growing literature on stock and bond comovement; see e.g. Ilmanen (2003), Connolly et al. (2005), Guidolin and Timmermann (2006), Christiansen and Rinaldo (2007), and Baele et al. (2010). Finally, the paper adds to the growing number of papers using quantile regression in finance. Quantile regression has, for example, been used to classify investment styles (Basset and Chen, 2001), test the capital asset pricing model (Barnes and Hughes, 2002), and directly model the quantile in value-at-risk models (Engle and Manganelli, 2004).

The rest of the paper is organized as follows. In Section 2, I outline the univariate and multivariate quantile regression models used throughout the empirical part of the paper. Section 3 describes the data. Section 4 and 5 present the empirical results based on univariate and multivariate models, respectively. Section 6 shows results from the portfolio study, and finally, Section 7 contains some concluding remarks. The Appendix describes the algorithm used to estimate the multivariate quantile regression model.

2 Modeling the return distribution

In the following, I outline both the univariate and multivariate models, and the estimation procedure. The univariate models follow from the seminal work by Koenker and Basset (1978), and can be estimated using linear programming as originally suggested by Koenker and d’Orey (1987). The concept of multivariate quantiles is far from trivial due to the lack of inherent ordering in the multidimension. Chaudhuri (1996) proposes a way to extend the notion of univariate quantiles to the multivariate case, and Chakraborty (2003) generalizes this to the regression setting. Estimating multivariate quantile models has proven very complicated, primarily due to problems with lack of equivariance of the parameter estimates under general nonsingular transformations of the response vectors. To account for this problem, Chakraborty (2003) proposes a transformation-retransformation procedure based on so-called ‘data-driven coordinate systems’. I will use the approach suggested by Chakraborty (2003) to estimate the multivariate models.

2.1 Univariate quantile regression models

The usual starting point in the literature on return predictability is models of the form

$$y_{t+1} = \mu_t + \sigma_t \varepsilon_{t+1},$$

where y_{t+1} is a scalar containing returns either in levels or in logs, and μ_t and σ_t denote the conditional mean and volatility, respectively. ε_{t+1} is a return innovation with mean zero and variance one. Based on this framework a large literature has explored whether the conditional mean or volatility of returns vary over time. In the literature on predictability

of the mean return it is common to use the following setup

$$y_{t+1} = \beta \mathbf{x}_t + e_{t+1},$$

where it is assumed that $\mu_t = \beta \mathbf{x}_t$, and \mathbf{x}_t is a vector of dimension k , with k denoting the number of state variables used to predict future returns (usually including a constant). This implies that the mean return forecast can be written as

$$E(y_{t+1} | \mathcal{F}_t) = \beta \mathbf{x}_t,$$

where \mathcal{F}_t denotes the time t information set. The disadvantage of using this approach is that we only obtain an estimate of the center of the return distribution. Consider the case with only one state variable that has a positive slope coefficient. If the state variable increases at time t we expect the return to increase at time $t + 1$. However, the approach does not tell us if this increase is associated with an increase in risk or if the entire return distribution simply shifts to the right. In other words, this approach does not provide clear insights into the risk-return relation that is key in, for example, understanding pricing of financial assets and determining the mix of risky assets in portfolio decisions. This can partially be accommodated by also predicting future volatility, σ_t . However, the risk associated with, for example, skewness and kurtosis is still unaccounted for when restricting attention to only the first two moments of the distribution.

In this paper, I want to go beyond predictability of the mean and variance in order to obtain a much more detailed picture of return predictability. The idea is to model the conditional α -quantile, $q_\alpha(y_{t+1} | \mathcal{F}_t)$, of the return distribution through the use of state variables. This implies the following model

$$q_\alpha(y_{t+1} | \mathcal{F}_t) = \beta_\alpha \mathbf{x}_t,$$

where the local effect of \mathbf{x}_t on the α -quantile is assumed to be linear, but since the parameters are allowed to vary across quantiles, the model is very flexible.¹ By considering a large number of quantiles, we can trace out the entire future return distribution $P(y_{t+1} \leq \hat{q}_{\alpha,t}) = \alpha$, where $\hat{q}_{\alpha,t} = \hat{\beta}_\alpha \mathbf{x}_t$ is the time t conditional quantile. In this way, we can obtain a very detailed picture of time-variations in the return distribution. In particular, by using quantile regression it is possible to obtain valuable insights into whether return predictability tracks time-varying expected returns, time-varying risk, or a combination of both. Time-varying risk in this model is not restricted to volatility but can be any higher-order moment of the return distribution, including skewness and kurtosis. As an example, consider the case with only one state variable. If the quantile slope coefficients are symmetric around zero and increasing as a function of the quantile, then an increase (decrease) in the state variable will lead to an increase (decrease) in volatility. If instead the lower quantile slope coefficients are very negative while the median and the upper quantile slope coefficients are close to zero, then an increase (decrease) in the

¹Based on the work by Engle and Manganelli (2004), Cenesizoglu and Timmermann (2008) include last period's conditional quantile and the absolute value of last period's return as predictor variables. To make the univariate results directly comparable to the multivariate results, I only include exogenous state variables as predictor variables.

state variable will lead to a more (less) negatively skewed return distribution. The use of quantile regression can in this way provide a very detailed picture of return predictability in terms of the entire future return distribution. In Section 4.1, I will use the empirical results to explain the different shapes of the distribution in more detail.

To illustrate a few special cases consider the setup with only one predictor variable

$$q_\alpha(y_{t+1} \mid \mathcal{F}_t) = \beta_{0,\alpha} + \beta_{1,\alpha}x_t.$$

If $\beta_{1,\alpha} = 0$ for all values of α , we get the special case with time-invariant quantiles. This corresponds to the 'prevailing mean' model often claimed to outperform time-varying expected returns in out-of-sample forecasts (see e.g. Welch and Goyal, 2008). This special case will be named the 'prevailing quantile' (PQ) model in the following. Another special case is where $\beta_{1,\alpha}$ is constant across all values of α , which corresponds to the standard prediction model where the state variable x_t simply shifts the conditional mean of the return distribution.

Koenker and Basset (1978) provided the seminal work on quantile regression models. Following their work, regression quantiles are defined as

$$\hat{\beta}_\alpha = \arg \min_{\beta_\alpha \in \mathbb{R}^{1 \times k}} \sum_{t=1}^{T-1} L_\alpha(y_{t+1} - \beta_\alpha \mathbf{x}_t),$$

where T is the sample size and

$$L_\alpha(y_{t+1} - \beta_\alpha \mathbf{x}_t) = (\alpha - \mathbf{1}\{y_{t+1} - \beta_\alpha \mathbf{x}_t < 0\}) (y_{t+1} - \beta_\alpha \mathbf{x}_t), \quad (1)$$

with $\mathbf{1}\{\cdot\}$ denoting the indicator function. The general idea is to replace the conventional quadratic loss function used in ordinary least squares to obtain the conditional mean function by the check loss function (1), which instead allows estimation of the conditional quantile function. In the special case where $\alpha = 0.5$, the check loss function simplifies to the absolute loss function, which is the appropriate loss function to use in median regression. An alternative representation of the check loss function (1) is

$$L_\alpha(y_{t+1} - \beta_\alpha \mathbf{x}_t) = |y_{t+1} - \beta_\alpha \mathbf{x}_t| + (2\alpha - 1)(y_{t+1} - \beta_\alpha \mathbf{x}_t).$$

This representation of the loss function is convenient to bear in mind when turning to the multivariate case.

Consistency and asymptotic normality of regression quantiles is well established in the literature, see e.g. Koenker (2005) who give a thorough exposition of quantile regression. β_α will be estimated using linear programming as originally suggested by Koenker and d'Orey (1987).² Standard errors will be obtained through (x, y) -bootstrapping, which entails drawing (\mathbf{x}_t, y_{t+1}) pairs with replacement from the $T - 1$ pairs of the original sample, each with equal probability. This form of the bootstrap has been widely used

²More specifically, I apply a modified version of the Koenker and d'Orey (1987) version of the Barrodale and Roberts (1974) simplex algorithm as implemented in the statistical software program EViews.

in applications of quantile regression and in contrast to the alternative of bootstrapping residuals, it does not require identically distributed error terms to yield a valid estimate of the asymptotic covariance matrix. This makes (x, y) -bootstrapping preferable in quantile regression since accounting for heteroscedasticity is exactly one of the motivations behind applying this tool. Buchinsky (1995) conducts a comprehensive Monte Carlo study of several estimators of the asymptotic covariance matrix and finds that the (x, y) -bootstrapping procedure yields the best results. The size of the bootstrap samples will be set equal to the original sample size. Buchinsky (1995) finds that for some estimators a bootstrap sample size smaller than the original sample size can yield more accurate results, but the performance of the (x, y) -bootstrapping procedure is robust to this choice.

2.2 Multivariate quantile regression models

In the multivariate setup, the aim is to estimate

$$\mathbf{y}_{t+1} = \beta \mathbf{x}_t + \mathbf{e}_{t+1}, \quad (2)$$

where \mathbf{y}_{t+1} is now a vector of dimension d , with d denoting the number of dependent variables. This implies that β now has dimension $d \times k$. In this case the conditional quantile function is given as

$$\mathbf{q}_\alpha(\mathbf{y}_{t+1} \mid \mathcal{F}_t) = \beta_\alpha \mathbf{x}_t,$$

where α now is a vector of quantiles.

To illustrate the setting, take the case with two dependent variables: y_1 and y_2 . By estimating (2) using quantile regression, we can trace out the joint distribution, $P(y_{1,t+1} \leq \hat{q}_{\alpha_1,t} \text{ and } y_{2,t+1} \leq \hat{q}_{\alpha_2,t})$, which only in the special case of independence is equal to the product of the two marginal probabilities, $P(y_{1,t+1} \leq \hat{q}_{\alpha_1,t}) \times P(y_{2,t+1} \leq \hat{q}_{\alpha_2,t}) = \alpha_1 \times \alpha_2$, which we can obtain from the univariate models. As an example consider the case where y_1 and y_2 denote the return on stocks and bonds. Using this model we will be able to analyze if prespecified state variables can predict, say, the joint outcome of stock returns in the lower tail and bond returns in the upper tail. Furthermore, we can obtain an estimate of the entire future joint return distribution of stocks and bonds needed in, for example, portfolio decisions. Two things are worth noticing here. First, for given values of α the slope coefficients (and thereby the quantile forecasts) are not necessarily identical in the univariate and the multivariate setup. This is due to possible dependence between the dependent variables. Second, in the univariate setup, we know that the probability of falling below a given α -quantile is α . In the multivariate setup, we only know this probability from the outset in the special case of independence as mentioned above. In the general case, we can obtain a measure of the joint probability by calculating an in-sample coverage probability.

The multivariate approach can, in general, yield important insights into the dependence structure of two or more variables. Often simple linear correlation is used as a measure of dependence, and given the variables are multivariate normally distributed,

this measure also completely characterizes the dependence structure. However, in the more general case where the distribution, for example, is affected by asymmetries, this measure no longer gives an adequate description of the dependence structure. The need to model the dependence between two or more variables also in the presence of asymmetries has spurred the use of, for example, copulas (see e.g. Patton, 2004) and regime-switching models using mixtures of multivariate normal distributions (see e.g. Ang and Bekaert, 2002, and Guidolin and Timmermann, 2007). The multivariate quantile approach presented in this paper complements the literature on multivariate modeling by proposing a non-parametric alternative to the existing methods. Furthermore, given simple linear correlation actually does provide an adequate measure of the dependence structure, we still only know the 'average' dependence. The measure provides no insight into which outcomes generate the correlation. Again, take the example with stock and bond returns and assume these are bivariate normally distributed implying that the linear correlation gives a complete description of the dependence structure. We know stocks and bonds are correlated, but which return outcomes generate this correlation? Is it, for example, stock returns in the lower tail of the return distribution jointly with bond returns in the upper tail? Or stock and bond returns jointly in the lower or upper tail? By choosing a sufficiently fine grid of quantiles, we can trace out the joint distribution and compare in-sample coverage probabilities to the probabilities given independence (the product of the quantiles), and thereby obtain insights into which return outcomes generate the correlation between stocks and bonds.

In contrast to the univariate case there is no inherent ordering in the multidimension, which presents a big challenge in performing multivariate quantile regression. Chaudhuri (1996) proposes a way to extend the notion of univariate quantiles to the multivariate case. The idea is to index multivariate geometric quantiles, based on Euclidean distances, using the elements of the d -dimension open unit ball. The corresponding quantiles not only give the idea of 'extreme' or 'central' observations but also about their orientation in a multivariate data cloud. Chaudhuri (1996) establishes existence and uniqueness of multivariate geometric quantiles, and furthermore he proves consistency and asymptotic normality.

Chakraborty (2003) generalizes this idea and defines regression quantiles in the multivariate linear model (2) as

$$\widehat{\beta}_\alpha = \arg \min_{\beta_\alpha \in \mathbb{R}^{d \times k}} \sum_{t=1}^{T-1} L_\alpha(\mathbf{y}_{t+1} - \beta_\alpha \mathbf{x}_t),$$

where the loss function is given as

$$L_\alpha(\mathbf{y}_{t+1} - \beta_\alpha \mathbf{x}_t) = \|\mathbf{y}_{t+1} - \beta_\alpha \mathbf{x}_t\| + \mathbf{u}'(\mathbf{y}_{t+1} - \beta_\alpha \mathbf{x}_t). \quad (3)$$

$\|\cdot\|$ denotes the Euclidean norm and the index vector \mathbf{u} is an element of the open unit ball $B^{(d)} = \{\mathbf{u} \mid \mathbf{u} \in \mathbb{R}^d, \|\mathbf{u}\| < 1\}$. Note here that a vector \mathbf{u} for which $\|\mathbf{u}\|$ is close to one corresponds to an 'extreme' quantile, while a vector \mathbf{u} for which $\|\mathbf{u}\|$ is close to zero corresponds to a 'central' quantile. Often it is of interest to compare the multivariate estimates to the corresponding univariate estimates. Furthermore, it is often more intuitive

to choose the quantiles for which to estimate the model (i.e. α) and then map these into the open unit ball (i.e. \mathbf{u}) than to choose the index vector \mathbf{u} directly. To accommodate this, note that we can obtain a 1-1 mapping from the open square $(0, 1)^d$ in which α is defined to the d -dimensional open unit ball as

$$\mathbf{u} = \frac{\|\mathbf{g}\|_\infty}{\|\mathbf{g}\|} \times \mathbf{g}, \quad (4)$$

where $\|\mathbf{g}\|_\infty = \max\{|g_1|, \dots, |g_d|\}$ and $\mathbf{g} = 2\alpha - \iota$, with ι denoting a d -dimensional vector of ones. Hence, after choosing the quantiles, α , for which we want to estimate the multivariate model, we can map these into the open unit ball using (4), and then proceed with the estimation from there.

However, as Chakraborty (2003) notes the geometric regression quantiles obtained by minimizing (3) are not equivariant under arbitrary nonsingular transformations of the response vectors and they are not even equivariant under coordinatewise scale transformations. This implies that regression quantiles obtained by minimizing (3) directly are very dependent on the choice of coordinate system, which is not desirable. Chakraborty (2003) introduces a transformation-retransformation procedure to resolve the problem of lack of equivariance. The fundamental idea in this procedure is, first, to form an appropriate 'data-driven coordinate system' (i.e. transform the response vectors), and then to formulate the model in terms of that coordinate system.³ Next, the idea is to estimate the model based on the transformed response vectors before, finally, retransforming the parameter estimates so as to express everything in terms of the original coordinate system. Chakraborty (2003) shows that the estimates following this procedure are equivariant, and he also proves existence and uniqueness as well as consistency and asymptotic normality. The Appendix contains the algorithm used in the present paper to estimate the multivariate quantile model, including how to choose the 'data-driven coordinate system'; for additional details, see Chakraborty (2003). As in the univariate case, standard errors will be obtained through (x, y) -bootstrapping.

Recently, a series of papers show that several state variables have a non-linear effect on the conditional mean of future stock and bond returns (see e.g. Rapach and Wohar, 2005, McMillan and Wohar, 2009, and Guidolin et al., 2009). Non-linearity can also be present in conditional quantile models and does not necessarily imply added difficulty in the estimation procedure of univariate models compared to the linear case; see e.g. Koenker and Park (1996) and De Gooijer and Zerom (2003). However, I restrict my attention to linear models in order to have a common setup in the univariate and multivariate case. To my knowledge the problem of estimating non-linear multivariate quantile models has not yet been addressed.

³Chaudhuri and Sengupta (1993) introduced the idea of 'data-driven coordinate systems'. Chakraborty and Chaudhuri (1996) used this idea and introduced a transformation-retransformation procedure to construct the multivariate median, while Chakraborty (1999) considered the regression analog.

3 Data

I examine predictability of stock and bond return distributions using a set of nine state variables, all of which have previously been used in the literature on predictability of the conditional mean and volatility. I use US monthly data from 1941:5 to 2008:12, where the starting date is dictated by data availability. Stock returns are obtained as the return (including dividends) on the S&P500 Index and bond returns are measured as the return on 5-year Treasury bonds. The 30-day T-bill rate is subtracted from stock and bond returns to obtain excess returns. The data on stocks, bonds, and bills is obtained from the Center for Research in Security Prices (CRSP).

The set of state variables contains two valuation ratios. The dividend-price ratio (DP) is calculated as the 12-month moving sum of dividends paid on the S&P Index divided by ultimo price. Likewise, the earnings-price ratio (EP) is calculated as the 12-month moving sum of earnings on the S&P Index divided by ultimo price. The data used to construct these valuation ratios is obtained from Robert Shiller's website. Furthermore, the set of predictor variables contains two corporate finance variables. The dividend payout ratio (DE) is the ratio of dividends to earnings, while net equity expansion (NTIS) is the ratio of 12-month moving sums of net issues by NYSE listed stocks divided by the total end-of-year market capitalization of NYSE stocks. Stock variance (SVAR) computed as the sum of squared daily returns on the S&P500 Index is also included as a state variable. NTIS and SVAR are obtained from Ivo Welch's website. The set of state variables also contains three bond yield measures. The Treasury-bill rate (TBL) is the *3-month Treasury Bill: Secondary Market Rate* from the economic research data base at the Federal Reserve Bank at St. Louis (FRED). The term spread (TMS) is the difference between the long-term government bond yield from Ibbotson's *Stocks, Bonds, Bills and Inflation Yearbook* and TBL. The default spread (DFS) is the difference between yields on BAA and AAA-rated corporate bonds obtained from FRED. Finally, inflation (INFL) is the return on the Consumer Price Index (All Urban Consumers) from the Bureau of Labor Statistics. This state variable is a common broad macroeconomic indicator.

Many papers have used a similar set of predictor variables to examine stock return predictability. For example, Welch and Goyal (2008) provide a comprehensive study of these variables' ability to predict the conditional mean stock return (both in- and out-of-sample).⁴ Other studies use similar state variables to predict both the conditional mean and volatility; see e.g. Marquering and Verbeek (2004) and Lettau and Ludvigson (2010). Cenesizoglu and Timmermann (2008) also use a broad set of state variables to examine predictability of the entire stock return distribution. The literature on bond return predictability is much less voluminous. Examples of papers examining predictability of the conditional mean bond return include Ilmanen (1995) and Kirby (1997). In a recent paper, Viceira (2010) examines predictability of both the conditional mean and volatility using the short rate and the yield spread as predictor variables. Finally, using vector autoregressions a number of papers have examined predictability of the conditional mean

⁴See Welch and Goyal (2008) for a non-exhaustive list of papers examining predictability of the conditional mean stock return using the individual state variables.

of both stock and bond returns; see e.g. Campbell et al. (2003), Engsted and Pedersen (2010), and Viceira (2010).

4 Univariate return distributions

In this section, I report results on predictability of stock and bond returns in a univariate setup. Section 4.1 contains in-sample results using both quantile regression and least squares estimated over the entire sample period. The results obtained using quantile regression yield insights into predictability of different parts of the return distributions, while the results obtained using least squares only contain information on predictability of the conditional mean return. Section 4.2 contains an out-of-sample analysis of forecasts based on the quantile regression framework.

4.1 In-sample return predictability

Table 1 shows the slope coefficients when excess stock returns are regressed on each of the nine state variables using either least squares (Panel A) or quantile regression (Panel B). To obtain a sufficiently detailed picture of the return distribution, I perform quantile regression for $\alpha = \{0.05, 0.10, 0.20, \dots, 0.90, 0.95\}$. t -statistics are shown in parentheses. For least squares the t -statistics are based on White (1980) heteroscedasticity-consistent standard errors, and for quantile regression they are based on (x, y) -bootstrapped standard errors as outlined in Section 2.1. For ease of readability, slope coefficients with an associated t -statistic larger than 1.96 or smaller than -1.96 are boldfaced. For least squares a slope coefficient significantly different from zero means that the state variable in question significantly predicts an increase or decrease in next period's *mean* excess return. For quantile regression a significant slope coefficient for a given α -quantile means that the state variable significantly predicts an increase or decrease in next period's α -*quantile* of the return distribution. If the α -quantile slope coefficients are significantly different it implies time-variation in higher-order moments of the return distribution. This can be tested using the slope equality test proposed by Koenker and Basset (1982). Panel C shows the Wald statistics from this test with p-values in brackets.

Panel A shows that five of the nine state variables predict the mean excess stock return. The slope coefficient on the dividend-price ratio and earnings-price ratio is significantly positive, while it is significantly negative for stock variance, T-bill rate, and inflation. These results are consistent with evidence from the existing literature and economic theory. Turning to the results based on quantile regression, Panel B shows that eight of the nine state variables predict at least some part of the return distribution. For example, the dividend-price ratio predicts the center and upper shoulder of the return distribution, while inflation predicts the center and lower shoulder. For visual evaluation the slope coefficients and associated 95% confidence bands using both least squares and quantile regression are displayed in Figure 1. Ignoring statistical significance

for a while and just focusing on the value of the slope coefficients, we observe that the individual state variables predict the future return distribution in quite different ways. For example, the slope coefficients for the dividend-price ratio are all positive and of similar magnitude. This implies that an increase in the dividend-price ratio leads to an increase of roughly same magnitude in all return quantiles, i.e. an upward location-shift of the return distribution with no changes in volatility, skewness or any other higher-order moment. In contrast, the slope coefficients on stock variance are mainly negative except for the highest quantiles. Furthermore, the median slope coefficient is of similar magnitude as the least squares slope coefficient indicating symmetry in the future return distribution. In this case, an increase in the state variable leads to a downward shift in the return distribution *and* an increase in volatility since the upper quantiles will increase and the lower quantiles decrease. Testing for slope equality across the quantiles (Panel C) reveals that only stock variance captures time-variation in higher-order moments of the stock return distribution. The remaining eight state variables either do not capture time-variation in any moment of the return distribution or only of the first moment.

Table 2 and Figure 2 show the corresponding results for excess bond returns, which are quite different from those obtained for stock returns.⁵ While only four of the nine state variables predict mean bond excess return, all state variables predict at least some part of the return distribution. Furthermore, the slope equality test shows that seven of the nine state variables significantly capture time-variation in higher-order moments of the bond return distribution. The state variables predict the future bond return distribution in quite different ways. For example, the T-bill rate significantly predicts the entire distribution except the center. The slope coefficients for the upper part of the distribution are positive, while they are negative for the lower part. Furthermore, the slope coefficient for the median is close to zero and roughly equal to the least squares slope coefficient indicating symmetry in the return distribution. Comparing slope coefficients for the $1 - \alpha$ and α -quantile, we also find evidence of symmetry in the distribution. These results imply that an increase in the T-bill rate does not shift the excess bond return distribution neither upwards nor downwards, but it increases the distribution's dispersion, i.e. volatility increases. A similar result is observed for the default spread. However, in this case the increase in dispersion is accompanied by an upward location-shift as seen from the slope coefficient for both the median and the mean. Another example is inflation. This state variable captures time-variation in the third moment. The slope coefficients are negative for all quantiles as well as the mean indicating a downward shift in the return distribution when inflation increases. However, only for the median and the lower quantiles are the slope coefficients significantly negative. Furthermore, the slope coefficients for the lower quantiles are much larger in absolute value than the slope coefficients for the upper quantiles. This implies that an increase in inflation leads to a more negatively skewed bond return distribution since the lower quantiles will decrease a lot while the upper quantiles will remain roughly the same (or decrease slightly). As a final example, consider the term spread. This state variable is generally considered to

⁵From Figure 1 and 2 it is clear that, in general, the quantile slope coefficients are estimated with high precision. In some cases (especially for bonds) also even more so than the least squares slope coefficients. Only in the extreme tails do we for some state variables observe a noticeable decrease in precision.

be the 'best' predictor of future bond returns. From Panel A it is clear that the least squares slope coefficient is positive, which is in line with theory and existing evidence, but it is not statistically significant. In contrast, the slope coefficients for the center and upper shoulder are all highly significant. Hence, the conditional median is predictable, but the conditional mean is not. Here it is important to note that in contrast to quantile regression, least squares is highly sensitive to outliers, which can lead to unpredictable conditional means although the conditional median is in fact predictable.

In general, the results presented in this section show that lack of predictability of the mean return does not necessarily imply that other parts of the return distribution are unpredictable. Predictability can be present in the first, second, third, or any other higher-order moment of the return distribution. The results also illustrate how the different state variables predict future asset returns in quite different ways in terms of location-shifts, volatility and skewness etc. These findings complement our existing knowledge of return predictability in the traditional sense, i.e. of the conditional mean. For example, an expected increase in future stock returns due to an increase in the dividend-price ratio is not associated with changes in risk, while an expected decrease in future bond returns due to an increase in inflation is associated with an increase in downside risk. These results yield important insights into the risk-return relation that is key in, for example, understanding the pricing of financial assets.

The risk-return trade-off for a given asset is usually measured by its conditional Sharpe ratio

$$SR_t = \frac{E_t R_{t+1}}{E_t V_{t+1}},$$

where $E_t R_{t+1}$ is the mean excess return from time t to $t + 1$ conditional on information available at time t and $E_t V_{t+1}$ is a measure of the standard deviation of excess return, again conditional on time t information. The Sharpe ratio is an intuitively appealing characterization of the price of risk, since it measures how much return an investor can get per unit of asset volatility. In classic asset pricing models such as the capital asset pricing model by Sharpe (1964), the expected risk premium varies proportionally with expected volatility, which implies that the Sharpe ratio should be constant over time. The expected risk premium is allowed to vary over time, but it must be perfectly positively correlated with expected volatility, i.e. the amount of risk must also be time-varying. More recent asset pricing models such as the habit model by Campbell and Cochrane (1999) allow the price of risk to vary over time, which implies that the expected risk premium and volatility need not be perfectly positively correlated.

During the last 20 years numerous papers have established that the expected risk premium on stocks varies over time, but researchers are still debating whether this is due to time-varying price of risk or time-varying amount of risk. Based on the dividend-price ratio, earnings-price ratio, T-bill rate, and inflation, the results in Table 1 provide evidence of time-varying price of stock market risk, since the conditional mean excess return based on these state variables is time-varying (Panel A), but the distribution including volatility (amount of risk) does not change over time (Panel C). For stock variance both the conditional mean excess return and volatility vary over time, but the

price of stock market risk is still time-varying. If stock variance increases, the numerator in the Sharpe ratio decreases while the denominator increases, and vice versa. Hence, the expected risk premium and volatility are negatively correlated, resulting in a time-varying Sharpe ratio.

Table 2 reveals evidence of time-varying price of risk also for bonds, but in a different way than for stocks. For example, based on the dividend-price ratio, earnings-price ratio, payout ratio, and T-bill rate the amount of risk varies over time (Panel C), but the expected risk premium is constant (Panel A). Hence, the denominator in the Sharpe ratio varies over time, but the numerator is time-invariant. Based on net equity expansion, stock variance, and default spread both the expected risk premium and the amount of risk vary over time, and in contrast to the case with stocks and stock variance as state variable, they are positively correlated. However, it is not directly evident from Table 2 if they are perfectly correlated implying constant bond market price of risk.

Risk is not confined to be measured by volatility. Similar to the Sharpe ratio, we can construct a characterization of the price of risk, which measures how much return an investor can get per unit of, say, asset skewness or kurtosis. The use of quantile regression allows for a robust and straightforward way to calculate the third and fourth moment of the return distribution (see e.g. Cenesizoglu and Timmermann, 2008), which in turn can provide more detailed insights into the risk-return relation than captured by the Sharpe ratio. An in-depth analysis of the risk-return relation is outside the scope of this paper, and hence I leave it for future research.

4.2 Out-of-sample forecasts

A key issue in the return predictability literature is out-of-sample evaluation. It is very likely that empirical forecasting models are, to some extent, almost always misspecified in the sense that the functional form of the model may be incorrect and/or that all the relevant state variables are not included compared to the true data-generating process. However, it is still relevant to examine the relative accuracy of competing forecasts. In the present context the relevant question is whether a given state variable is better at forecasting a given return quantile than the other state variables. For each quantile, I perform pairwise comparisons based on the check loss function (1) and the testing procedure proposed by Diebold and Mariano (1995). The test is very simple and entails calculating the loss differential

$$d_{t+1,\alpha} = (\alpha - \mathbf{1} \{y_{t+1} - \hat{q}_{\alpha,t}^i < 0\}) (y_{t+1} - \hat{q}_{\alpha,t}^i) - (\alpha - \mathbf{1} \{y_{t+1} - \hat{q}_{\alpha,t}^j < 0\}) (y_{t+1} - \hat{q}_{\alpha,t}^j),$$

at each point in time in the out-of-sample evaluation period. $\hat{q}_{\alpha,t}^i$ and $\hat{q}_{\alpha,t}^j$ are quantile forecasts for state variable i and j , respectively, based on information only up to time t . The null of equal forecast accuracy can then be tested as

$$t_\alpha = \frac{\bar{d}_\alpha}{\sqrt{\text{var}(\bar{d}_\alpha)}} \sim \mathcal{N}(0, 1), \quad (5)$$

where \bar{d}_α is the average of $d_{t+1,\alpha}$ over the out-of-sample period. I apply a recursive scheme in generating out-of-sample forecasts. Initially, I estimate the models based on data from the start of the sample and up to 1969:12. I then generate forecasts for return quantiles in 1970:01 and compare these to realized returns in the same period. Next, I add the data from 1970:01, reestimate the models, and generate forecasts for return quantiles in 1970:02, which I then compare to realized returns in that period. This recursive scheme is repeated up to the end of the sample period resulting in a total of 468 out-of-sample forecasts for the period 1970:01 to 2008:12.

Table 3 shows t -statistics from the pairwise loss differential tests for the lower tail (Panel A), center (Panel B), and upper tail (Panel C) of the return distribution. In each panel the results based on stock returns are shown below the main diagonal while the results based on bond returns are shown above the main diagonal. The loss differential is calculated as the loss from the variable on the vertical axis minus the loss from the variable on the horizontal axis. Hence, a positive (negative) entry in the table implies that there is higher (lower) loss associated with the variable on the vertical axis than the variable on the horizontal axis, i.e. the variable on the horizontal (vertical) axis yields a relatively better forecast for that particular quantile. Again, for ease of readability t -statistics larger than 1.96 or smaller than -1.96 are boldfaced. In addition to the nine state variables, Table 3 also includes results based on the prevailing quantile (PQ) model, i.e. a model without state variables. This is the quantile equivalent to the 'prevailing mean' model often claimed to outperform time-varying expected returns in out-of-sample forecasts.

Regarding stock returns, the forecast accuracy is only significantly different in three pairwise comparisons. The default spread forecasts the lower tail significantly better than the payout ratio, and inflation forecasts the center significantly better than both stock variance and the term spread. Ignoring statistical significance for a while and just focusing on the sign of the loss differential, a number of interesting results appear. For example, in the lower tail the average loss associated with stock variance is lower than the average loss associated with the other eight state variables and PQ. In contrast, in the center and the upper tail this is the case for inflation and the default spread, respectively. Hence, the relative predictive accuracy of the individual state variables differs across the return distribution. State variables that perform relatively well in the lower tail might not perform equally well in the upper tail and vice versa. Likewise for the state variables that perform relatively worse than the other state variables. In the center the term spread is associated with higher average loss than the other eight state variables and PQ, and in the upper tail this is the case for the T-bill rate.

In the literature on stock return predictability, it is often argued that state variables do not predict the future return better than the prevailing mean (see e.g. Welch and Goyal, 2008). The results in Table 3 support a similar argument concerning the entire return distribution. None of the state variables predict the lower, center, or upper tail of the stock return distribution significantly better than the prevailing quantile. In fact, with exception of the dividend-price ratio, stock variance, and default spread in the lower tail, inflation in the center, and default spread in the upper tail, PQ is always associated

with the relatively lowest loss.

Turning to bond returns, the loss differential is in many cases statistically significant. For example, in both the lower and upper tail, the default spread is a significantly better forecaster than all the other state variables (except stock variance in the upper tail and the T-bill rate in both tails). Inflation also predicts the lower tail significantly better than the majority of other state variables. If we just focus on the sign of the loss differential we see that the T-bill rate is a relatively good forecaster of both the lower and upper tail, while it performs relatively worse than all the other state variables in predicting the center of the distribution. In contrast, the term spread is a relatively good forecaster of the center, but it performs worse than all the other state variables in the tails of the distribution. The term spread is often considered to be the 'best' forecaster of future bond returns. The results in Table 3 show that the relatively good predictive ability of this state variable is confined to the center of the distribution, and completely vanishes in the tails.

5 Multivariate return distributions

Often we need an estimate of a future *joint* return distribution. For example, in portfolio choice problems, where the investor is allowed to invest in more than one risky asset, or in various risk management problems. In general, this calls for a multivariate model. Only in the special case of independence between the relevant assets, can we rely on a combination of univariate models. In this section, I use the multivariate quantile model presented in Section 2.2 to analyze the joint distribution between stock and bond returns. In Section 5.1, I use the multivariate quantile model to examine if stocks and bonds are independent, and if not, in which parts of the distribution the dependence is present. If stocks and bonds are independent, the joint quantile forecasts are simply equal to the quantile forecasts from the univariate models. Next, Section 5.2 examines the joint in-sample predictability of stock and bond returns. Finally, Section 5.3 contains an out-of-sample analysis. I calculate the out-of-sample loss differential between quantile forecasts from the univariate and joint models. This can be used to evaluate if it matters to take the dependence structure between stocks and bonds into account in forecasting. I also perform pairwise out-of-sample loss differential tests similar to those shown in Table 3. Similar to the univariate case, these tests can help identify if any of the state variables are better than the others in predicting various parts of the *joint* stock and bond return distribution out-of-sample.

5.1 Are stocks and bonds independent?

Dependence can be evaluated by calculating in-sample coverage probabilities. These probabilities are calculated by, first, estimating the quantile model over the entire sample period. Next, based on the estimated coefficients, in-sample quantile forecasts are

constructed. Finally, in-sample coverage probabilities in the univariate and bivariate case, respectively, are calculated as

$$CP_\alpha = \frac{1}{T-1} \sum_{t=1}^{T-1} \mathbf{1}(y_{t+1} \leq \hat{q}_{\alpha,t}),$$

$$CP_{\alpha_1, \alpha_2} = \frac{1}{T-1} \sum_{t=1}^{T-1} \mathbf{1}(y_{1,t+1} \leq \hat{q}_{\alpha_1,t} \text{ and } y_{2,t+1} \leq \hat{q}_{\alpha_2,t}).$$

In the univariate case, we know that $CP_\alpha = \alpha$ by construction. In the bivariate case, however, we only know the in-sample coverage probability from the outset if the variables are independent, in which case $CP_{\alpha_1, \alpha_2} = \alpha_1 \times \alpha_2$. Hence, if the joint coverage probability is different from the product of the two quantiles, we can conclude that the variables are dependent. Furthermore, by choosing a sufficiently fine grid of quantile combinations, we can determine in which part of the joint distribution the dependence occurs.

Table 4 shows in-sample coverage probabilities for 11 different quantile combinations (α_s, α_b) with α_s denoting the quantile for stocks and α_b the quantile for bonds. The quantile combinations (0.05,0.05), (0.50,0.50), and (0.95,0.95) denote the joint lower tail, center, and upper tail of the stock-bond return distribution, while (0.05,0.95) denote the joint outcome of stock returns in the lower tail and bond returns in the upper tail, and vice versa for (0.95,0.05). The other quantile combinations can be interpreted similarly. By choosing a sufficiently fine grid of quantile combinations, we can trace out the entire joint distribution. In the final column the coverage probabilities given independence are shown. The first thing worth noticing in Table 4 is that, as expected, the results are quite robust across state variables and the prevailing quantile model. Hence, the overall conclusions about the dependence between stocks and bonds do not depend on specific state variables or the assumption about time-varying return distributions.

From Table 4 it is clear that stocks and bonds are not independent. The in-sample coverage probabilities are generally quite different from the coverage probabilities given independence. From the first five quantile combinations where $\alpha_s = \alpha_b$ we see that the joint stock-bond distribution contains more probability mass in the joint tails compared to the case with independence. For example, roughly 20% of the observations lie above (0.95,0.95). If stocks and bonds are independent, this number should only be 10%. Most noticeable is, however, the coverage probability associated with the quantile combination (0.05,0.95). Given independence only 5% of the observations should be in this part of the distribution, but the coverage probability is roughly 20%. Moreover, since the coverage probability associated with the quantile combination (0.05, 0.50) does not reveal a similar overrepresentation (compared to the case with independence), we can conclude that part of the dependence is due to stock returns falling in the lower tail jointly with bond returns in the upper shoulder. In other words, the joint outcome of very low stock returns and high bond returns appears more frequently than if stocks and bonds were independent. These observations are also clearly visible in Figure 3, where each quadrant shows the coverage probability for the relevant quantile combinations based on the prevailing quantile model (with probabilities given independence in parentheses).

Based on the entire sample period, the correlation between stocks and bonds is roughly 0.10. This is an 'average' measure of dependence and it does not provide insights into which return outcomes are responsible for the comovement. This insight can, however, be obtained based on coverage probabilities as in Figure 3. Note that outcomes along the main upward sloping diagonal generate positive correlation, while outcomes along the main downward sloping diagonal generate negative correlation. A growing literature documents time-varying correlation between stocks and bonds (see e.g. Ilmanen, 2003, Connolly et al., 2005, Christiansen and Rinaldo, 2007, and Baele et al., 2010), with correlation being both positive and negative. According to Figure 3 the positive correlation is mainly due to the joint outcome of very high stock returns and moderately high bond returns. The joint outcome of low stock and bond returns do not appear to contribute to the positive correlation. Figure 3 also shows that the negative correlation is primarily driven by the joint outcome of very low stock returns and moderately high bond returns. Although this helps us understand which return outcomes generate the correlation between stocks and bonds, it is also clear from Figure 3 that the stock-bond distribution is affected by asymmetries, which basically implies that simple linear correlation does not give an adequate measure of dependence between stocks and bonds. This is consistent with Guidolin and Timmermann (2006), who based on a regime switching model find evidence of nonlinear dynamics in the joint distribution of stock and bond returns.

By choosing a more fine grid of quantile combinations, it is possible to obtain a more detailed picture of the dependence structure between stocks and bonds. For the purpose of this paper, however, the results in Table 4 and Figure 3 are sufficient to conclude that stocks and bonds are not independent, and hence, we need to resort to the multivariate setup when an estimate of the future joint return distribution is required.

5.2 Joint in-sample return predictability

Table 5 shows the slope coefficients and associated t -statistics estimated in the multivariate setting for stock (Panel A) and bond (Panel B) returns for five different quantile combinations. As in Tables 1 and 2, slope coefficients with an associated t -statistic larger than 1.96 or smaller than -1.96 are boldfaced. From Table 5 it is clear that at least part of the joint stock-bond distribution is predictable as a function of economic state variables. Consider, for example, the default spread. This state variable significantly predicts both the joint lower and upper tail of the stock-bond distribution, as well as the joint outcome of stock returns in the lower tail and bond returns in the upper tail, and vice versa. The implication of the results in Table 5 is that the joint distribution between stocks and bonds is time-varying. Furthermore, based on the size and sign of the coefficients it is clear that this time-variation is not restricted to only location-shifts, but also concerns the shape of the distribution.

According to the in-sample coverage probabilities in Table 4 and Figure 3, stocks and bonds are dependent. Another way to evaluate dependence is to compare the jointly estimated slope coefficients with those estimated in the univariate setup. Consider, for example, stock variance in Table 5, Panel A. For $\alpha_s = 0.95$, the slope coefficient for

stock returns is equal to 0.711 and 1.132 for $\alpha_b = 0.05$ and 0.95, respectively. If stocks and bonds are independent, the slope coefficient for stock returns for a given α_s -quantile should be constant across the α_b -quantile and equal to the slope coefficient in the univariate setup. This coefficient is estimated to be 3.336, and thus these results illustrate that we need to account for the dependence structure between stocks and bonds when we model their joint distribution. Figure 4 illustrates this point more clearly. The solid line gives the univariately estimated slope coefficient for the 0.95-quantile for stock returns as a function of stock variance, while the line with circles gives the corresponding estimates in the multivariate setup for different values of the α_b -quantile. Figure 4 shows that for $\alpha_b = 0.50$ and 0.60, the slope coefficient in the multivariate setting is higher than the slope coefficient in the univariate setting, while it is lower for the remaining α_b -quantiles.

These results indicate that the conditional joint distribution between stocks and bonds cannot be captured by simply combining the univariately estimated distributions. Whenever an estimate of the joint distribution between stocks and bonds is required, we thus need to resort to the multivariate model.

5.3 Out-of-sample joint forecasts

The fact that stocks and bonds are dependent does not necessarily imply that multivariate quantile forecasts perform better out-of-sample than a combination of univariately estimated quantile forecasts. To evaluate if it matters to take the dependence into account in out-of-sample forecasting, I calculate out-of-sample loss based on (3) and the parameter estimates from both the univariate and multivariate setup using the recursive scheme outlined in Section 4.2. Based on these loss series, I calculate the differential test (5). The results are shown in Table 6. Again, for ease of readability t -statistics larger than 1.96 or smaller than -1.96 are boldfaced. A positive (negative) entry implies that the out-of-sample loss associated with forecasts from the univariate setup is higher (lower) than the loss associated with forecasts from the multivariate setup. From Table 6 it is clear that in some cases, a combination of quantile forecasts constructed based on the univariate setup performs best out-of-sample, although the difference is not significant. However, the vast majority of test statistics are positive and many of these are also statistically significant. Hence, the conclusion obtained from Table 6 is that it clearly matters to take the dependence between stocks and bonds into account, also in out-of-sample forecasts.

The next natural question is: Do any of the state variables perform better in out-of-sample joint forecasting than others? To answer that question, I perform pairwise loss differential tests similar to those in Table 3, but now based on the multivariate loss function (3). Results for five different quantile combinations are shown in Table 7. In Panel A, the results for the joint lower (upper) tail are shown below (above) the main diagonal, while Panel B shows the results for the joint center of the distribution. Finally, below (above) the main diagonal Panel C shows the joint outcome of stock returns in the lower (upper) tail and bond returns in the upper (lower) tail. From Table 7 it is clear that many of the loss differentials are statistically significant in the 'tails' of the

distribution, but not in the center.

In terms of the state variables' relative out-of-sample performance, Table 7 reveals clear differences between the different parts of the joint distribution. In the joint lower tail (0.05,0.05), the T-bill rate performs significantly better than all the other state variables, while the term spread performs worse. In the joint upper tail (0.95,0.95), the relatively best forecaster is the default spread, while the worst is the dividend-price ratio. In both cases the differences are statistically significant. Turning to Panel C, we see that for (0.05,0.95) the results are very similar to those for (0.05,0.05). Again, the T-bill rate is the relatively best forecaster, while the term spread is the worst. Hence, these results seem robust for stock returns in the lower tail, irrespective of bond returns falling in the lower or upper tail. Likewise, for (0.95,0.05) we find that the default spread is the relatively best forecaster similar to the case with (0.95,0.95). However, for (0.95,0.05) the T-bill rate is now the relatively worst forecaster. Regarding the center of the distribution, inflation is the relatively best forecaster, while the term spread is the worst.

The results in Table 7 are in general consistent with those from the univariate setting as shown in Table 3. The relative importance of stocks and bonds differs, however, in the different parts of the joint distribution. The relative good and bad performance in the center of the joint distribution by inflation and the term spread, respectively, can be ascribed to stocks. Likewise, the relative bad performance by the T-bill rate in (0.95,0.05) seems to be due to its poor ability to forecast the upper tail of the stock return distribution. In contrast, the relatively good and bad performance of the T-bill rate and the term spread, respectively, in (0.05,0.05) and (0.05,0.95) can be ascribed to bonds.

6 A portfolio study

To examine if it bears any economic significance to take the empirical return distribution into account, this section contains a portfolio study. In general, investors require an estimate of the entire distribution of future returns to make their portfolio decisions. Only with certain utility and/or distributional assumptions, knowledge of the first two moments of the return distribution is sufficient (see e.g. Campbell and Viceira, 2002). I perform an out-of-sample portfolio study based on an investor with power utility defined over next period's wealth and under unknown return distributions estimated using quantile regression. I compare these results to the special case where returns are assumed to be lognormally distributed. In this case, we obtain a simple closed-form solution to the optimal portfolio choice, which is consistent with mean-variance analysis, and hence only requires knowledge of the first two moments of the return distribution.

6.1 Optimal portfolio choice

I consider both a univariate setting, where the investor only has access to one risky asset (either stocks or bonds) and a risk free asset, and a multivariate setting, where the investor has access to two risky assets (both stocks and bonds) as well as the risk free asset. First, I describe the setting in the univariate case. In this case the budget constraint is given as

$$\begin{aligned} W_{t+1} &= W_t \left(1 + w_t \tilde{R}_{t+1} + (1 - w_t) R_{f,t} \right) \\ &= 1 + w_t R_{t+1} + R_{f,t}, \end{aligned}$$

where w_t denotes the fraction of wealth invested in stocks and bonds, respectively, at time t . The remainder $1 - w_t$ is invested in the risk free asset. \tilde{R}_{t+1} denotes the return on the risky asset at time $t + 1$, $R_{f,t}$ denotes the return on the risk free asset at time $t + 1$ but known at time t , and R_{t+1} denotes the excess return on the risky asset at time $t + 1$ calculated as $R_{t+1} = \tilde{R}_{t+1} - R_{f,t}$. Without loss of generality, W_t is set equal to 1.

The investor is assumed to have standard power utility (CRRA) defined over next period's wealth

$$U(W_{t+1}) = \frac{W_{t+1}^{1-\gamma}}{1-\gamma},$$

where γ denotes the relative risk aversion parameter. To obtain the optimal portfolio choice, the investor solves the following optimization problem

$$w_t^* = \arg \max_{w_t} \int \frac{(1 + w_t R_{t+1} + R_{f,t})^{1-\gamma}}{1-\gamma} f(R_{t+1} | \mathcal{F}_t) dR_{t+1}, \quad (6)$$

where $f(R_{t+1} | \mathcal{F}_t)$ is the conditional probability distribution of future excess returns based on the information set at time t . Hence, in the general case, the investor needs an estimate of the entire future return distribution. To solve for the optimal portfolio choice, we can either impose assumptions about $f(R_{t+1} | \mathcal{F}_t)$ or discretize the integral in (6). To make the approach simple and consistent across the univariate and multivariate setup, I choose to discretize the integral in the following way

$$\begin{aligned} w_t^* &= \arg \max_{w_t} \frac{(1 + w_t \hat{q}_{0.05,t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \hat{p}_{(0,0.05)} \\ &\quad + \sum_{\alpha=0.1}^{0.9} \frac{(1 + w_t \hat{q}_{\alpha,t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \hat{p}_{(\alpha-0.05, \alpha+0.05)} \\ &\quad + \frac{(1 + w_t \hat{q}_{0.95,t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \hat{p}_{(0.95,1)}, \end{aligned}$$

using the quantile grid $\alpha = \{0.05, 0.10, 0.20, \dots, 0.90, 0.95\}$. $\hat{q}_{\alpha,t}$ is the one-period α -quantile forecast based on time t information, and $\hat{p}_{(\alpha^1, \alpha^2)}$ is the in-sample coverage probability between α^1 and α^2 . With exception of the lower and upper tails, the α -quantile

forecasts are assumed to be located at the midpoint between the α^1 - and α^2 -quantiles. For example, for $\widehat{q}_{0.1,t}$ the associated probability is calculated as the proportion of in-sample observations that fall between the 0.05- and 0.15-quantile. Note, this assumption basically implies that the model must be estimated for 19 quantiles at each point in time. In the univariate case, $\widehat{p}_{(\alpha^1, \alpha^2)} = \alpha^2 - \alpha^1$, and can thus be determined without calculating in-sample coverage probabilities. Due to dependence between the variables, the relevant probabilities in the multivariate case need to be determined based on in-sample coverage probabilities (see Table 4). To maintain a consistent notation between the univariate and multivariate setup, I apply the above notation also in the univariate case. Maximization is done by grid search over the interval 0 to 1 with stepsizes of 0.01. Hence, no-short sale and no-borrowing constraints are imposed from the outset.

An alternative to estimating the entire return distribution is to impose a distributional assumption. I consider the case where returns are assumed to be lognormally distributed, which implies the following optimal portfolio choice (cf. Campbell and Viceira, 2002)

$$w_t^* = \frac{E_t \widetilde{r}_{t+1} - r_{f,t} + \frac{\sigma_t^2}{2}}{\gamma \sigma_t^2},$$

where returns are now measured in logs. $E_t \widetilde{r}_{t+1}$ denotes the mean log return on the risky asset at time $t + 1$ conditional on time t information, and σ_t^2 denotes the return variance also conditional on time t information. I consider both the case where $E_t \widetilde{r}_{t+1}$ is estimated using least squares and the same state variables as in the 'quantile' setup and the case where only the historical average is used to forecast future returns (denoted by C in Table 8). σ_t^2 is estimated as the sample variance based on time t information.⁶ To make the results comparable to the case where the empirical return distribution is applied, w_t^* is restricted to lie between 0 and 1. If the optimal w_t is larger than 1, it is set to be equal to 1, and if the optimal w_t is smaller than 0, it is set to be equal to 0.

In the multivariate setup where the investor has access to both stocks and bonds, the budget constraint is given as

$$W_{t+1} = 1 + w_t^s R_{s,t+1} + w_t^b R_{b,t+1} + R_{f,t},$$

where w_t^s and w_t^b denote the portfolio weight in stocks and bonds, respectively. The remainder $1 - w_t^s - w_t^b$ is invested in the risk free asset. This implies the following maximization problem

$$\mathbf{w}_t^* = \arg \max_{w_t^s, w_t^b} \int \int \frac{(1 + w_t^s R_{s,t+1} + w_t^b R_{b,t+1} + R_{f,t})^{1-\gamma}}{1 - \gamma} f(R_{s,t+1}, R_{b,t+1} | \mathcal{F}_t) dR_{s,t+1} dR_{b,t+1},$$

where \mathbf{w}_t^* is a vector consisting of optimal stock and bond weights, and $f(R_{s,t+1}, R_{b,t+1} | \mathcal{F}_t)$ is the conditional *joint* probability distribution of future excess stock and bond returns

⁶I have also used a GARCH(1,1) model to forecast future volatility, but this makes no noticeable difference, and hence, for the sake of brevity these results are not reported.

based on the information set at time t . In discretized form, this can be rewritten as

$$\begin{aligned}
\mathbf{w}_t^* = & \arg \max_{w_t^a, w_t^b} \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(0.1,0.1),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(0,0,1)}^{(0,0,1)} \\
& + \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(0.1,0.9),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(0,0,1)}^{(0,9,1)} + \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(0.9,0.1),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(0,9,1)}^{(0,0,1)} \\
& + \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(0.9,0.9),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(0,9,1)}^{(0,9,1)} + \sum_{\alpha=0.3}^{0.7} \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(0.1,\alpha),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(0,0,1)}^{(\alpha-0.2,\alpha+0.2)} \\
& + \sum_{\alpha=0.3}^{0.7} \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(\alpha,0.1),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(\alpha-0.2,\alpha+0.2)}^{(0,0,1)} + \sum_{\alpha=0.3}^{0.7} \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(0.9,\alpha),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(0,9,1)}^{(\alpha-0.2,\alpha+0.2)} \\
& + \sum_{\alpha=0.3}^{0.7} \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(\alpha,0.9),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(\alpha-0.2,\alpha+0.2)}^{(0,9,1)} \\
& + \sum_{\alpha^1=0.3}^{0.7} \sum_{\alpha^2=0.3}^{0.7} \frac{(1 + \mathbf{w}'_t \widehat{\mathbf{q}}_{(\alpha_1,\alpha_2),t} + R_{f,t})^{1-\gamma}}{1-\gamma} \times \widehat{p}_{(\alpha_2-0.2,\alpha_2+0.2)}^{(\alpha_1-0.2,\alpha_1+0.2)},
\end{aligned}$$

where, for simplicity, I only use the quantile grid $\alpha = \{0.10, 0.30, 0.70, 0.90\}$. $\widehat{\mathbf{q}}_{(\alpha_1,\alpha_2),t}$ is a vector of jointly estimated quantile forecasts, where α_1 and α_2 are the quantiles for the first and the second variable, respectively. $\widehat{p}_{(\alpha^i,\alpha^j)}^{(\alpha^k,\alpha^l)}$ denotes the in-sample probability of the return on the first variable falling between the α^i - and α^j -quantile, while at the same time the second variable falls between the α^k - and α^l -quantile. This approach can easily be applied to a much finer quantile grid at the cost of more quantile computations.⁷ As in the univariate case, optimization is done by grid search over the interval 0 to 1 with stepsizes of 0.01. Furthermore, I impose the restriction that the sum of the portfolio weights are not allowed to be larger than one.

In the multivariate case the optimal portfolio weights using the lognormal assumption about returns are given as (cf. Campbell and Viceira, 2002)

$$\mathbf{w}_t^* = \frac{1}{\gamma} \Sigma_t^{-1} \left(E_t \widetilde{\mathbf{r}}_{t+1} - \iota r_{f,t} + \frac{\sigma_t^2}{2} \right),$$

where Σ_t^{-1} is the inverse of the variance-covariance matrix of returns, ι is a vector of ones, $\widetilde{\mathbf{r}}_{t+1}$ is the vector of log returns on the risky assets, and σ_t^2 is a vector consisting of the diagonal elements of Σ_t . The mean return vector and the variance-covariance matrix are estimated in the same way as in the univariate case. Again, to ensure results comparable to those based on the empirical return distribution, portfolio weights smaller than 0 are set to be equal to 0 and portfolio weights larger than 1 are set to be equal to 1. Furthermore, if both weights are positive and the sum of the weights is larger than 1,

⁷The grid applied here implies that at each point in time the multivariate model must be estimated for 21 quantile combinations. The more fine grid applied in the univariate case would result in 217 quantile combinations.

then both weights are reduced to ensure that the sum is 1, while still maintaining the ratio between the two weights.

The out-of-sample portfolio study follows the same recursive scheme used to evaluate out-of-sample forecast performance in Sections 4.2 and 5.3. The portfolio weights, w_t^* or \mathbf{w}_t^* , give rise to a realized utility next period of $U(W_{t+1}^*)$. Based on the realized utility we can assess the economic significance by calculating the certainty equivalent return (CER)

$$CER = \left((1 - \gamma) \frac{1}{T^*} \sum_{t=1}^{T^*} U(W_t^*) \right)^{1/(1-\gamma)} - 1,$$

where $1/T^* \sum_{t=1}^{T^*} U(W_t^*)$ is the mean realized utility and T^* is the total number of observations in the out-of-sample period. In the case with $\gamma = 1$ (log utility) the CER is calculated as

$$CER = \exp \left(\frac{1}{T^*} \sum_{t=1}^{T^*} U(W_t^*) \right) - 1.$$

The CER is defined as the certain return required by the investor for him to be indifferent between accepting the certain return and following the portfolio strategy, which can potentially give a higher return but at the cost of uncertainty. The higher the CER, the more attractive is the uncertain alternative. In the present context, the CER is measured after realization of the portfolio strategy, and thus provides a basis for assessing the economic significance. The CER can be compared directly across investment strategies.

6.2 Empirical results

Table 8 shows the per annum CER in the case where the investor besides the risk free asset can invest in stocks (Panel A), bonds (Panel B), and both stocks and bonds (Panel C). Results are shown for three values of the relative risk aversion coefficient representing low or medium risk aversion ($\gamma = 1$), high risk aversion ($\gamma = 5$), and very high risk aversion ($\gamma = 10$). Each panel gives the CER obtained using knowledge of the empirical return distribution captured through quantile regression and by assuming lognormally distributed returns. In each panel, the highest CER for each state variable across the empirical and lognormal distribution is boldfaced.

The average return on the 30-day T-bill, which represents the risk free alternative for an investor with a 1-month investment horizon, is approximately 4% per annum. From Table 8 it is clear that the CER is only below 4% for the very risk averse investor who is not restricted to invest in bonds. In all other cases, access to risky assets provide economic gains to the investor, in some cases of more than 6% per annum compared to the risk free alternative. Table 8 also shows that for the investor with low or medium risk aversion who is restricted to invest in either stocks or bonds, stocks is the most attractive investment asset, while bonds are preferable for the highly risk averse investor. For example, based on the term spread the CER for stocks is almost 3% higher than for

bonds based on $\gamma = 1$, while the CER for bonds is more than 4.5% higher than for stocks based on $\gamma = 10$.

From Panel A it is clear that nothing is gained by estimating the entire stock return distribution compared to assuming lognormality. In almost all cases, the CER based on the lognormal assumption is higher than the CER based on the empirical distribution. Losses obtained by estimating the entire stock return distribution are as large as almost 3% per annum. For bonds, however, the results are completely reversed. Gains obtained by taking the entire bond return distribution into account are as large as 1.5% per annum. The clear difference between stocks and bonds in terms of the losses or gains obtained by estimating the entire return distributions compared to imposing a lognormal assumption and only estimating the mean and the variance is not necessarily surprising. From the slope equality test in Table 1 it is clear that the state variables basically only capture location shifts in the stock return distribution. Hence, if stock returns are (close to) lognormally distributed the only effect from estimating the entire return distribution through quantile regression is increased uncertainty due to a higher number of parameters to be estimated. From the confidence bands in Figure 1 it is clear that this uncertainty can be quite large, especially in the tails of the distribution. In contrast, the results in Table 2 show that the state variables capture much more than location shifts in the bond return distribution. Hence, the use of quantile regression to capture the entire bond distribution contributes with information not captured by the lognormal distribution, and according to Panel B in Table 8 this additional information is economically significant.

Panel C reveals very mixed results regarding the economic significance of estimating the entire joint distribution compared to assuming that stocks and bonds are jointly lognormally distributed. Based on the evidence from Panels A and B, this result is not surprising. The difference in CER is in some cases very large. For the earnings-price ratio and $\gamma = 1$, the CER is almost 2.5% higher per annum for the empirical return distribution compared to the lognormal distribution. For the payout ratio this result is reversed and here the difference is more than 3% per annum. It is important to note that a very coarse quantile grid has been used in the multivariate portfolio study. Hence, the results in Panel C are affected both by differences across stocks and bonds in terms of the state variables predictive ability, and by a very coarse estimate of the entire joint distribution. However, despite these reservations, the results do illustrate that gains can be made by estimating the empirical return distribution instead of imposing a lognormal distributional assumption.

7 Concluding remarks

This paper provides detailed insights into predictability of the entire stock and bond return distribution through the use of quantile regression. The traditional focus on predictability of either the mean or volatility of returns is in many areas of financial economics insufficient. The use of quantile regression allows us to examine specific parts of the return distribution such as the tails or the center, and for a sufficiently fine grid of

quantiles, we can trace out the entire distribution. The idea of using quantile regression to examine predictability of individual return distributions can also be applied to joint return distributions.

An empirical analysis on US data shows that certain parts of both the stock and bond return distribution are predictable as a function of a set economic state variables, also even if the mean is unpredictable. The analysis also shows large differences between stocks and bonds in terms of predictability. Although certain parts of the stock return distribution are significantly predictable, the effect across the return distribution is, in general, not significantly different, which implies that the state variables only predict location shifts. This is in clear contrast to the bond return distribution, where the state variables predict changes in the entire distribution, including volatility and skewness. The empirical analysis also shows that, to a certain extent, the joint stock-bond return distribution is also predictable. The relative accuracy of the different state variables in predicting stocks and bonds out-of-sample varies across the return distribution. Some state variables perform relatively good in the tails of the distribution, while others perform relatively good in the center. A similar result is found for the joint return distribution between stocks and bonds.

Overall, quantile regression gives a very flexible empirical approach to model return distributions, free of the usual parametric assumptions. This applies in both the univariate and multivariate case. The quality of the empirical model in terms of capturing the actual return distribution both in- and out-of-sample is of course conditional on the economic state variable used to predict future returns. An out-of-sample portfolio study based on an investor with power utility shows that for bonds it bears economic significance to apply the empirical return distribution captured through quantile regression in portfolio decisions compared to imposing a lognormal distributional assumption about returns. This is not the case for stocks, which can be explained by the fact that for this asset the state variables capture only location shifts. Given stock returns are (almost) lognormally distributed, the use of quantile regression basically only increases uncertainty due to the higher number of parameters to be estimated. For the multivariate case where the investor can invest in both stocks and bonds, the results are mixed, which follows naturally from the differences between stocks and bonds in the individual cases. Another important limitation to the multivariate portfolio study is the very coarse quantile grid used to capture the joint distribution.

This paper examines the predictive ability of individual state variables on the stock and bond return distribution. The out-of-sample analysis shows that the relative predictive ability of the state variables varies a lot across the return distributions. This suggests that a combination of state variables would be fruitful in capturing the entire distribution compared to individual state variables. Cenesizoglu and Timmermann (2008) apply an equal-weighted combination of quantile forecasts across a set of individual state variables in order to incorporate information from all state variables without having to estimate additional parameters. The disadvantage of this approach is that the individual state variables' relative good forecasting ability in different parts of the distribution is not fully utilized. Instead a recursive procedure that continuously applies the relatively best state

variables to forecast the different return quantiles would be a way to ensure the best possible estimate of the future return distribution. I leave this as an interesting topic for future research.

8 References

Ang, A. and G. Bekaert (2002). International asset allocation with regime shifts. *Review of Financial Studies* 15, 1137-1187.

Baele, L., G. Bekaert, and K. Inghelbrecht (2010). The determinants of stock and bond return comovements. *Review of Financial Studies* 23, 2374-2428.

Barrodale, I. and F.D.K. Roberts (1974). Solution of an overdetermined system of equations in the l_1 norm. *Communications of the ACM* 17, 319-320.

Barnes, M.L. and A.W. Hughes (2002). A quantile regression analysis of the cross section of stock market returns. Working paper, Federal Reserve Bank of Boston.

Basset, G. and H. Chen (2001). Portfolio style: Return-based attribution using quantile regression. *Empirical Economics* 26, 293-305.

Buchinsky, M. (1995). Estimating the asymptotic covariance matrix for quantile regression models: A Monte Carlo study. *Journal of Econometrics* 68, 303-338.

Campbell, J.Y., Y.L. Chan, and L.M. Viceira (2003). A multivariate model of strategic asset allocation. *Journal of Financial Economics* 67, 41-80.

Campbell, J.Y. and J.H. Cochrane (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 205-251.

Campbell, J.Y. and L.M. Viceira (2002). *Strategic Asset Allocation - Portfolio Choice for Long-Term Investors*. Oxford University Press.

Cenesizoglu, T. and A. Timmermann (2008). Is the distribution of stock returns predictable? Working paper, University of California San Diego.

Chakraborty, B. (1999). On multivariate median regression. *Bernoulli* 5, 683-703.

Chakraborty, B. (2003). On multivariate quantile regression. *Journal of Statistical Planning and Inference* 110, 109-132.

Chakraborty, B. and P. Chaudhuri (1996). On a transformation and re-transformation technique for constructing affine equivariant multivariate median. *Proceedings of the American Mathematical Society* 124, 2539-2547.

Chaudhuri, P. (1996). On a geometric notion of quantiles for multivariate data. *Journal of the American Statistical Association* 91, 862-872.

Chaudhuri, P. and D. Sengupta (1993). Sign tests in multidimension: Inference based on the geometry of the data cloud. *Journal of the American Statistical Association* 88, 1363-1370.

Christiansen, C. and A. Rinaldo (2007). Realized bond-stock correlation: macroeconomic announcement effects. *Journal of Futures Markets* 27, 439-469.

Connolly R., C. Stivers, and L. Sun (2005). Stock market uncertainty and the stock-bond return relation. *Journal of Financial and Quantitative Analysis* 40, 161-194.

Corradi, V. and N.R. Swanson (2006). Predictive density evaluation. In: Elliott, G., C.W.J. Granger, and A. Timmermann (Eds.), *Handbook of Economic Forecasting* 1, 197-284.

De Gooijer, J.G. and D. Zerom (2003). On additive conditional quantiles with high-dimensional covariates. *Journal of the American Statistical Association* 98, 135-146.

Diebold, F.X. and R.S. Mariano (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253-263.

Dittmar, R.F. (2002). Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *Journal of Finance* 57, 369-403.

Engle, R.F. and S. Manganelli (2004). CAViaR: Conditional autoregressive value at risk by regression quantiles. *Journal of Business and Economic Statistics* 22, 367-381.

Engsted, T. and T.Q. Pedersen (2010). Return predictability and intertemporal asset allocation: Evidence from a bias-adjusted VAR model. Working paper, Aarhus University.

Guidolin, M., S. Hyde, D. McMillan, and S. Ono (2009). Non-linear predictability in stock and bond returns: When and where is it exploitable? *International Journal of Forecasting* 25, 373-399.

Guidolin, M. and A. Timmermann (2006). An econometric model of nonlinear dynamics in the joint distribution of stock and bond returns. *Journal of Applied Econometrics* 21, 1-22.

Guidolin, M. and A. Timmermann (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control* 31, 3503-3544.

Harvey, C.R. and A. Siddique (2000). Conditional skewness in asset pricing tests. *Journal of Finance* 55, 1263-1295.

Ilmanen, A. (1995). Time-varying expected returns in international bond markets. *Journal of Finance* 50, 481-506.

Ilmanen, A. (2003). Stock-bond correlations. *Journal of Fixed Income* 13, 55-66.

Kirby, C. (1997). Measuring the predictable variation in stock and bond returns. *Review of Financial Studies* 10, 579-630.

Koenker, R. (2005). *Quantile Regression*. Cambridge University Press.

Koenker, R. and G. Basset (1978). Regression quantiles. *Econometrica* 46, 33-50.

Koenker, R. and G. Basset (1982). Robust tests for heteroskedasticity based on regression quantiles. *Econometrica* 50, 43-62.

Koenker, R. and V. d'Orey (1987). Computing regression quantiles. *Journal of the Royal Statistical Society (C)* 36, 383-393.

Koenker, R. and B.J. Park (1996). An interior point algorithm for nonlinear quantile regression. *Journal of Econometrics* 71, 265-283.

Lettau, M. and S.C. Ludvigson (2010). Measuring and modeling variation in the risk-return trade-off. In: Ait-Shalia, Y. and L.P. Hansen (Eds.), *Handbook of Financial Econometrics* 1, 617-690.

Marquering, W. and M. Verbeek (2004). The economic value of predicting stock index return and volatility. *Journal of Financial and Quantitative Analysis* 39, 407-429.

McMillan, D.G. and M.E. Wohar (2009). Stock return predictability and dividend-price ratio: A nonlinear approach. *International Journal of Finance and Economics* (forthcoming).

Patton, A.J. (2004). On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. *Journal of Financial Econometrics* 2, 130-168.

Paye, B. (2009). Which variables forecast aggregate stock market volatility? Working paper, Rice University.

Rapach, D.E. and M.E. Wohar (2005). Valuation ratios and long-horizon stock price predictability. *Journal of Applied Econometrics* 20, 327-344.

Sharpe, W. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance* 19, 425-444.

Tay, A.S. and K.F. Wallis (2002). Density forecasting: a survey. In: Clements, M.P. and D.F. Hendry (Eds.), *A Companion to Economic Forecasting*, 45-68.

Viceira, L.M. (2010). Bond risk, bond return volatility, and the term structure of interest rates. *International Journal of Forecasting* (forthcoming).

Welch, I. and A. Goyal (2008). A comprehensive look at the empirical performance of the equity premium prediction. *Review of Financial Studies* 21, 1455-1508.

White, H. (1980). A heteroscedasticity-consistent covariance matrix estimator and a direct test for heteroscedasticity. *Econometrica* 48, 817-838.

9 Appendix: Algorithm to carry out multivariate quantile regression

The aim is to estimate

$$\mathbf{y}_{t+1} = \beta \mathbf{x}_t + \mathbf{e}_{t+1}, \quad 1 \leq t \leq T - 1 \quad (7)$$

where \mathbf{y}_{t+1} is a $(d \times 1)$ response vector, \mathbf{x}_t is a $(k \times 1)$ vector of regressors, β is a $(d \times k)$ coefficient matrix, and \mathbf{e}_{t+1} is a $(d \times 1)$ error vector.

1. Estimate (7) using least squares to obtain an estimate, $\widehat{\Sigma}$, of the covariance matrix Σ associated with the distribution of the error vector \mathbf{e} from the data.
2. Define S to be the set of all subsets of $d + k$ (number of response variables plus number of regressors) indices from the set $\{1, 2, \dots, T - 1\}$, i.e.

$$\begin{aligned} A &= \{\mathbf{a} : \mathbf{a} \subset \{1, 2, \dots, T - 1\} \text{ and } \#\{i : i \in \mathbf{a}\} = d\}, \\ B &= \{\mathbf{b} : \mathbf{b} \subset \{1, 2, \dots, T - 1\} \text{ and } \#\{i : i \in \mathbf{b}\} = k\}, \\ S &= \{\theta = \mathbf{a} \cup \mathbf{b} : \mathbf{a} \in A, \mathbf{b} \in B, \mathbf{a} \cap \mathbf{b} = \emptyset\}. \end{aligned}$$

3. Fix $\theta = \{i_1, \dots, i_k, j_1, \dots, j_d\} \in S$.
4. Construct $W(\theta)$ to be a $k \times k$ matrix with $\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_k}$ as columns.
5. Construct $H(\theta)$ to be a $d \times k$ matrix with $\mathbf{y}_{i_1}, \dots, \mathbf{y}_{i_k}$ as columns.
6. Construct $P(\theta)$ to be a $d \times d$ matrix with $\mathbf{y}_{j_1} - H(\theta)[W(\theta)]^{-1} \mathbf{x}_{j_1}, \dots, \mathbf{y}_{j_d} - H(\theta)[W(\theta)]^{-1} \mathbf{x}_{j_d}$ as columns.
7. Compute

$$V(\theta) = [P(\theta)]' \widehat{\Sigma}^{-1} P(\theta).$$

8. Compute

$$n(\theta) = \frac{[\text{trace}\{V(\theta)\}]/2}{[\det\{V(\theta)\}]^{1/2}},$$

and note that $n(\theta) \geq 1$.

9. Minimize $n(\theta)$ with respect to $\theta \in S$ and call this $\widehat{\theta}$. According to Chakraborty (2003), we can reduce the amount of computation time required for searching for the optimal θ by stopping whenever $n(\theta)$ is sufficiently close to 1. Steps 1-9 represent the search for the optimal 'data-driven coordinate system'.
10. Repeat steps 4-6 using $\widehat{\theta}$ to obtain $P(\widehat{\theta})$.
11. Choose the values in the index vector, \mathbf{u} , which is an element of the open unit d -dimensional ball $B^{(d)} = \{\mathbf{u} \mid \mathbf{u} \in \mathbb{R}^d, \|\mathbf{u}\| < 1\}$. $\|\cdot\|$ denotes the Euclidian norm.

12. Form the transformed index vector

$$v(\hat{\theta}) = \begin{cases} \frac{[P(\hat{\theta})]^{-1} \mathbf{u}}{\|[P(\hat{\theta})]^{-1} \mathbf{u}\|} \|\mathbf{u}\| & \text{if } \mathbf{u} \neq \mathbf{0} \\ 0 & \text{if } \mathbf{u} = \mathbf{0} \end{cases}.$$

13. Estimate the multivariate regression quantiles as

$$\hat{\beta}_\alpha = \arg \min_{\beta_\alpha \in \mathbb{R}^{d \times k}} \sum_{\substack{t=1 \\ t \notin \theta}}^{T-1} \left\{ \left\| [P(\hat{\theta})]^{-1} (\mathbf{y}_{t+1} - \beta_\alpha \mathbf{x}_t) \right\| + [v(\hat{\theta})]' [P(\hat{\theta})]^{-1} (\mathbf{y}_{t+1} - \beta_\alpha \mathbf{x}_t) \right\}. \quad (8)$$

$[P(\hat{\theta})]^{-1} \mathbf{y}_{t+1}$ represents the transformation of the response vector based on the 'data-driven coordinate system'. As we transform the response vector, we also need to modify the orientation of the index vector \mathbf{u} . This is done in step 12 as $[P(\hat{\theta})]^{-1} \mathbf{u}$, and to preserve the Euclidean norm of the vector \mathbf{u} , $[P(\hat{\theta})]^{-1} \mathbf{u}$ is rescaled by multiplying with $\|\mathbf{u}\| / \|[P(\hat{\theta})]^{-1} \mathbf{u}\|$. Hence, \mathbf{u} in (3) is replaced by $v(\hat{\theta})$ in (8). Finally, the $v(\hat{\theta})$ th regression quantiles computed using the transformed response vector are retransformed back to the original coordinate system. $[P(\hat{\theta})]^{-1} \beta_\alpha \mathbf{x}_t$ represents this step.

Chakraborty (2003) shows that Newton-Raphson like algorithms can be used to compute transformation-retransformation regression quantiles from multivariate observations, i.e. to perform the minimization in (8). I apply this approach combined with the BFGS method to avoid evaluating the Hessian directly.

10 Tables and figures

Table 1. Slope coefficient estimates for stock returns in the univariate setting.

	DP	EP	DE	NTIS	SVAR	TBL	TMS	DFS	INFL
Panel A: Slope coefficients based on least squares									
	0.302 (3.10)	0.123 (2.34)	0.007 (0.53)	-0.003 (-0.03)	-1.069 (-2.73)	-0.114 (-2.21)	0.200 (1.69)	0.185 (0.47)	-1.126 (-3.39)
Panel B: Slope coefficients based on quantile regression									
0.05	0.361 (1.47)	0.114 (0.94)	-0.009 (-0.26)	0.296 (1.25)	-6.220 (-3.77)	-0.149 (-0.96)	0.196 (0.56)	-0.735 (-0.88)	-1.600 (-1.50)
0.10	0.305 (1.54)	0.138 (1.32)	-0.007 (-0.16)	0.281 (1.22)	-5.483 (-2.45)	-0.175 (-1.91)	0.548 (2.03)	-1.239 (-1.43)	-1.361 (-1.51)
0.20	0.151 (1.23)	0.049 (0.87)	0.010 (0.57)	-0.050 (-0.34)	-2.074 (-1.04)	-0.169 (-2.48)	0.428 (2.81)	-0.436 (-0.79)	-1.416 (-2.90)
0.30	0.179 (1.22)	0.058 (0.70)	0.024 (1.13)	-0.066 (-0.69)	-1.077 (-1.05)	-0.184 (-2.59)	0.192 (1.17)	-0.438 (-0.89)	-1.740 (-4.04)
0.40	0.203 (1.36)	0.041 (0.50)	0.020 (1.21)	-0.049 (-0.43)	-1.205 (-1.55)	-0.186 (-2.81)	0.225 (1.72)	-0.568 (-1.00)	-1.759 (-4.68)
0.50	0.226 (2.14)	0.092 (1.50)	0.006 (0.41)	-0.060 (-0.80)	-1.025 (-1.12)	-0.181 (-2.91)	0.118 (1.21)	0.149 (0.33)	-1.415 (-4.02)
0.60	0.285 (2.57)	0.110 (2.07)	0.012 (0.79)	0.035 (0.38)	-0.384 (-0.35)	-0.113 (-1.85)	0.069 (0.50)	0.399 (1.02)	-1.378 (-3.57)
0.70	0.433 (2.89)	0.128 (1.67)	0.011 (0.65)	0.026 (0.17)	-0.487 (-0.32)	-0.107 (-1.30)	0.046 (0.26)	0.525 (0.98)	-0.658 (-1.25)
0.80	0.403 (3.70)	0.141 (2.69)	0.022 (1.21)	-0.086 (-0.74)	1.826 (1.17)	-0.109 (-1.76)	0.146 (1.02)	0.748 (1.20)	-0.427 (-0.72)
0.90	0.292 (2.38)	0.154 (2.02)	-0.004 (-0.16)	-0.261 (2.19)	3.269 (2.12)	-0.047 (-0.67)	0.200 (1.02)	1.049 (1.57)	-0.407 (-0.79)
0.95	0.258 (1.26)	0.212 (1.52)	0.007 (0.20)	-0.246 (-1.13)	3.336 (0.81)	0.016 (0.12)	0.410 (1.53)	2.320 (2.03)	-0.646 (-0.96)
Panel C: Wald test statistic from slope equality test									
	6.45 [0.776]	4.53 [0.921]	5.48 [0.857]	10.83 [0.371]	22.08 [0.015]	4.80 [0.904]	11.59 [0.313]	9.69 [0.468]	9.06 [0.526]

Notes: In Panel B, the numbers in the first column denote the quantiles for which the regression is carried out. In Panel A and B, numbers in parentheses denote t-statistics. In Panel C, numbers in brackets denote p-values. Slope coefficients for which the t-statistic is higher than 1.96 or lower than -1.96 are boldfaced.

Table 2. Slope coefficient estimates for bond returns in the univariate setting.

	DP	EP	DE	NTIS	SVAR	TBL	TMS	DFS	INFL
Panel A: Slope coefficients based on least squares									
	-0.002 (-0.08)	-0.012 (-0.61)	0.008 (1.86)	-0.082 (-2.66)	0.442 (2.21)	0.010 (0.36)	0.097 (1.69)	0.330 (1.98)	-0.242 (-1.98)
Panel B: Slope coefficients based on quantile regression									
0.05	0.007 (0.04)	-0.063 (-0.81)	0.047 (2.37)	0.061 (0.73)	0.234 (0.44)	-0.238 (-3.59)	0.064 (0.45)	-1.428 (-4.69)	-1.586 (-5.96)
0.10	0.155 (2.66)	-0.018 (-0.41)	0.030 (3.45)	0.049 (0.86)	0.217 (0.53)	-0.226 (-8.33)	0.051 (0.63)	-1.164 (-2.96)	-1.140 (-5.65)
0.20	0.107 (3.84)	0.015 (0.56)	0.020 (4.51)	0.014 (0.26)	0.123 (0.56)	-0.151 (-5.40)	0.033 (0.62)	-0.526 (-2.33)	-0.545 (-3.48)
0.30	0.076 (4.28)	0.018 (1.44)	0.014 (3.46)	-0.046 (-1.10)	0.077 (0.33)	-0.089 (-4.61)	0.085 (1.87)	-0.145 (-0.97)	-0.343 (-2.68)
0.40	0.030 (2.10)	0.008 (1.13)	0.010 (3.64)	-0.067 (-2.29)	0.209 (0.82)	-0.050 (-3.99)	0.140 (3.92)	0.140 (1.17)	-0.122 (-1.99)
0.50	0.000 (0.02)	-0.002 (-0.28)	0.003 (1.12)	-0.069 (-3.69)	0.516 (2.01)	0.004 (0.25)	0.139 (4.43)	0.323 (3.56)	-0.076 (-2.41)
0.60	-0.047 (-3.02)	-0.020 (-2.62)	-0.003 (-0.70)	-0.083 (-3.31)	0.602 (2.76)	0.052 (3.27)	0.168 (4.25)	0.631 (4.20)	-0.057 (-0.90)
0.70	-0.105 (-4.35)	-0.039 (-3.15)	-0.006 (-1.26)	-0.126 (-4.46)	0.709 (2.65)	0.112 (6.65)	0.215 (4.32)	1.032 (6.12)	-0.103 (-1.43)
0.80	-0.148 (-5.00)	-0.053 (-2.12)	-0.004 (-0.62)	-0.165 (-4.04)	0.628 (1.87)	0.160 (6.51)	0.199 (2.80)	1.286 (5.00)	-0.173 (-0.97)
0.90	-0.145 (-1.88)	-0.034 (-0.73)	0.007 (0.87)	-0.168 (-2.59)	1.199 (1.71)	0.246 (6.17)	0.115 (1.29)	1.763 (9.35)	-0.305 (-0.96)
0.95	0.056 (0.34)	0.053 (0.79)	-0.002 (-0.12)	-0.124 (-1.31)	1.896 (2.45)	0.292 (5.05)	0.002 (0.02)	1.706 (3.38)	-0.009 (-0.02)
Panel C: Wald test statistic from slope equality test									
	97.21 [0.000]	31.44 [0.001]	39.41 [0.000]	15.57 [0.113]	10.22 [0.421]	185.23 [0.000]	20.08 [0.029]	116.54 [0.000]	43.79 [0.000]

Notes: In Panel B, the numbers in the first column denote the quantiles for which the regression is carried out. In Panel A and B, numbers in parentheses denote t-statistics. In Panel C, numbers in brackets denote p-values. Slope coefficients for which the t-statistic is higher than 1.96 or lower than -1.96 are boldfaced.

Table 3. Loss differential in the univariate setting.

	DP	EP	DE	NTIS	SVAR	TBL	TMS	DFS	INFL	PQ
Panel A: $\alpha = 0.05$										
DP		-0.25	0.75	0.24	1.44	1.45	-0.64	2.70	2.59	2.33
EP	0.78		0.94	0.46	1.14	1.67	-0.54	3.24	3.05	1.53
DE	1.55	1.24		-0.70	-0.36	1.35	-1.35	2.38	2.13	-0.12
NTIS	0.23	-0.22	-0.96		1.13	1.46	-0.78	2.93	2.85	1.98
SVAR	-1.46	-1.54	-1.73	-1.65		1.27	-1.13	2.63	2.41	1.92
TBL	0.20	0.01	-0.31	0.09	1.38		-1.91	0.29	-0.18	-1.18
TMS	0.42	0.27	-0.02	0.33	1.39	0.27		3.19	3.28	1.39
DFS	-0.28	-1.06	-2.07	-0.46	1.39	-0.33	-0.54		-0.66	-2.48
INFL	0.20	0.00	-0.34	0.08	1.33	-0.01	-0.28	0.32		-2.23
PQ	0.25	-0.49	-1.62	-0.09	1.48	-0.16	-0.41	0.85	-0.15	
Panel B: $\alpha = 0.50$										
DP		1.64	0.68	2.65	0.95	-0.83	1.75	1.62	3.16	2.05
EP	0.60		-0.17	2.50	0.54	-1.08	1.58	1.40	3.01	1.49
DE	-0.08	-0.26		2.02	0.57	-1.27	1.67	1.16	1.56	0.59
NTIS	0.66	0.49	0.53		-1.57	-2.18	0.23	-0.79	-1.73	-2.27
SVAR	0.79	0.64	0.76	0.23		-1.21	1.20	0.78	0.22	-0.32
TBL	-0.04	-0.16	0.01	-0.36	-0.52		2.13	1.36	1.70	1.29
TMS	0.85	0.70	0.94	0.40	0.15	0.68		-0.73	-1.23	-1.49
DFS	0.74	0.61	0.58	0.03	-0.21	0.36	-0.41		-0.70	-1.16
INFL	-1.48	-1.65	-1.74	-1.88	-1.96	-1.09	-2.19	-1.85		-3.18
PQ	-0.39	-0.89	-0.18	-1.04	-1.18	-0.11	-1.22	-1.32	1.68	
Panel C: $\alpha = 0.95$										
DP		0.12	0.11	0.27	1.25	1.18	-0.75	2.53	0.15	1.43
EP	0.00		0.04	0.13	1.07	1.26	-0.73	2.79	0.04	0.88
DE	-0.50	-0.46		0.08	1.13	1.41	-0.88	2.65	-0.00	0.72
NTIS	0.01	0.01	0.57		0.84	1.07	-0.76	2.21	-0.09	0.53
SVAR	-0.02	-0.02	0.23	-0.03		0.62	-1.71	1.85	-1.19	-0.72
TBL	0.51	0.42	1.25	0.48	0.29		-1.64	0.58	-1.27	-0.90
TMS	-0.03	-0.03	0.76	-0.05	0.00	-0.81		2.48	0.92	1.46
DFS	-0.85	-0.89	-0.50	-1.16	-0.43	-1.11	-1.15		-2.48	-2.16
INFL	-0.73	-0.61	-0.73	-0.75	-0.32	-1.67	-1.08	0.25		1.13
PQ	-0.79	-0.74	-1.80	-0.90	-0.39	-1.43	-1.16	0.17	-0.32	

Notes: The table shows t-statistics from pairwise out-of-sample loss comparisons. t-statistics higher than 1.96 or lower than -1.96 are boldfaced. In each panel, numbers below the main diagonal give the results for stocks, while numbers above the main diagonal give the results for bonds.

Table 4. In-sample coverage probabilities in the multivariate setting.

(α_s, α_b)	DP	EP	DE	NTIS	SVAR	TBL	TMS	DFS	INFL	PQ	Ind.
(0.05,0.05)	0.006	0.007	0.007	0.006	0.007	0.003	0.004	0.007	0.005	0.007	0.003
(0.10,0.10)	0.021	0.022	0.023	0.023	0.023	0.016	0.023	0.023	0.019	0.023	0.010
(0.50,0.50)	0.269	0.276	0.273	0.275	0.269	0.269	0.277	0.276	0.270	0.275	0.250
(0.90,0.90)	0.683	0.675	0.677	0.682	0.676	0.683	0.681	0.671	0.682	0.676	0.810
(0.95,0.95)	0.787	0.797	0.799	0.799	0.794	0.793	0.794	0.788	0.799	0.797	0.903
(0.05,0.95)	0.195	0.195	0.197	0.200	0.205	0.200	0.191	0.207	0.190	0.195	0.048
(0.95,0.05)	0.023	0.021	0.021	0.023	0.022	0.015	0.020	0.020	0.019	0.022	0.048
(0.05,0.50)	0.009	0.007	0.005	0.004	0.009	0.007	0.005	0.011	0.004	0.009	0.025
(0.50,0.05)	0.011	0.011	0.010	0.010	0.010	0.005	0.009	0.009	0.006	0.010	0.025
(0.50,0.95)	0.471	0.464	0.472	0.477	0.470	0.476	0.475	0.480	0.475	0.472	0.475
(0.95,0.50)	0.539	0.540	0.543	0.529	0.522	0.515	0.528	0.522	0.539	0.541	0.475

Notes: For each quantile combination, the table gives the in-sample coverage probabilities calculated over the entire sample period, 1941:5-2008:12. The final column gives the coverage probability given independence (the product of the quantiles).

Table 5. Slope coefficient estimates in multivariate quantile regression.

(α_s, α_b)	DP	EP	DE	NTIS	SVAR	TBL	TMS	DFS	INFL
Panel A: Slope coefficient estimates for stock returns									
(0.05,0.05)	0.111 (0.96)	-0.018 (-0.25)	0.034 (2.07)	-0.042 (-0.38)	-1.796 (-2.02)	-0.320 (-4.75)	0.328 (2.61)	-1.169 (-2.84)	-2.202 (-4.72)
(0.50,0.50)	0.269 (3.39)	0.118 (2.37)	0.015 (1.02)	-0.040 (-0.44)	-0.677 (-0.70)	-0.176 (-3.36)	0.205 (2.22)	-0.036 (-0.08)	-1.410 (-5.60)
(0.95,0.95)	0.422 (3.24)	0.224 (2.94)	-0.007 (-0.43)	-0.166 (-1.31)	1.132 (0.68)	0.098 (1.72)	-0.025 (-0.16)	1.852 (3.74)	-0.414 (-0.94)
(0.05,0.95)	0.284 (2.04)	0.095 (1.33)	0.015 (0.62)	0.156 (1.20)	-3.362 (-2.96)	-0.161 (-2.38)	0.283 (1.93)	-1.234 (-2.74)	-1.623 (-3.19)
(0.95,0.05)	0.267 (2.66)	0.130 (2.13)	-0.005 (-0.34)	-0.099 (-1.09)	0.711 (0.53)	-0.074 (-1.22)	0.168 (1.32)	0.988 (2.12)	-0.708 (-1.57)
Panel B: Slope coefficient estimates for bond returns									
(0.05,0.05)	-0.076 (-0.80)	-0.109 (-1.76)	0.039 (3.40)	-0.033 (-0.44)	-0.697 (-0.88)	-0.191 (-3.45)	0.189 (1.54)	-1.291 (-5.02)	-1.256 (-2.77)
(0.50,0.50)	-0.004 (-0.30)	-0.007 (-1.02)	0.006 (1.95)	-0.069 (-2.96)	0.440 (1.77)	-0.002 (-0.14)	0.138 (3.42)	0.362 (3.07)	-0.128 (-1.75)
(0.95,0.95)	0.100 (0.88)	0.088 (1.27)	-0.006 (-0.39)	-0.212 (-2.84)	2.295 (1.89)	0.233 (5.00)	-0.050 (-0.40)	2.388 (5.61)	0.434 (0.95)
(0.05,0.95)	0.044 (0.46)	0.058 (1.14)	-0.005 (-0.31)	-0.211 (-4.18)	2.694 (2.62)	0.196 (3.82)	0.040 (0.38)	2.024 (5.80)	0.329 (0.79)
(0.95,0.05)	-0.049 (-0.66)	-0.082 (-1.75)	0.032 (2.93)	-0.039 (-0.79)	-0.510 (-1.00)	-0.200 (-3.54)	0.164 (1.65)	-1.221 (-4.15)	-1.344 (-3.46)

Notes: For each quantile combination, the table gives the slope coefficients for stock and bond returns. Numbers in parentheses denote t-statistics. Slope coefficients for which the t-statistic is higher than 1.96 or lower than -1.96 are boldfaced.

Table 6. Loss differential between joint and univariate quantile estimation.

(α_s, α_b)	DP	EP	DE	NTIS	SVAR	TBL	TMS	DFS	INFL	PQ
(0.05,0.05)	1.99	1.84	1.55	0.96	2.42	3.60	0.28	1.81	1.46	1.79
(0.10,0.10)	0.65	0.20	0.48	-0.74	0.48	2.26	-1.54	0.04	0.44	0.12
(0.50,0.50)	-0.09	1.60	-0.46	-1.70	0.84	0.48	-0.30	1.82	-1.38	0.24
(0.90,0.90)	-0.14	1.40	1.90	2.84	2.93	-0.22	3.39	2.90	2.04	2.31
(0.95,0.95)	1.49	3.14	3.21	4.64	4.79	1.56	4.74	5.02	3.22	3.83
(0.05,0.95)	1.59	1.12	1.01	0.00	1.55	2.80	-0.61	0.40	0.80	0.91
(0.95,0.05)	-0.16	1.76	1.77	3.65	3.58	-0.13	3.81	3.28	2.28	2.55
(0.05,0.50)	0.80	1.25	0.60	0.66	-0.42	-1.69	2.53	0.64	1.55	1.41
(0.50,0.05)	13.79	16.25	14.03	15.75	12.38	13.37	16.84	14.58	16.93	16.46
(0.50,0.95)	0.67	1.06	-1.70	-1.45	-1.67	-0.09	0.07	-0.60	-0.03	-0.33
(0.95,0.50)	15.85	17.90	14.19	13.88	16.12	11.94	17.75	15.39	17.57	17.51

Notes: For each quantile combination, the table gives the t-statistics from out-of-sample loss comparisons between joint and univariate quantile estimation. t-statistics higher than 1.96 or lower than -1.96 are boldfaced.

Table 7. Loss differential in the multivariate setting.

	DP	EP	DE	NTIS	SVAR	TBL	TMS	DFS	INFL	PQ
Panel A: $\alpha = (0.05, 0.05)$ and $\alpha = (0.95, 0.95)$										
DP		4.17	3.69	4.50	3.88	3.14	2.86	6.06	2.37	4.31
EP	5.04		2.20	3.34	2.90	1.86	1.47	5.34	0.82	2.46
DE	-1.23	-3.72		1.96	1.89	0.03	-0.60	3.42	-2.41	-0.09
NTIS	4.21	0.94	3.88		0.33	-1.34	-2.06	2.62	-2.94	-2.08
SVAR	-1.17	-3.72	0.17	-3.64		-1.56	-2.12	2.10	-2.75	-1.98
TBL	-3.01	-4.34	-3.12	-4.51	-2.38		-0.46	3.51	-1.36	-0.06
TMS	2.82	1.15	4.08	0.69	3.23	5.77		3.64	-1.10	0.63
DFS	0.03	-4.93	1.38	-3.75	1.27	3.32	-3.02		-3.89	-3.55
INFL	0.29	-1.79	1.45	-2.11	0.90	4.47	-2.90	0.30		2.34
PQ	2.18	-3.71	2.69	-3.41	2.41	3.66	-2.49	1.77	0.44	
Panel B: $\alpha = (0.50, 0.50)$										
DP										
EP	0.32									
DE	0.11	-0.03								
NTIS	0.73	0.73	0.56							
SVAR	0.80	0.73	0.67	0.01						
TBL	-0.06	-0.12	-0.16	-0.44	-0.49					
TMS	0.62	0.56	0.63	0.06	0.05	0.53				
DFS	0.26	0.14	0.13	-0.64	-0.65	0.17	-0.60			
INFL	-1.44	-1.54	-2.16	-2.03	-2.04	-1.06	-2.03	-1.84		
PQ	-0.09	-0.35	-0.23	-1.00	-1.06	0.03	-0.84	-0.67	1.97	
Panel C: $\alpha = (0.05, 0.95)$ and $\alpha = (0.95, 0.05)$										
DP		3.09	4.12	4.58	4.19	-0.27	3.01	6.74	3.56	4.88
EP	5.04		2.81	3.88	3.38	-0.99	2.19	6.61	2.46	3.67
DE	-1.87	-4.03		1.50	1.41	-3.48	0.38	3.16	0.26	0.08
NTIS	1.94	-0.78	2.67		-0.34	-3.25	-0.95	1.52	-1.28	-1.82
SVAR	-2.33	-3.58	-1.10	-3.06		-3.23	-0.65	1.90	-1.07	-1.64
TBL	-3.62	-4.87	-3.52	-4.14	-1.13		3.05	3.99	3.86	2.98
TMS	2.47	1.19	4.13	1.46	3.52	6.44		2.30	-0.25	-0.35
DFS	-1.03	-4.30	0.84	-2.51	2.00	3.11	-3.01		-2.54	-3.57
INFL	-0.45	-2.22	1.17	-1.44	1.44	4.55	-3.06	0.17		-0.19
PQ	0.91	-3.78	3.28	-1.45	2.62	4.23	-2.49	1.73	0.91	

Notes: The table shows t-statistics from pairwise out-of-sample loss comparisons. t-statistics higher than 1.96 or lower than -1.96 are boldfaced. In Panel A, numbers below the main diagonal give the results for the quantile combination (0.05,0.05), while numbers above the main diagonal give the results

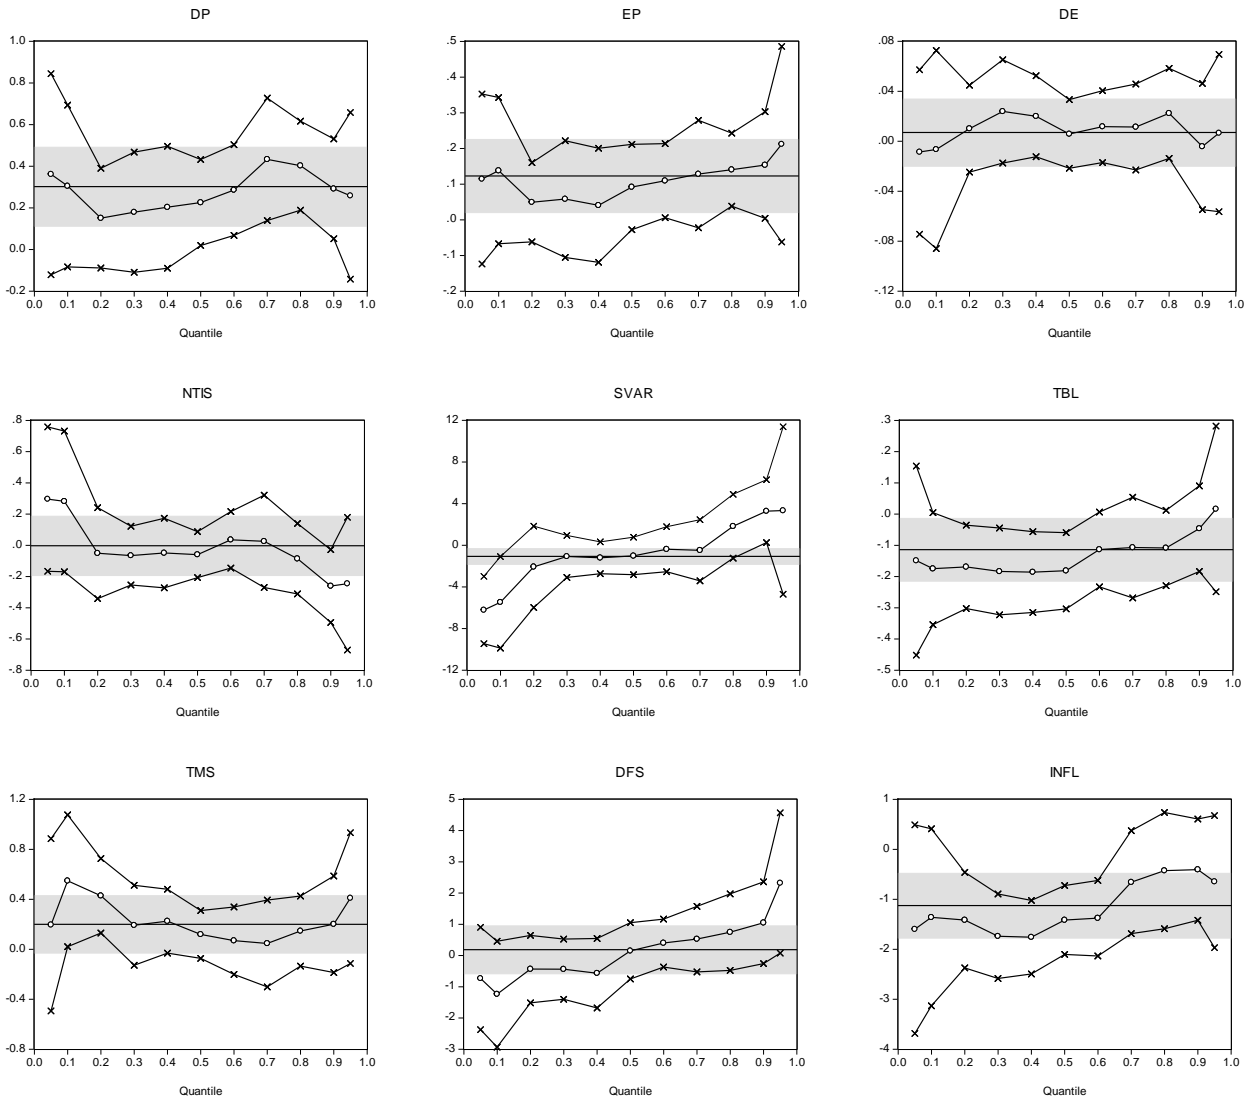
for the quantile combination (0.95,0.95). In Panel C, numbers below the main diagonal give the results for the quantile combination (0.05,0.95), while numbers above the main diagonal give the results for the quantile combination (0.95,0.05).

Table 8. Certainty equivalent return.

γ	DP	EP	DE	NTIS	SVAR	TBL	TMS	DFS	INFL	PQ/C	
Panel A: Stocks											
Emp. dist.	1	8.79	9.19	8.77	8.97	9.64	8.49	10.57	9.19	9.86	9.19
	5	5.49	5.43	4.19	4.39	4.99	5.72	6.60	4.02	6.74	4.09
	10	4.85	3.78	3.80	0.92	4.37	3.79	2.27	3.12	4.18	3.56
Logn. dist.	1	8.90	9.55	9.81	10.14	10.35	8.71	10.21	9.14	10.51	10.11
	5	7.29	7.12	5.59	5.80	6.39	5.76	7.22	5.80	7.88	6.02
	10	6.48	6.43	5.49	3.79	5.82	5.33	3.87	5.06	5.63	5.77
Panel B: Bonds											
Emp. dist.	1	7.73	7.80	7.43	7.25	7.68	7.64	7.79	7.76	8.37	7.69
	5	6.85	6.88	7.02	6.60	6.95	7.11	7.31	7.02	7.73	7.08
	10	6.08	6.15	6.72	6.26	6.27	6.56	6.92	6.15	7.18	6.45
Logn. dist.	1	6.70	6.41	6.85	7.25	7.05	7.61	6.80	7.07	6.85	6.85
	5	6.42	6.32	6.59	6.98	6.73	6.69	6.48	6.73	6.50	6.48
	10	6.19	6.12	6.24	6.67	6.45	6.23	6.28	6.29	6.28	6.26
Panel C: Stocks and bonds											
Emp. dist.	1	9.24	10.66	7.11	10.04	9.17	8.92	10.27	9.16	9.16	9.19
	5	6.91	6.50	4.02	6.17	6.53	6.29	7.03	5.30	6.72	4.72
	10	4.99	4.06	3.37	2.94	4.89	3.57	2.31	3.48	3.44	3.57
Logn. dist.	1	9.00	8.31	10.27	8.59	9.29	8.41	10.18	7.66	10.21	9.07
	5	6.67	5.63	6.93	5.50	6.99	7.50	7.75	5.08	8.48	6.56
	10	3.83	4.29	5.08	4.15	5.50	5.77	4.44	4.73	6.34	5.03

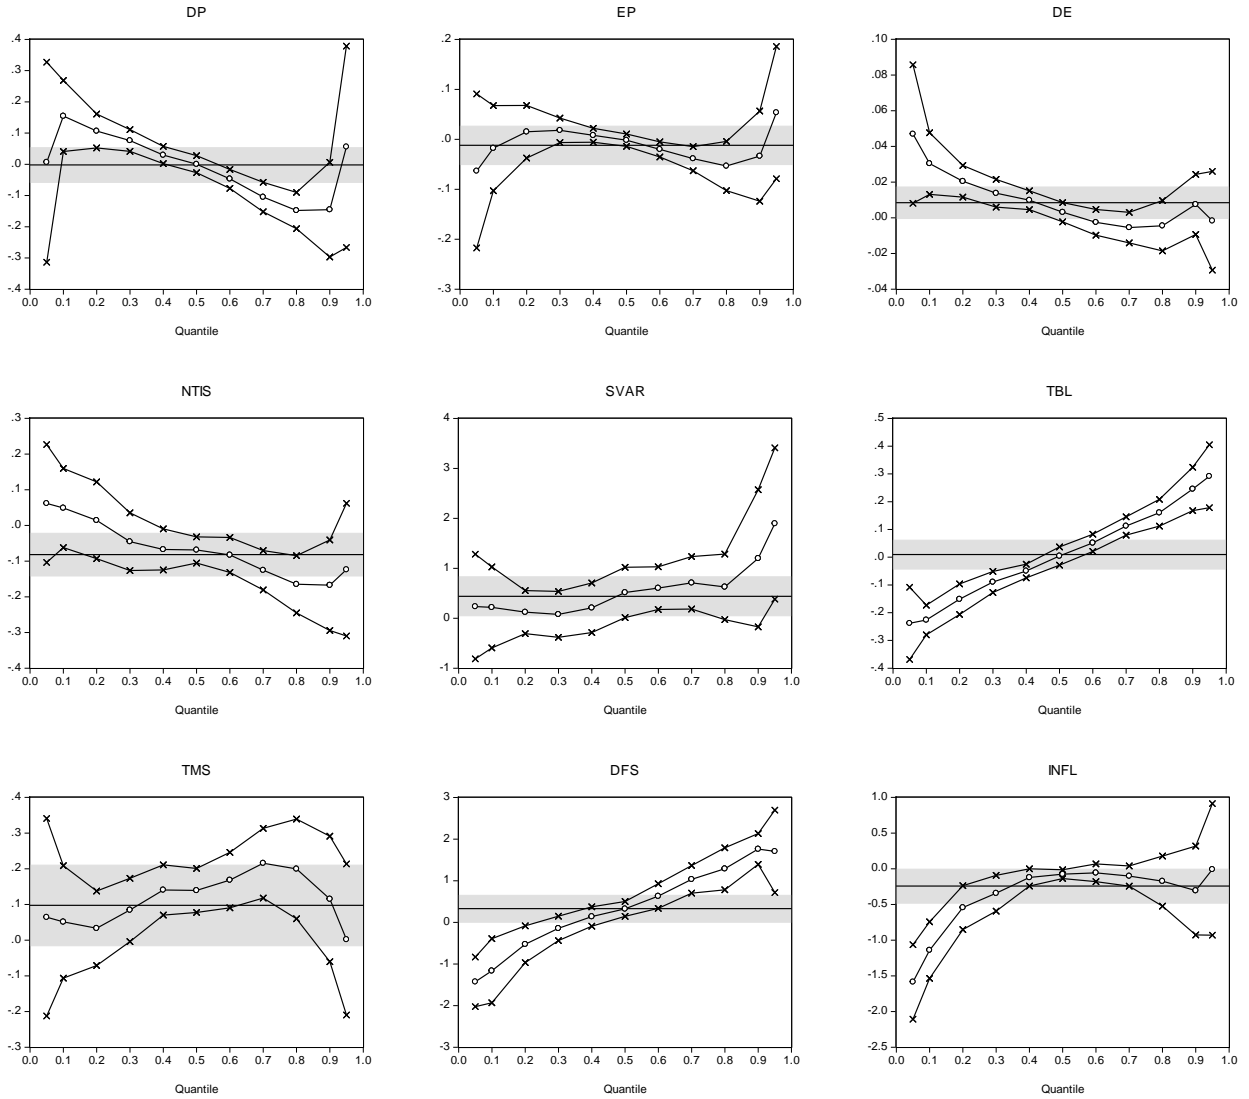
Notes: The table gives the per annum certainty equivalent return (in per cent) for an investor with power utility and coefficient of relative risk aversion γ . In Panel A the investor can only invest in stocks and a risk free asset, in Panel B the investor can only invest in bonds and a risk free asset, and in Panel C the investor can invest in both stocks and bonds as well as a risk free asset. Each panel shows the results when quantile regression is used to obtain an estimate of the distribution (Emp. dist.) and when returns are assumed to be lognormally distributed from the outset (Logn. dist.). For each state variable the highest certainty equivalent return across Emp. dist. and Logn. dist. is boldfaced.

Figure 1. Slope coefficient estimates for stock returns in the univariate setting.



Notes: The solid horizontal line gives the least squares slope coefficient with the shaded area indicating the 95% confidence interval. The line with circles gives the slope coefficients from quantile regression, with the lines with crosses indicating the 95% confidence interval.

Figure 2. Slope coefficient estimates for bond returns in the univariate setting.




Notes: The solid horizontal line gives the least squares slope coefficient with the shaded area indicating the 95% confidence interval. The line with circles gives the slope coefficients from quantile regression, with the lines with crosses indicating the 95% confidence interval.

Figure 3. Joint distribution between stocks and bonds.

		0.004 (0.003)	0 (0.023)	0.005 (0.023)	0.009 (0.003)
α_b	0.95	0.186 (0.023)	0.011 (0.203)	0.059 (0.203)	0.182 (0.023)
	0.50	0.002 (0.023)	0.263 (0.203)	0.254 (0.203)	0.003 (0.023)
	0.05	0.007 (0.003)	0.003 (0.023)	0.012 (0.023)	0 (0.003)
		0.05	0.50	0.95	
			α_s		

Notes: Each quadrant gives the in-sample coverage probability for the relevant quantile combination based on the prevailing quantile model. Numbers in parentheses give the coverage probability given independence.

Figure 4. Slope coefficient estimates for stock returns as a function of stock variance.



Notes: The solid horizontal line gives the univariately estimated 0.95-quantile slope coefficient for stock returns as a function of stock variance. The line with circles gives the corresponding estimates in the multivariate setup for different values of the bond quantile.

Research Papers 2010



- 2010-24: Søren Johansen and Morten Ørregaard Nielsen: Likelihood inference for a fractionally cointegrated vector autoregressive model
- 2010-25: Tom Engsted and Bent Nielsen: Testing for rational bubbles in a co-explosive vector autoregression
- 2010-26: Robinson Kruse: On European monetary integration and the persistence of real effective exchange rates
- 2010-27: Sanne Hiller and Robinson Kruse: Milestones of European Integration: Which matters most for Export Openness?
- 2010-28: Robinson Kruse: Forecasting autoregressive time series under changing persistence
- 2010-29: Nikolaus Hautsch and Mark Podolskij: Pre-Averaging Based Estimation of Quadratic Variation in the Presence of Noise and Jumps: Theory, Implementation, and Empirical Evidence
- 2010-30: Martin M. Andreasen: Non-linear DSGE Models and The Central Difference Kalman Filter
- 2010-31: Morten Ørregaard Nielsen and Per Frederiksen: Fully Modified Narrow-Band Least Squares Estimation of Weak Fractional Cointegration
- 2010-32: Mogens Bladt and Michael Sørensen: Simple simulation of diffusion bridges with application to likelihood inference for diffusions
- 2010-33: Fernando Baltazar-Larios and Michael Sørensen: Maximum likelihood estimation for integrated diffusion processes
- 2010-34: Leonidas Tsiaras: The Forecast Performance of Competing Implied Volatility Measures: The Case of Individual Stocks
- 2010-35: Leonidas Tsiaras: Dynamic Models of Exchange Rate Dependence Using Option Prices and Historical Returns
- 2010-36: Robinson Kruse and Rickard Sandberg: Linearity Testing in Time-Varying Smooth Transition Autoregressive Models under Unknown Degree of Persistency
- 2010-37: Tom Engsted and Thomas Q. Pedersen: The log-linear return approximation, bubbles, and predictability
- 2010-38: Thomas Q. Pedersen: Predictable return distributions