Essays in Applied Econometrics of High Frequency Financial Data

Ilya Archakov

Thesis submitted for assessment with a view to obtaining the degree of Doctor of Economics of the European University Institute

Florence, 19 December, 2016

European University Institute
Department of Economics

Essays in Applied Econometrics of High Frequency Financial Data

Ilya Archakov

Thesis submitted for assessment with a view to obtaining the degree of
Doctor of Economics of the European University Institute

Examining Board

Prof. Peter Reinhard Hansen, Supervisor, University of North Carolina
Prof. Juan José Dolado, EUI
Prof. Christian Brownlees, Universitat Pompeu Fabra
Prof. Asger Lunde, Aarhus University

European University Institute

DEPARTMENT OF ECONOMICS

Researcher declaration to accompany the submission of written work

I Ilya Archakov certify that I am the author of the work *Essays in Applied Econometrics of High Frequency Financial Data* I have presented for examination for the PhD thesis at the European University Institute. I also certify that this is solely my own original work, other than where I have clearly indicated, in this declaration and in the thesis, that it is the work of others.

I warrant that I have obtained all the permissions required for using any material from other copyrighted publications.

I certify that this work complies with the *Code of Ethics in Academic Research* issued by the European University Institute (IUE 332/2/10 (CA 297).

Statement of inclusion of previous work (if applicable):

I confirm that chapter <1> was jointly co-authored with Peter Hansen and Asger Lunde and I contributed 50% of the work.

I confirm that chapter <2> was jointly co-authored with Bo Laursen and I contributed 50% of the work.

I confirm that chapter <4> was jointly co-authored with Peter Hansen, Guillaume Horel and Asger Lunde and I contributed 25% of the work.

I confirm that chapter <4> draws upon a book chapter appeared in *The Fascination of Probability, Statistics and their Applications: In Honour of Ole E. Barndorff-Nielsen*

Signature and Date:
Ilya Archakov, 01/12/2016

Abstract

In the first chapter, co-authored with Peter Hansen and Asger Lunde, we suggest a novel approach to modeling and measuring systematic risk in equity markets. We develop a new modeling framework that treats an asset return as a dependent variable in a multiple regression model. The GARCH-type dynamics of conditional variances and correlations between the regression variables naturally imply a temporal variation of regression coefficients (betas). The model incorporates extra information from the realized (co-)variance measures extracted from high frequency data, which helps to better identify the latent covariance process and capture its changes more promptly. The suggested structure is consistent with the broad class of linear factor models in the asset pricing literature.

We apply our framework to the famous three-factor Fama-French model at the daily frequency. Throughout the empirical analysis, we consider more than 800 individual stocks as well as style and sectoral exchange traded funds from the U.S. equity market. We document an appreciable cross-sectional and temporal variation of the model-implied risk loadings with the especially strong (though short-lived) distortion around the Financial Crisis episode. In addition, we find a significant heterogeneity in a relative explanatory power of the Fama-French factors across the different sectors of economy and detect a fluctuation of the risk premia estimates over time. The empirical evidence emphasizes the importance of taking into account dynamic aspects of the underlying covariance structure in asset pricing models.

In the second chapter, written with Bo Laursen, we extend the popular dynamic Nelson-Siegel framework by introducing time-varying volatilities in the factor dynamics and incorporating the realized measures to improve the identification of the latent volatility state. The new model is able to effectively describe the conditional distribution dynamics of a term structure variable and can still be readily estimated with the Kalman filter.

We apply our framework to model the crude oil futures prices. Using more than 150,000,000 transactions for the large panel of contracts we carefully construct the realized volatility measures corresponding to the latent Nelson-Siegel factors, estimate the model at daily frequency and evaluate it by forecasting the conditional density of futures prices. We document that the time-varying volatility specification suggested in our model strongly outperforms the constant volatility benchmark. In addition, the use of realized measures provides moderate, but systematic gains in density forecasting.

In the third chapter, I investigate the rate at which information about the daily asset volatility level arrives with the transaction data in the course of the trading day. The contribution of this analysis is three-fold. First, I gauge how fast (after the market opening) the reasonable projection of the new daily volatility level can be constructed. Second, the framework provides a natural experimental field for the comparison of the small sample properties of different types of estimators as well as their (very) short-run forecasting capability. Finally, I outline an adaptive modeling framework for volatility dynamics that attaches time-varying weights to the different predictive signals in response to the changing stochastic environment.

In the empirical analysis, I consider a sample of assets from the Dow Jones index. I find that the average precision of the ex-post daily volatility projections made after only 15 minutes of trading (at 9:45a.m. EST) amounts to 65% (in terms of predictive R^2) and reaches up to 90% before noon. Moreover, in conjunction with the prior forecast, the first 15 minutes of trading are able to predict about 80% of the ex-post daily volatility. I document that the predictive content of the realized measures that use data at the transaction frequency is strongly superior as compared to the estimators that use sparsely sampled data, but the difference is getting

negligible closer to the end of the trading day, as more observations are used to construct a projection.

In the final chapter, joint with Peter Hansen, Guillaume Horel and Asger Lunde, we introduce a multivariate estimator of financial volatility that is based on the theory of Markov chains. The Markov chain framework takes advantage of the discreteness of high-frequency returns and suggests a natural decomposition of the observed price process into a martingale and a stationary components. The new estimator is robust to microstructural noise effects and is positive semidefinite by construction. We outline an approach to the estimation of high dimensional covariance matrices. This approach overcomes the curse of dimensionality caused by the tremendous number of observed price transitions (normally, exceeding 10,000 per trading day) that complicates a reliable estimation of the transition probability matrix for the multivariate Markov chain process.

We study the finite sample properties of the estimator in a simulation study and apply it to high-frequency commodity prices. We find that the new estimator demonstrates a decent finite sample precision. The empirical estimates are largely in agreement with the benchmarks, but the Markov chain estimator is found to be particularly well with regards to estimating correlations.

Acknowledgements

I'm deeply indebted to my supervisor Peter Reinhard Hansen for his phenomenal support and guidance during all these years. I'm grateful to Massimiliano Marcellino and Juan Dolado for their constant encouragement and helpful advice. I would like to thank Asger Lunde for the invaluable collaboration and feedback. WIthout it this thesis would hardly have been possible. A special gratitude goes to Roberto Renó and Christian Brownlees for their passion and curiosity as well as for many insightful discussions.

I'm grateful to each and everyone with whom I was lucky to study and work together. I wish to especially thank Pavel, Moritz, Reinhard, Gabriela, and Bo for all the numerous enjoyable moments and exciting collaborations.

I am very thankful to the exceptional EUI team for their permanent assistance. I particularly thank Jessica Spataro, Julia Valerio, Marcia Gastaldo, Lucia Vigna, Françoise Thauvin, Thomas Bourke, Riccardo Mongiu, Anne Banks, Sarah Simonsen, Sonia Sirigu and Loredana Nunni for their help and kindness.

I would like to acknowledge the financial support by the Italian Ministry of Foreign Affairs. I wouldn't be able to complete the thesis without it.

Finally, I thank my family for an outstanding support. I extend my deepest gratitude to Maria for her incredible patience and love that helped me travel through this journey.

Contents

Chapter 1

A Factor Model with Realized Measures

Joint with Peter Hansen and Asger Lunde

1.1 Introduction

The economic environment is constantly changing. Among other things, it is reflected in time-varying business conditions and permanently fluctuating interdependencies among the single companies as well as the whole sectors of the economy. This, in turn, manifests in changing volatilities and correlations between financial asset returns over time. Ignoring such variation might lead to inaccurate or wrong implications in macroeconomic and finance research. This matter is also of critical significance for practitioners since the covariance between financial securities is in the heart of a portfolio allocation, risk management, and hedging. An availability of high frequency data gives a promise of a more accurate measurement and estimation of unobservable covariation variables and suggests new frontiers for econometric modeling.

In this paper we develop a flexible framework that properly accounts for temporal changes in the latent covariance structure of multivariate asset returns, incorporates realized measures of volatilities and correlations based on high frequency data, nests several modeling approaches, simple in estimation, and can be useful for a range of research and practical applications. More particularly, the design allows to carefully model the conditional second moments of multivariate and univariate returns, simplifies the treatment of high dimensional return vectors and is consistent with the broad class of linear factor models in asset pricing.

Our framework is closely related to the Realized GARCH model (Hansen et al. (2012)) and extends it along the several directions. Similarly, together with asset returns we explicitly model the dynamics of realized measures of variances and correlations which timely provide us with precise signals about the latent conditional covariance. The central element of the modeling strategy is the assumption of a multiple regression structure for an asset return which implies that a given financial asset has an exposure to a set of regressors (factors) and to an idiosyncratic component. Importantly, the model implies dynamics for the whole covariance structure, i.e. both within the set of regressors and between the regressors and a dependent variable.

We follow an hierarchical modeling framework suggested in Hansen et al. (2014b). At first, we build the marginal model, or the core model, for the set of regressors that represents a newly developed multivariate realized GARCH model. Then, conditionally on the variables filtered from the core part, we set the models for individual asset returns handling them as dependent regression variables. When the regressors are common factors for a variety of assets, the core model is relevant for the conditional dynamics of many individual returns. An assumption of a factor structure together with an hierarchical modeling makes it feasible to deal with conditional covariance matrices of potentially large dimensions as long as we suppose that the cross-correlations are predominantly explained by a handful of common factors.

Furthermore, the suggested framework is perfectly adapted for an empirical analysis of asset pricing linear factor models if regressors are treated as factors of systematic risk. The rich dynamics of the conditional covariance structure implied by the framework induces a temporal variation in regression coefficients, or factor loadings (betas). As a consequence, the use of high frequency data is helpful not only to identify the unobservable covariation components, but also directly leads to potential efficiency gains in estimating conditional factor loadings and risk premia.

From the methodological side, the paper is related to several branches of the econometric literature. The core model for multiple factors contributes to the large class of multivariate GARCH models[1]. Similar to the DCC model (Engle (2002)) or the model in Tse and Tsui (2002), we split the covariance structure into the variances and correlations and specify separately the corresponding conditional dynamics. In contrast to the DCC, instead of modeling the conditional correlation matrix directly we work with the vector of elements from its logarithmic transform. Such strategy is attractive because it renders the filtered correlation matrix valid, preserving positive definiteness at each step. In addition, all correlation elements are endowed with their own non-trivial dynamics, unlike in the CCC (Bollerslev (1990)) or the DECO (Engle and Kelly (2011)) models. However, the most important feature that differs our modeling approach from most of models in the literature is the use of realized measures of volatilities and correlations. Thus, signals from both low frequency returns and high frequency observations are incorporated into the model in order to better capture the dynamics of unobservable covariance.

Modeling of dynamic variances and covariances with realized measures became a popular research direction after the works of Andersen and Bollerslev (1998a), Andersen et al. (2001b), Barndorff-Nielsen and Shephard (2002b), Andersen et al. (2003a), Barndorff-Nielsen and Shephard (2004), etc., where the theoretical grounds for the realized variance estimator and the practical benefits from its applications were establshed[2]. The literature suggests several dominant approaches to modeling the multivariate volatility with realized measures. They include purely reduced-form models where the realized covariance elements are described as time series processes (e.g., Liu (2009), Chiriac and Voev (2011)), semi-reduced-form models where the reduced-form time series of realized covariances are supplemented by measurement equations (see Golosnoy et al. (2012), Bauwens et al. (2012)), and the models where realized covariance measures are modeled jointly with the multivariate return vector (Noureldin et al. (2012)). In the realized GARCH paradigm we not only model return vector together with realized measures, but also incorporate the measurement equations which complete the model by connecting observed realized and latent conditional covariances. Note that an implicit advantage of the measurement equations relates to the increased stability of model estimation due to an additional likelihood component specified for the realized measures.

The modeling framework developed in the paper extends the growing family of the realized GARCH models. First, the core part of the model can be viewed as a direct multivariate generalization of Hansen et al. (2012) and Hansen and Huang (2016). Apart from the alternative multivariate Realized GARCH models proposed in Balter et al. (2015) and Hansen et al. (2016), we ensure positive definiteness of the covariance process and disentangle the volatility dynamics from the correlation dynamics, thus making the model formulation more convenient and tractable. Second, the conditional model for individual returns represents an extension of the Realized Beta GARCH framework (Hansen et al. (2014b)) for the case of multiple conditioning factors. Finally, our model can incorporate the leverage effects, spillover effects for volatilities and correlations, as well as the heavy-tailed distributional aspects (see Banulescu et al. (2014)). This endows the conditional return density with flexible dynamics and makes a promise of a decent empirical fit.

An appealing feature of the developed framework relates to its interpretation as a linear factor model. In particular, the possibility to readily extract dynamic factor loadings as a function of conditional covariation variables suggests a promising contribution to the literature on the empirical evaluation of dynamic asset pricing models. The factor loadings, or the factor betas, are the benchmark indicators of non-diversifiable systematic

[1]For more references, see surveys in Laurent et al. (2006) and Silvennoinen and Terasvirta (2008).

[2]For more references on the estimation of realized volatility measures, modeling, and forecasting see reviews in McAleer and Medeiros (2008), Andersen et al. (2010), Hansen and Lunde (2011).

risk for a given asset or portfolio. A number of studies documented empirical evidence in favour of dynamic nature of betas in the asset pricing factor models (see Bollerslev et al. (1988), Ferson and Harvey (1991), Jagannathan and Wang (1996), and others). Thus, an ability to accurately and promptly estimate local factor loadings may improve an assessment of asset pricing models and their cross-section pricing implications.

The factor loadings implied by the specification of our framework are intrinsically time-varying since they depend on the dynamic conditional variances and correlations incorporated into the model. This is unlike to the rolling regression approach commonly used in the literature. Although the rolling regressions allow a certain degree of time variation in betas, they are not capable to capture most of the dynamic features that are possibly embedded in factor loadings, especially at higher frequencies. The problem of imprecise beta estimation when long-window rolling regressions are used has often been pointed out (see Fama and French (1997), Lewellen and Nagel (2006), etc.).

The dynamic betas produced within our framework are expected to have an advanced precision. On the one hand, since our model enjoys the superior information content from high frequency realized measures, the model-implied betas are more efficient and supposed to more rapidly adjust with the changes in the true unobservable loadings. As a consequence, we may presume an advantage over the dynamic betas estimated by more standard multivariate GARCH frameworks where only the information from low frequency returns is exploited (Bekaert and Wu (2000), Engle (2012)). On the other hand, as opposed to the reduced form models (e.g., Bollerslev and Zhang (2003), Andersen et al. (2006)) for realized betas (Barndorff-Nielsen and Shephard (2004), Bandi and Russell (2005)), our approach is based on the multivariate GARCH structure for daily returns. Therefore, conditional betas are primarily driven by the dynamics of variances and correlations between actual returns and risk factors, rather than by the corresponding realized measures solely.

In order to examine our novel modeling framework empirically we apply it to the celebrated Fama-French three factor model (Fama and French (1992), Fama and French (1993)) at a daily frequency. At first, we construct intra-daily Fama-French factors using high frequency data on the style exchange traded funds (ETFs henceforth), which represent portfolios of assets selected according to the different investment styles. The constructed intra-daily factor series allow us to obtain realized measures of variances and correlations between the factors. At second, we estimate the core model for the three factors and then estimate conditional Fama-French factor models for different test assets. For this purpose we use large cross-section of more than 800 stocks from the U.S. equity market and a range of ETFs that reflect different investment styles and sectoral indexes. At third, we produce a descriptive empirical analysis of the model-implied dynamic betas, implement an approximate decomposition of asset return variances along the risk factors and, finally, assess the pricing implications of the Fama-French factors through the prism of our modeling framework.

We found that although the betas demonstrate an appreciable variability in dynamics, such variability is highly heterogeneous across the sectors of the U.S. economy. Our findings reveal a large extent of temporal variation in the cross-section distributions of factor loadings, especially for the book-to-market factor. We, thus, corroborate the results from Hansen et al. (2014b) obtained for the market betas, but now controlling for the two additional risk factors. We document that the variance decomposition of stock and ETF returns, implied by the dynamic Fama-French model, significantly changes over time. In particular, we found that the value (book-to-market) factor is reflected mostly in the return variances of the financial companies around the Great Recession and in the aftermath. Despite that the cross-section risk premia estimates, computed with the model-implied factor loadings, are very noisy at high frequencies, the results suggest an evidence of an appreciable time variation in the premia for all three risk factors.

Our empirical findings confirm an ability of the modeling framework to evaluate multi-factor models for asset returns. It motivates to investigate further both statistical implications and the economic significance of the model-implied variables. In particular, it might be interesting and important to examine the model performance in estimation and forecasting large-scale dynamic covariance matrices given a proper choice of the common factors, to assess risk-return benefits from the portfolio selection strategies based on the filtered and predicted risk factor loadings, evaluate the implications of different asset pricing models on the conditional density of returns, etc. These routes are the subject of the ongoing and future research.

We proceed as follows. In Section 1.2 we present the framework and discuss the estimation methodology. In Section 1.3 we describe the data, provide the details on the construction of realized measures and give a short discourse of the model estimation results for the Fama-French factors. We discuss the properties of the filtered dynamic betas and provide a variance decomposition analysis in Section 1.4. In Section 1.5 we refer to asset pricing implications of our modeling framework. Section 1.6 concludes and provides an extended outline. The detailed estimation algorithm, the model evaluation, and the thorough description of the ETF selection for the Fama-French factors replication can be found in Appendices.

1.2 Realized GARCH for Multiple Factor Models

In this section, we describe the modeling strategy in details. We begin with the brief description of the general modeling assumptions and introduce a notation. Then, we proceed with the multivariate core model for the factors. After, we present the conditional factor model for individual asset returns. Finally, we discuss the assumptions that are needed to apply our framework to a large system of individual assets and outline the estimation strategy.

1.2.1 Notation and Preliminaries

We cast the model in a discrete time and index t stands for a generic time unit which is typically a trading day. We also suppose that there are many individual assets in the asset universe, $i = 1, ..., N$. Let r_t^i denotes a return on an individual asset i and $\mathbf{r}_t^c$ denotes a $k \times 1$ vector of returns associated with k factors. Additionally, we observe realized covariance measures for the vector of factor returns $\mathbf{r}_t^c$. The information contained in a realized covariance matrix may be represented through $k \times 1$ vector of realized factor variances $\mathbf{x}_t^c$ and the $k \times k$ realized correlation matrix Y_t^c. Also, for each asset i we possess realized variance measures of individual asset returns, x_t^i, as well as $k \times 1$ vectors $\mathbf{y}_t^i$ of realized correlations between asset and factor returns, r_t^i and $\mathbf{r}_t^c$. Naturally, all realized variances are strictly positive and all pairwise realized correlations lie in the open interval $(-1, 1)$.

We introduce two types of information sets. The first type is a tuple which contains all observable variables related solely to the multivariate factor process, $\mathcal{X}_t^c = \left(\mathbf{r}_t^c, \mathbf{x}_t^c, Y_t^c\right)'$. The second information tuple consists of observables associated with an individual asset, $\mathcal{X}_t^i = \left(r_t^i, x_t^i, \mathbf{y}_t^i\right)'$, and is available for all $i = 1, ..., N$. We also introduce two types of natural filtrations. The first is generated only by the information sets $\mathcal{X}_t^c$ and can be written as the following σ-algebra

$$\mathcal{F}_t^c = \sigma\left(\mathcal{X}_t^c, \mathcal{X}_{t-1}^c, ...\right)$$

The second type is based on the richer information where the variables related to an asset i are also incorporated,

$$\mathcal{F}_t^{c,i} = \sigma\left(\mathcal{X}_t^c, \mathcal{X}_t^i, \mathcal{X}_{t-1}^c, \mathcal{X}_{t-1}^i, ...\right)$$

In accordance with the Realized GARCH framework, we model a conditional joint density of observable returns and realized measures, $p(\mathcal{X}_t^c, \mathcal{X}_t^i | \mathcal{F}_{t-1}^{c,i})$. We follow an hierarchical modeling approach previously suggested in Hansen et al. (2014b). This approach is based on the assumption that the conditional joint density of the observable variables related to the factors, $p(\mathcal{X}_t^c | \mathcal{F}_{t-1}^{c,i})$, does not depend on observations related to any individual asset i[3], that is

$$p(\mathcal{X}_t^c | \mathcal{F}_{t-1}^{c,i}) = p(\mathcal{X}_t^c | \mathcal{F}_{t-1}^c) \tag{1.1}$$

As a consequence, we may factorize the joint conditional density as follows

[3]This assumption is natural, especially in the current context of a factor model. Whereas an individual asset returns are supposed to depend on factors, the opposite is usually not assumed.

$$p(\mathcal{X}_t^c, \mathcal{X}_t^i | \mathcal{F}_{t-1}^{c,i}) = p(\mathcal{X}_t^c | \mathcal{F}_{t-1}^c) p(\mathcal{X}_t^i | \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}) \tag{1.2}$$

Since the conditional distribution of asset variables depends on factor variables, but not vice versa, the modeling strategy for $p(\mathcal{X}_t^c, \mathcal{X}_t^i | \mathcal{F}_{t-1}^{c,i})$ naturally consists of two steps. At first, the model for $p(\mathcal{X}_t^c | \mathcal{F}_{t-1}^c)$ is formulated. Note, that this model can be set up independently from individual assets due to (1.1). We will refer to this model as to the core model, since it is on the upper level of the hierarchical framework. In the second step, the dynamics for $p(\mathcal{X}_t^i | \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i})$ is specified conditionally on the estimated core model. The corresponding models (on the lower level of the modeling hierarchy) will represent conditional factor models for individual assets $i = 1, ..., N$.

1.2.2 Multivariate Realized GARCH Model for Factors

We begin with the core model for the factors and their realized measures of covariance. As we have noted before, this model is associated with the conditional density $p(\mathcal{X}_t^c | \mathcal{F}_{t-1}^c)$. We now introduce the version of the multivariate Realized GARCH model where the dynamics of conditional distribution of the return vector is modeled together with the realized measures of its variances and correlations.

1.2.2.1 Conditional Distribution of Returns

We suppose that the $k \times 1$ vector of factors is conditionally distributed as a multivariate Normal random vector[4]

$$\mathbf{r}_t^c | \mathcal{F}_{t-1}^c \sim N\left(\mu^c, H_t^c R_t^c H_t^c\right) \tag{1.3}$$

where μ^c is a $k \times 1$ constant vector, H_t^c is an unobservable diagonal matrix of conditional volatilities, and R_t^c is a latent conditional correlation matrix of $\mathbf{r}_t^c$:

$$H_t^c = \begin{pmatrix} \sqrt{h_{1,t}^c} & 0 & \cdot & 0 \\ 0 & \sqrt{h_{2,t}^c} & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & \cdot & \sqrt{h_{k,t}^c} \end{pmatrix} \quad \text{and} \quad R_t^c = \begin{pmatrix} 1 & \diamond & \cdot & \diamond \\ \rho_{21,t}^c & 1 & \cdot & \diamond \\ \cdot & \cdot & \cdot & \cdot \\ \rho_{k1,t}^c & \rho_{k2,t}^c & \cdot & 1 \end{pmatrix}$$

As long as we assume that R_t^c is an invertible matrix, we may rewrite (1.3) in the equation form

$$\mathbf{r}_t^c = \mu^c + H_t^c \mathbf{z}_t^c \tag{1.4}$$

where $\mathbf{z}_t^c$ is a $k \times 1$ standardized factor return distributed as a multivariate Normal, $\mathbf{z}_t^c \sim N(0, R_t^c)$. In our model, we follow the principles of the Dynamic Conditional Correlation (DCC) framework (see Engle (2002), Engle and Sheppard (2001), etc.) in sense that we disentangle the dynamics of conditional variances from the dynamics of conditional correlations. In contrast to the DCC, we incorporate realized measures of variances and correlations into the conditional covariance process.

1.2.2.2 Conditional Variances

At first, lets describe the conditional dynamics of variances. For expositional convenience we introduce a $k \times 1$ vector of conditional variances $\mathbf{h}_t^c = (h_{1,t}^c, h_{2,t}^c, ..., h_{k,t}^c)'$. The realized analogue of this vector is the $k \times 1$ vector $\mathbf{x}_t^c = (x_{1,t}^c, x_{2,t}^c, ..., x_{k,t}^c)'$ consisted of realized variances of k factors obtained using the high frequency data. In the spirit of the Realized GARCH model, we formulate a measurement equation to establish a contemporaneous relation between the conditional variances and their realized measures

[4]Note that the distributional choice may be different. We use the Normal assumption and rely on the QML argument in case of a possible misspecification.

$$\log \mathbf{x}_t^c = \Phi_0^c + \Phi_1^c \log \mathbf{h}_t^c + \delta_c(\mathbf{z}_t^c) + \mathbf{u}_t^c \tag{1.5}$$

where Φ_0^c is a $k \times 1$ vector and Φ_1^c is a $k \times k$ matrix of constant coefficients, which accounts for a bias in realized measures[5]. A random vector $\mathbf{u}_t^c \sim N\left(0, \Omega_u^c\right)$ introduces a stochastic measurement error[6]. Note that measurement errors of realized variances corresponding to different factor returns may be correlated and, therefore, covariance matrix Ω_u is generically non-diagonal. The measurement error $\mathbf{u}_t^c$ is assumed to be independent from $\mathbf{z}_t^c$. Although this assumption is quite restrictive, the presence of the leverage component $\delta_c(\mathbf{z}_t^c)$ provides some compensation for that. Following Hansen et al. (2012), we introduce a parametric leverage function which reads

$$\delta_c(\mathbf{z}_t^c) = D_1^c \mathbf{z}_t^c + D_2^c(\mathbf{z}_t^c \circ \mathbf{z}_t^c - \iota_k)$$

where ι_k stands for the $k \times 1$ unit vector, D_1^c and D_2^c are $k \times k$ coefficient matrices. The leverage function defined as above represents a multivariate version of the leverage function introduced in Hansen et al. (2012). This function captures an asymmetry in the relation between returns and volatility dynamics and is motivated by the stylized empirical facts[7]. Hansen et al. (2012) found that a parsimonious leverage function written as a second-order Hermite polynomial makes a good job in capturing the aforementioned asymmetry and appreciably improves the model fit. If we set the matrices D_1^c and D_2^c to be diagonal, then, as it follows from (1.5), we come up with k independent asymmetry effects between k factors and k corresponding conditional variances. In this case, each element of vector $\delta_c(\mathbf{z}_t^c)$ gets the same functional form as a univariate leverage function from Hansen et al. (2012). In contrast, if D_1^c or D_2^c are non-diagonal, the leverage spillovers from some factors to the volatilities of other factors are incorporated.

We now may formulate the GARCH-type dynamic equation for the conditional variances

$$\log \mathbf{h}_t^c = \Gamma_0^c + \Gamma_1^c \log \mathbf{h}_{t-1}^c + \Gamma_2^c \log \mathbf{x}_{t-1}^c + \tau_c(\mathbf{z}_{t-1}^c) \tag{1.6}$$

where Γ_0^c is a constant $k \times 1$ vector, Γ_1^c and Γ_2^c are $k \times k$ coefficient matrices, and, finally, $\tau_c(\cdot)$ is an additional leverage function which has the same form as in the previous case

$$\tau_c(\mathbf{z}_t^c) = T_1^c \mathbf{z}_t^c + T_2^c(\mathbf{z}_t^c \circ \mathbf{z}_t^c - \iota_k)$$

with T_1^c and T_2^c being $k \times k$ coefficient matrices. The equation (1.6) is the Realized EGARCH dynamic equation from Hansen and Huang (2016) written in a vector form. A key feature is the presence of a realized variance measure, which improves the reaction of the modeled conditional volatility process to the latent changes in the volatility level. Another important aspect is the presence of an additional leverage component $\tau_c(\cdot)$, which adds an extra flexibility in modeling the dependence between return shocks and volatility shocks and was proven to be beneficial for the empirical model performance. In both the measurement equation (1.5) and in the transition equation (1.6), we use vectors with log-transformed entries. This choice is in line with the EGARCH modeling framework and ensures that the elements from $\mathbf{h}_t^c$ always stay positive[8].

[5] There are several reasons that may cause a bias in the realized measures of volatility. One of the main reasons is that realized measures do not account for an overnight variation since they are computed using intra-daily returns observed only during the trading hours. Therefore, if one takes $\mathbf{r}_t^c$ to be close-to-close returns, realized measures will consistently underestimate the return variance. Other reasons include a small sample bias of particular high frequency realized estimators, the presence of jumps and outliers, etc.

[6] Another interpretation of $\mathbf{u}_t^c$ is a volatility shock (see Hansen and Huang (2016)). Since $\mathbf{h}_t^c$ can be treated as a forecast of the latent variance that is made at $t-1$, the difference between the realized variance $\mathbf{x}_t^c$ and $\mathbf{h}_t^c$ represents a forecast error, or variance innovation with respect to the predicted volatility level.

[7] Conventionally, the leverage effect is accociated with a (usually negative) correlation between the return volatility and the return shocks, or news. See, e.g., Black (1976), Christie (1982), Engle and Ng (1993).

[8] In addition, the use of log-transformed variances in the measurement equation (1.5) can be motivated by the empirical studies of Andersen et al. (2001a) and Barndorff-Nielsen and Shephard (2002c), who found that the logarithm of the realized variance can be well approximated by the Normal distribution.

1.2.2.3 Conditional Correlations

We now turn to the conditional dynamics of correlations. A full-rank real valued matrix R^c needs to satisfy 2 properties in order to be a correlation matrix: it has to be positive definite and each element from the main diagonal must be equal to one. These restrictions on the matrix structure complicate modeling of the conditional correlation dynamics[9]. In contrast to the standard DCC approach, we will not formulate a matrix process, but suggest a vector process for the elements of the transformed conditional correlation matrix.

We base our modeling strategy on the matrix logarithmic and exponential transformations. Applying the spectral decomposition to a positive definite matrix A we get $A = V\Lambda V'$, where V is an orthonormal matrix of eigenvectors and Λ is a diagonal matrix with the eigenvalues of A on the main diagonal. Then, log-matrix transform of A implies $\text{logm}(A) = V\log(\Lambda)V'$, where $\log(\Lambda)$ applies logarithmic transformations to all diagonal elements of the matrix Λ. The matrix exponential represents an inverse operation and reads $\text{expm}(A) = V\exp(\Lambda)V'$, where $\exp(\Lambda)$ takes the exponent of all diagonal elements in Λ.

The useful property of log/exp-matrix operations is that they provide a simple mapping from the space of positive definite matrices to the space of symmetric matrices and back. This property is especially useful when is applied to modeling matrices which have to be positive definite, such as covariance matrices[10]. In what follows, we will also exploit this feature, but in the context of modeling correlation matrices and their realized measures. To begin with, we provide a lemma

Lemma 1. *Suppose that R is a non-singular correlation matrix. Then, the diagonal elements of $\text{logm}(R)$ are bounded above, but the off-diagonal elements are unbounded on $\mathbb{R}$.*

Lets denote the log-transformed matrices of conditional and realized correlations as

$$LR_t^c = \text{logm}(R_t^c) = \begin{pmatrix} l\rho_{11,t}^c & \diamond & \cdot & \diamond \\ l\rho_{21,t}^c & l\rho_{22,t}^c & \cdot & \diamond \\ \cdot & \cdot & \cdot & \cdot \\ l\rho_{k1,t}^c & l\rho_{k2,t}^c & \cdot & l\rho_{kk,t}^c \end{pmatrix} \quad \text{and} \quad LY_t^c = \text{logm}(Y_t^c) = \begin{pmatrix} ly_{11,t}^c & \diamond & \cdot & \diamond \\ ly_{21,t}^c & ly_{22,t}^c & \cdot & \diamond \\ \cdot & \cdot & \cdot & \cdot \\ ly_{k1,t}^c & ly_{k2,t}^c & \cdot & ly_{kk,t}^c \end{pmatrix}$$

By $l\rho_t^c = (l\rho_{21,t}^c, l\rho_{31,t}^c, l\rho_{32,t}^c, ..., l\rho_{kk-1}^c)' = \text{vech}_{od}LR_t^c$ we denote the $\frac{1}{2}k(k-1)\times 1$ vector with the stacked distinct off-diagonal elements from LR_t^c. The realized analogue of $l\rho_t^c$ is denoted by $l\mathbf{y}_t^c = ((y_{21,t}^c, ly_{31,t}^c, ly_{32,t}^c, ..., ly_{kk-1}^c)' = \text{vech}_{od}LY_t^c$. We also make the following conjecture

Conjecture 1. *For any vector $x \in \mathbb{R}^{k(k-1)/2}$ there exists at most unique $k \times k$ correlation matrix R such that $x = vech_{od}(LR)$.*

Note that the conjecture does not claim that for any real vector such correlation matrix necessarily exists. We formulate measurement and dynamic equations, analogous to (1.5) and (1.6), for the vectors $l\mathbf{y}_t^c$ and $l\rho_t^c$. Implications of Lemma 1 and Conjecture 1 make these vectors convenient for doing it. At first, the elements of both vectors are not bounded, which simplifies a formulation of the measurement equation. At second, vectors $l\mathbf{y}_t^c$ and $l\rho_t^c$ contain all information that is needed to restore the reciprocal correlation matrices Y_t^c and R_t^c.

We introduce realized measures by means of the measurement equation similarly to the case of realized variances,

$$l\mathbf{y}_t^c = \Psi_0^c + \Psi_1^c\, l\rho_t^c + \mathbf{v}_t^c \tag{1.7}$$

[9]Several other examples (apart from the original DCC model), where the dynamic matrix process for conditional correlations is proposed, include the consistent DCC model of Aielli (2013) as well as the modified DCC version suggested by Shephard et al. (2008). In the context of modeling dynamics of realized correlation matrix, Bauwens et al. (2012) developed an extension of the Conditional Autoregressive Wishart model of Golosnoy et al. (2012). They disentangle the dynamics of realized variances from realized correlations with the dynamic equation for the latter being closely related to the specification used in the DCC models.

[10]An idea to use matrix logarithm for modeling covariance matrices was introduced in Chiu et al. (1996). It was then exploited to model dynamic covariance matrices in the context of the multivariate GARCH (Kawakatsu (2006)), Stochastic Volatility (Ishihara et al. (2014)) and DCC models (Asai and So (2015)). Matrix logarithm has been also applied to study the dynamics of realized covariances. In particular, Bauer and Vorkink (2011) used log-matrix transform for modeling and forecasting realized covariance matrices. The matrix logarithm and exponential were considered as cases of the more general Box-Cox transformation and its inverse and then were applied to modeling realized covariance matrices in Weigand (2014).

where Ψ_0^c is a $\frac{1}{2}k(k-1)\times 1$ vector and Ψ_1^c is a $\frac{1}{2}k(k-1)\times\frac{1}{2}k(k-1)$ matrix of constant coefficients to capture a possible bias in realized measures. The vector of measurement errors $\mathbf{v}_t^c \sim N\big(0,\Omega_v^c\big)$ is supposed to be independent from $\mathbf{z}_t^c$, but may correlate with $\mathbf{u}_t^c$. The concatenated vector of all measurement errors is then distributed Normally with zero mean and variance

$$\Omega^c = \mathrm{Var}\begin{pmatrix}\mathbf{u}_t^c\\ \mathbf{v}_t^c\end{pmatrix} = \begin{pmatrix}\Omega_u^c & (\Omega_{uv}^c)'\\ \Omega_{uv}^c & \Omega_v^c\end{pmatrix}$$

The dynamic equation for $l\rho_t^c$ is given by the following autoregressive process

$$l\rho_t^c = \Lambda_0^c + \Lambda_1^c\, l\rho_{t-1}^c + \Lambda_2^c\, l\mathbf{y}_{t-1}^c \tag{1.8}$$

where Λ_0^c is a constant $\frac{1}{2}k(k-1)\times 1$ vector, Λ_1^c and Λ_2^c are coefficient matrices with the conformable size. Note that the elements from the vector $l\rho_t^c$ are not directly interpretable because each element in $l\rho_t^c$ is a function of the whole correlation matrix R_t^c. If we suppose that the true dynamics of the conditional correlations is sufficiently smooth and persistent, we may expect a similar regularity for $l\rho_t^c$ since the log-matrix transformation is continuously differentiable.

The estimation of the core model requires to have available the filtered matrix R_t^c at each time period t. Note that only the off-diagonal elements of LR_t^c are known from $l\rho_t^c$. In order to find the diagonal elements of LR_t^c we recall that $R_t^c = \mathrm{expm}(LR_t^c)$ has to be a correlation matrix and use Conjecture 1. Namely, by means of simple numerical methods the diagonal elements of LR_t^c can be uniquely identified from the condition that diagonal elements of $R_t^c = \mathrm{expm}(LR_t^c)$ are all equal to ones. As long as we have restored the matrix LR_t^c, we can find the conditional correlation matrix as $R_t^c = \mathrm{expm}(LR_t^c)$.

We note that for the case $k = 2$ the logarithm of a correlation matrix has a simple analytical form. It can be found (see Williams (1999), Bauer and Vorkink (2011)) that in this case

$$LR_t^c = \frac{1}{2}\begin{pmatrix}\log(1-(\rho_{21,t}^c)^2) & 2\,\mathrm{atanh}(\rho_{21,t}^c)\\ 2\,\mathrm{atanh}(\rho_{21,t}^c) & \log(1-(\rho_{21,t}^c)^2)\end{pmatrix}$$

where an off-diagonal element, $l\rho_t^c = \mathrm{atanh}(\rho_{21,t}^c)$, is the Fisher transform which maps from $(-1,1)$ to $\mathbb{R}$ and which was used in Hansen et al. (2014b), inter alia, for modeling the dynamics of the conditional correlation process in a bivariate GARCH framework. Thus, the Realized Beta GARCH model of Hansen et al. (2014b) is a special case of our Multivariate Realized GARCH model for $k = 2$. For larger k, the closed form of LR_t^c is becoming far more complicated.

1.2.3 Conditional Factor Model for Individual Assets

We formulate the conditional model for individual asset returns and corresponding realized measures of volatility and correlations with the factors. This model deals with the conditional density $p(\mathcal{X}_t^i|\mathcal{X}_t^c,\mathcal{F}_{t-1}^{c,i})$ from (1.2). Naturally, the conditional model represents an extension of the univariate Realized GARCH model to the dynamic factor framework with the factor dynamics described by the multivariate Realized GARCH from the previous subsection.

1.2.3.1 Conditional Distribution of an Individual Return

According to the conditional model, an individual asset return r_t^i is given by

$$r_t^i = \mu_i + \sqrt{h_t^i}z_t^i \tag{1.9}$$

where μ^i is the constant trend parameter, h_t^i is the unobservable conditional variance, and z_t^i is the standardized return, which is conditionally jointly distributed with the standardized factors. Namely,

$$\begin{pmatrix} \mathbf{z}_t^c \\ z_t^i \end{pmatrix} \Bigg| \mathcal{F}_{t-1} \sim N\left(\begin{pmatrix} \mathbf{0}_{k\times 1} \\ 0_{1\times 1} \end{pmatrix}, R_t^{c,i} \right) \tag{1.10}$$

where the $(k+1)\times(k+1)$ joint conditional correlation matrix $R_t^{c,i}$ is

$$R_t^{c,i} = \begin{pmatrix} R_t^c & \rho_t^i \\ (\rho_t^i)' & 1 \end{pmatrix} \tag{1.11}$$

The matrix R_t^c in (1.11) is the $k \times k$ factor correlation matrix filtered from the core model and ρ_t^i is a $k \times 1$ vector of correlations between an individual asset return r_t^i and the factors $\mathbf{r}_t^c$.

It follows from assumptions (1.10) and (1.11) that

$$z_t^i | \mathbf{z}_t^c \sim N\big((\rho_t^i)'(R_t^c)^{-1}\mathbf{z}_t^c,\ 1-(\rho_t^i)'(R_t^c)^{-1}\rho_t^i\big) \tag{1.12}$$

1.2.3.2 Conditional Variance and Correlations

In analogy with (1.5) and (1.6) we introduce the measurement and GARCH equations for the conditional variance h_t^i of an asset return r_t^i

$$\log x_t^i = \phi_0^i + \phi_1^i \log h_t^i + \delta_i(z_t^i) + u_t^i \tag{1.13}$$

$$\log h_t^i = \gamma_0^i + \gamma_1^i \log h_{t-1}^i + \gamma_2^i \log x_{t-1}^c + \tau_i(z_{t-1}^i) \tag{1.14}$$

where x_t^i is a realized measure of h_t^i, u_t^i is a measurement error, and leverage functions are defined as before

$$\delta_i(z_t^i) = \delta_1^i z_t^i + \delta_2^i((z_t^i)^2 - 1) \quad \text{and} \quad \tau_i(z_t^i) = \tau_1^i z_t^i + \tau_2^i((z_t^i)^2 - 1)$$

For the conditional correlations ρ_t^i and the corresponding realized measures $\mathbf{y}_t^i$ we suggest similar specifications

$$F(\mathbf{y}_t^i) = \Psi_0^i + \Psi_1^i F(\rho_t^i) + \mathbf{v}_t^i \tag{1.15}$$

$$F(\rho_t^i) = \Lambda_0^i + \Lambda_1^i F(\rho_{t-1}^i) + \Lambda_2^i F(\mathbf{y}_{t-1}^i) \tag{1.16}$$

where $F(\cdot)$ applies the element-wise Fisher transform to the vectors $\mathbf{y}_t^i$ and ρ_t^i in order to map each correlation element from the closed interval $(-1,1)$ onto $\mathbb{R}$. Similar to the core model, the measurement equation (1.15) defines the contemporaneous relationship between the $k \times 1$ vector of latent conditional correlations ρ_t^i and its realized measures from $\mathbf{y}_t^i$. Matrices Ψ_0^i and Ψ_1^i of the proper dimensions account for a possible measurement bias, whereas $\mathbf{v}_t^i$ stands for the $k \times 1$ vector of stochastic measurement errors. The equation (1.16) describes a transition of conditional correlations by imposing the vector autoregressive dynamics on the transformed conditional correlation, so it depends on its own lags and the lagged realized measures, as it is standard for the Realized GARCH framework.

As in the core model, we assume that the measurement errors from (1.13) and (1.15) are zero mean and Gaussian and we allow them to be correlated

$$\Omega^i = var\begin{pmatrix} u_t^i \\ \mathbf{v}_t^i \end{pmatrix} = \begin{pmatrix} \Omega_u^i & (\Omega_{uv}^i)' \\ \Omega_{uv}^i & \Omega_v^i \end{pmatrix}$$

Furthermore, we suppose that for each considered asset i the measurement errors from the core model (in (1.5) and (1.7)) have a joint Normal distribution with the measurement errors from the conditional model for this individual asset, therefore

$$\begin{pmatrix} \mathbf{u}_t^c \\ \mathbf{v}_t^c \\ u_t^i \\ \mathbf{v}_t^i \end{pmatrix} \sim N\left(\begin{pmatrix} \mathbf{0}_{k(k+1)/2} \\ \mathbf{0}_{k+1} \end{pmatrix}, \begin{pmatrix} \Omega^c & (\Omega^{c,i})' \\ \Omega^{c,i} & \Omega^i \end{pmatrix} \right) \tag{1.17}$$

Importantly, we keep an assumption that the measurement errors are independent from the return shocks $\mathbf{z}_t^c$ and z_t^i.

However, equations (1.15)-(1.16) require specific constraints on the parameter matrices and the measurement error specification in order to guarantee that the conditional correlations ρ_t^i, being filtered from the model, comply with the required positive definiteness of the joint correlation matrix (1.11) for all t. The following lemma imposes the necessary and sufficient condition on the elements from vector ρ_t^i to ensure that $R_t^{c,i}$ remains a correlation matrix.

Lemma 2. *Suppose that R is a $(k+1)\times(k+1)$ symmetric matrix with ones on the main diagonal. Let the upper left $k\times k$ block of R consists of non-singular correlation matrix R^c. Denote the upper right $k\times 1$ block of R by p. Then, R is a non-singular correlation matrix if and only if $p'(R^c)^{-1}p<1$.*

Fisher transforms in (1.15) and (1.16) only ensure that elements of ρ_t^i are less than 1 in absolute values, which is necessary, but not sufficient for the positive definiteness of $R_t^{c,i}$ as long as $k>1$. Nonetheless, we apply the current specification in our empirical analysis due to its simplicity and convenience, but check the condition $(\rho_t^i)'(R_t^c)^{-1}\rho_t^i<1$ at each $t=1,...,T$ for the estimated model. In case it is violated, we look for an another set of parameters that solves the likelihood maximization problem and generates admissible series ρ_t^i. In the subsection devoted to the model estimation, we describe this procedure in more details. On top of that, in Appendix A.2 we outline an alternative modeling strategy for the correlation process ρ_t^i via the spherical coordinate transform which automatically ensures a positive definiteness of $R_t^{c,i}$.

1.2.4 Multiple Regression Representation and Key Quantities

An asset or portfolio return r_t^i from (1.9) can be alternatively represented as a linear multiple regression

$$r_t^i = \alpha_t^i + (\beta_t^i)'\mathbf{r}_t^c + \varepsilon_t^i \tag{1.18}$$

where β_t^i is a $k\times 1$ vector of conditional regression coefficients (betas) that is defined as

$$\beta_t^i = \mathrm{Var}\big(\mathbf{r}_t^c\big|\mathcal{F}_{t-1}^{c,i}\big)^{-1}\mathrm{Cov}\big(r_t^i,\mathbf{r}_t^c\big|\mathcal{F}_{t-1}^{c,i}\big)$$

and can be written in terms of the model variables as

$$\beta_t^i = \big(H_t^c R_t^c H_t^c\big)^{-1}H_t^c\rho_t^i\sqrt{h_t^i} = \big(H_t^c\big)^{-1}\big(R_t^c\big)^{-1}\rho_t^i\sqrt{h_t^i} \tag{1.19}$$

where matrices H_t^c and R_t^c are filtered from the core model, whereas the conditional correlations ρ_t^i and the variance h_t^i are extracted from the individual factor model related to asset i. The intercept coefficient in (1.18) is formally expressed as $\alpha_t^i=\mu^i-(\beta_t^i)'\mu^c$ and compensates for possible non-zero factor means. The term ε_t^i is an idiosyncratic Gaussian error term with zero mean.

The multiple regression representation (1.18) is conformable with the conditional factor framework for asset returns (see Engle et al. (1990), Braun et al. (1995), Bollerslev and Zhang (2003) among many others) with time-varying factor loadings depending on the conditional covariance structure. Due to orthogonality of ε_t^i, the conditional idiosyncratic variance of r_t^i is readily obtained using (1.3) and (1.12), so we can write

$$\varepsilon_t^i = \sqrt{h_t^i(1-(\rho_t^i)'(R_t^c)^{-1}\rho_t^i)}\,\omega_t^i \tag{1.20}$$

where $\omega_t^i\sim$ i.i.d.$N(0,1)$ is a standardized idiosyncratic shock.

1.2.5 Large System of Assets

Consider now a sample of assets, $i = 1, ..., N$ and assume that the factors are common for all these assets. The factors may represent, for example, the sources of a systematic risk for asset returns as it is commonly implied in the Arbitrage Pricing Theory. The factor framework described above allows to study some dynamic properties of the joint conditional distribution of many individual assets. In doing so, we closely follow the framework developed in Hansen et al. (2014b). We first introduce a natural filtration that comprises all the information that is available up to day t related to the common factors and an asset universe

$$\mathcal{F}_t^{c,I} = \sigma\big(\mathcal{X}_t^c, \{\mathcal{X}_t^i\}_{i=1}^N, \mathcal{X}_{t-1}^c, \{\mathcal{X}_{t-1}^c\}_{i=1}^N, ...\big)$$

In order to apply our hierarchical model to a high dimensional vector of individual returns, we additionally assume that

$$p(\mathcal{X}_t^c|\mathcal{F}_{t-1}^{c,I}) = p(\mathcal{X}_t^c|\mathcal{F}_{t-1}^c) \quad \text{and} \quad p(\mathcal{X}_t^i|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,I}) = p(\mathcal{X}_t^i|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}) \tag{1.21}$$

for all $i = 1, ..., N$. Thus, both the core and the conditional factor models are not affected by the fact that the information from many assets is observed. Consequently, the models for individual assets can be set up and estimated independently from each other conditionally on the estimated core model for the factors.

An exposure to the common factors is the only channel for the conditional correlations between the considered assets in the model. Though this assumption is restrictive to a certain extent, it does simplify dramatically the modeling and estimation of large systems of assets.

Using (1.18), the vector of individual returns $\mathbf{r}_t = (r_{1,t}, r_{2,t}, ..., r_{N,t})'$ can be expressed as

$$\mathbf{r}_t = \alpha_t + B_t \mathbf{r}_t^c + \varepsilon_t$$

where B_t is $N \times k$ matrix with asset betas in the corresponding rows, $\alpha_t = (\alpha_t^1, \alpha_t^2, ..., \alpha_t^N)'$ is a vector of intercepts and $\varepsilon_t = (\varepsilon_t^1, \varepsilon_t^2, ..., \varepsilon_t^N)'$ is a vector of idiosyncratic asset returns. It then follows that the conditional variance of a N-dimensional vector of asset returns reads

$$\mathrm{Var}\big(\mathbf{r}_t|\mathcal{F}_{t-1}^{c,I}\big) = B_t \mathrm{Var}\big(\mathbf{r}_t|\mathcal{F}_{t-1}^{c,I}\big) B_t' + \Sigma_t \tag{1.22}$$

where $\Sigma_t = \mathrm{E}\big(\varepsilon_t \varepsilon_t'|\mathcal{F}_{t-1}^{c,I}\big)$ is a conditional covariance matrix of the idiosyncratic return components. Since we do not model explicitly the joint conditional density $p(\{\mathcal{X}_t^i\}_{i=1}^N|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,I})$, the dynamic properties of the residual correlations in Σ_t are left uncovered in our modeling framework[11].

Equation (1.22) illustrates the factor structure embedded in the conditional variance of the high dimensional return vector. Although the common factors are supposed to account for an amount of correlations within the system, we emphasize that some residual correlations may be contained in the idiosyncratic component Σ_t. Such correlations may come from the common factors or group (e.g., industry-specific) factors that are not included in the specification of the core model. In case the factors explain a common return variation sufficiently well, the parsimoniously formulated covariance (1.22) can significantly simplify the estimation, modeling and forecasting of large covariance matrices.

1.2.6 Estimation

Due to the structural similarities between the presented modeling framework and the Realized Beta GARCH model, the estimation principles are largely the same as in Hansen et al. (2014b). The model implies an observation driven dynamics for the time-varying volatility and correlation parameters, which is an appealing feature shared by the GARCH models. It permits a simple estimation by maximizing the likelihood function that

[11] Hansen et al. (2014b) found a significant amount of the residual correlation that is left unexplained in the one-factor (CAPM) model. On top of that, they found an evidence that such residual correlation is time-varying.

is analytically available. In addition, the hierarchical structure of the model suggests a convenient estimation procedure even if we deal with a large system of assets. We now briefly outline the estimation scheme and refer a more detailed description to Appendix A.3.

We begin with the core model. It describes the conditional distribution dynamics of a multivariate return vector. It is related to the density $p(\mathcal{X}_t^c|\mathcal{F}_{t-1}^c)$ from decomposition (1.2) that can be further factorized as

$$p(\mathcal{X}_t^c|\mathcal{F}_{t-1}^c) = p(\mathbf{r}_t^c, \mathbf{x}_t^c, Y_t^c|\mathcal{F}_{t-1}^c) = p(\mathbf{r}_t^c|\mathcal{F}_{t-1}^c)p(\mathbf{x}_t^c, Y_t^c|\mathbf{r}_t^c, \mathcal{F}_{t-1}^c)$$

where $p(\mathbf{r}_t^c|\mathcal{F}_{t-1}^c)$ is the marginal part related to the factors and $p(\mathbf{x}_t^c, Y_t^c|\mathbf{r}_t^c, \mathcal{F}_{t-1}^c)$ is the conditional part related to the corresponding realized covariation measure, taken conditionally on the factors. The associated log-likelihood function is

$$\mathcal{L}^c = \sum_{t=1}^{T} log\ p\big(\mathbf{r}_t^c|\mathcal{F}_{t-1}^c; \theta^c\big) + \sum_{t=1}^{T} log\ p\big(\mathbf{x}_t^c, Y_t^c|\mathbf{r}_t^c, \mathcal{F}_{t-1}^c; \theta^c\big) \tag{1.23}$$

where θ^c is a vector with the static parameters of the core model. The model is estimated by maximizing (1.23) with respect to θ^c.

Secondly, the conditional factor model for an individual asset return is related to $p(\mathcal{X}_t^i|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i})$ from (1.2). The subsequent decomposition of this conditional density gives

$$p(\mathcal{X}_t^i|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}) = p(r_t^i, x_t^i, \mathbf{y}_t^i|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}) = p(r_t^i|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i})p(x_t^i, \mathbf{y}_t^i|r_t^i, \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i})$$

where $p(r_t^i|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i})$ is a density of the daily asset return evaluated conditionally on the variables from the core model. The term $p(x_t^i, \mathbf{y}_t^i|r_t^i, \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i})$ is a density related to the realized measure conditionally on the daily asset return as well as on the variables from the core model. The corresponding log-likelihood is given by

$$\mathcal{L}^{i|c} = \sum_{t=1}^{T} log\ p(r_t^i|\mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}; \theta^i) + \sum_{t=1}^{T} log\ p(x_t^i, \mathbf{y}_t^i|r_t^i, \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}; \theta^i) \tag{1.24}$$

where θ^i stands for the parameter vector of the conditional factor model formulated for an asset i.

In order to estimate the model for a large system of N assets, it is then needed to estimate the core model and separately estimate N individual models for $i = 1, ..., N$ conditionally on the variables filtered from the core model.

The one-step ahead forecasting of the return distributions from the model is trivial due to the fact that all dynamic variables are specified in the observation driven manner, i.e. are the functions of lagged variables. Thus, given all relevant variables at day t, the time-varying parameters at $t+1$ are determined. The multi-step ahead forecasting can be accomplished via simulation methods or by bootstrapping. Such forecasting schemes designed for the Realized GARCH class of models are described and investigated, for instance, in Lunde and Olesen (2014) and Hansen et al. (2014b). We notice that bootstrap method is more preferable in the current context since it is more robust to possible distribution misspecifications.

1.3 Application to the Fama-French Model

We apply the realized factor framework developed in the previous section to the conditional version of the three-factor model of Fama and French. We describe the data that is used to replicate the Fama-French risk factors on a daily and an intra-daily frequencies as well as the data on the test assets. Then, we discuss the construction of the realized measures of variances and correlations that we use in our empirical analysis. We estimate the multi-factor Realized GARCH model for the large cross-section of stocks and then we apply the model to style and sectoral ETFs. Finally, we report and discuss the model estimation results.

1.3.1 Three-factor Fama-French Model

A conditional version of the Fama-French model (Fama and French (1992), Fama and French (1993)) implies that an asset return has the following generating process

$$r_t^i = \alpha_t^i + \beta_t^{mkt} r_t^{mkt} + \beta_t^{smb} r_t^{smb} + \beta_t^{hml} r_t^{hml} + \varepsilon_t^i \tag{1.25}$$

where r_t^i is an excess return on some asset i. Returns r_t^{mkt}, r_t^{smb} and r_t^{hml} are the market (in excess of risk-free rate), small-minus-big (market capitalization), and high-minus-low (book-to-market ratio) factors. Vector $\beta_t = \left(\beta_t^{mkt}, \beta_t^{smb}, \beta_t^{hml}\right)'$ is $\mathcal{F}_{t-1}$-measurable vector of conditional factor loadings (betas) and ε_t^i is an idiosyncratic shock with zero mean. The conditional mean (alpha) α_t^i is also allowed to be time-varying and represents a measure of an abnormal risk-adjusted asset performance.

The design of the multi-factor Realized GARCH developed in the previous section intends to model the dynamic conditional covariance of the three factors together with an individual return. In particular, the core part of the model represents a multivariate GARCH model for the three Fama-French factors. Equations (1.3) and (1.4) read

$$\begin{pmatrix} r_t^{mkt} \\ r_t^{smb} \\ r_t^{hml} \end{pmatrix} = \mu^c + H_t^c \begin{pmatrix} z_t^{mkt} \\ z_t^{smb} \\ z_t^{hml} \end{pmatrix}; \quad \begin{pmatrix} z_t^{mkt} \\ z_t^{smb} \\ z_t^{hml} \end{pmatrix} \Bigg| \mathcal{F}_{t-1}^c \sim N\left(0, R_t^c\right)$$

The conditional variances H_t^c and correlations R_t^c are modeled jointly with the corresponding realized measures. Thus, the core model filters the latent covariance structure of the Fama-French factors using signals from both low frequency daily factor returns and high frequency intra-daily realized covariance measures.[12]

A conditional model for an individual asset return represents a univariate GARCH model where the standardized return is assumed to be distributed jointly with the Fama-French factors, thus, directly implying a factor structure for a given return. Equations (1.9) and (1.10) read

$$r_t^i = \mu^i + \sqrt{h_t^i} z_t^i, \quad \left(z_t^{mkt}, z_t^{smb}, z_t^{hml}, z_t^i\right)' \big| \mathcal{F}_{t-1}^{c,i} \sim N\left(0, R_t^{c,i}\right)$$

The conditional variance h_t^i and the conditional correlations ρ_t^i are modeled jointly with the corresponding realized measures taking the conditional covariances between the factors, filtered from the core model, as given.

1.3.2 Data

The time period selected for the analysis begins on January 3, 2005 and lasts until December 31, 2013 spanning 2260 trading days in total. As it is implied by the design of the modeling framework, we use the data observed at distinct frequencies. Along with the factors and asset returns obtained at a daily frequency, we make use of intra-daily factors and asset price observations in order to construct corresponding realized measures of volatilities and correlations.

The close-to-close daily returns on the three Fama-French factors are extracted from the Kenneth French data library[13]. These factor returns are constructed using value-weight asset and portfolio returns formed on the size and book-to-market characteristics. Since almost all stocks listed on the NYSE, AMEX and NASDAQ are exploited, the factor returns from the Kenneth French library are highly representative for the broad U.S. equity market and are widely recognized as a benchmark in the empirical research.

The realized measures of variances and correlations related to the daily Fama-French factors are constructed with the high frequency data on exchange traded funds. Broad market and style ETFs can be treated as market and style portfolios and used to replicate risk factors. In Appendix A.5 we discuss in details both the selection of the proper ETFs and the scheme of how we construct the ETF-based risk factors. Since an intra-daily data

[12] We note that the process of the conditional factor covariance is an interesting object by itself. In particular, it is one of the central elements of the stochastic discount factor (SDF) implied by the linear factor model (e.g., see Cochrane (2001)).

[13] http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

on ETFs is available, it allows to construct realized measures for the Fama-French factors replicated by means of these ETFs at an intra-daily frequency. The intraday transaction data on ETFs is extracted from the Wharton TAQ database.

We estimate our model for a universe of individual assets. We select the stocks traded on NYSE, AMEX and NASDAQ throughout the whole length of our estimation period. In addition, we require that for each chosen asset we observe at least 10 transactions for each trading day from the considered period. The resulting sample consists of 824 stocks. We construct close-to-close daily returns on the individual stocks with the close prices extracted from the CRSP US Stock Database. Such prices are adjusted for the stock splits and dividends. Furthermore, each stock is accompanied by the CRSP's unique and permanent issue identification number, PERMNO. These identificators allow to keep track the selected companies even in case some of these companies change their ticker symbols in the course of the considered period. The intra-daily transaction data for the pool of individual assets is obtained from the TAQ database. We clean this data beforehand in accordance with the methodology suggested in Barndorff-Nielsen et al. (2009).

Furthermore, we use 9 style ETFs and 9 ETFs which track sectoral indexes as individual test assets (see Table 1.1). The corresponding realized measures are constructed with the data from TAQ, whereas the adjusted close prices are obtained from the Yahoo Finance.

Fund's type	Ticker	Description	Fund's type	Ticker	Description
Style funds	IWD	large value	Sectoral funds	XLB	Materials
BlackRock, iShares	IWB	large neutral	*SSGA, SPDR*	XLE	Energy
	IWF	large growth		XLF	Financials
	IWS	mid value		XLI	Industrials
	IWR	mid neutral		XLK	Technology
	IWP	mid growth		XLP	Consumer staples
	IWN	small value		XLU	Utilities
	IWM	small neutral		XLV	Health care
	IWO	small growth		XLY	Consumer discretionary

Table 1.1: List of individual ETFs used in the empirical analysis.

1.3.3 Realized Measures

An obvious challenge relates to the fact that the Fama-French factors are unobserved at intra-daily frequency, so we can not directly compute the corresponding realized measures. A growing interest in the intra-daily analysis of factor models has appeared recently in the literature. Usually factors are replicated by means of intra-daily returns on portfolios formed from a sorted cross-section of stocks (see Ait-Sahalia et al. (2014), Fan et al. (2015), Kalnina and Tewou (2015)). We use a slightly different approach and exploit intra-daily ETF data in order to replicate the Fama-French factors at a high frequency. A similar approach of extracting realized measures for the Fama-French factors was implemented in Bannouh et al. (2012).

Since the ETF-based factors just approximately replicate the benchmark Fama-French factors, we may expect that our realized measures constructed with the high frequency ETF data may be inconsistent measures of volatilities and correlations for the benchmark factors. We assume, however, that the ETF-based realized measures are well correlated with the latent factor covariance structure and, thus, are useful instruments. A possible systematic bias is then supposed to be captured within the model by the parameters of the measurement equations.

1.3.3.1 Realized Covariance for the Fama-French Factors

Realized covariance measures for the multivariate factor process are estimated using intra-daily price observations on ETFs. Suppose that for a given trading day t we observe intra-daily transaction prices for a set of market and style ETFs (portfolios). Lets denote by $\mathbf{p}^{etf} = (p^m, p^{sv}, p^{sn}, p^{sg}, p^{bv}, p^{bn}, p^{bg})'$ the vector with log-prices on market, small value, small neutral, small growth, big value, big neutral, and big growth ETFs

respectively. Denote by $s \in [0,1]$ an intra-daily time index. In agreement with the original definition of the Fama-French factors from Fama and French (1993), the factor "log-prices" can be obtained as the following linear combinations of ETF log-prices:

$$\begin{pmatrix} p_s^{mkt} \\ p_s^{smb} \\ p_s^{hml} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & -\frac{1}{3} & -\frac{1}{3} & -\frac{1}{3} \\ 0 & \frac{1}{2} & 0 & -\frac{1}{2} & \frac{1}{2} & 0 & -\frac{1}{2} \end{pmatrix} \mathbf{p}_s^{etf}, \quad s \in \mathcal{T}_t^c \tag{1.26}$$

where p^{mkt}, p^{smb}, and p^{hml} are "log-prices" on the market, the size, and the value Fama-French factors respectively. The set of time indexes $\mathcal{T}_t^c$ is determined by the refresh time sampling applied to the asynchronous intra-daily prices of all 7 funds observed at day t. We use the multivariate Realized Kernel estimator of Barndorff-Nielsen et al. (2011a) to obtain realized covariance measures, $RC^c_{MRK,t}$, of the three Fama-French factors from the intra-daily factor "log-prices" defined in (1.26). The realized correlation matrix that is used in the model is

$$Y_t^c = \text{diag}\big(RC^c_{MRK,t}\big)^{-\frac{1}{2}} RC^c_{MRK,t} \text{diag}\big(RC^c_{MRK,t}\big)^{-\frac{1}{2}}$$

where $\text{diag}\big(RC^c_{MRK,t}\big)$ is a diagonal matrix with the diagonal elements from $RC^c_{MRK,t}$.

Although the elements from $\text{diag}\big(RC^c_{MRK,t}\big)$ can be directly used as realized measures of factor variances $\mathbf{x}_t^c$, we will suggest more efficient volatility measures. The refresh time sampling in (1.26) leads to appreciable information losses, since we discard a lot of available intra-daily observations in order to bring all 7 asynchronous ETF transaction series to the same grid, $\mathcal{T}_t$. An improvement can be made by specifying 3 univariate factor "log-prices" instead of the one multivariate as in (1.26). It gives

$$\begin{aligned} p_s^{mkt} &= [1,0,0,0,0,0,0]\mathbf{p}_s^{etf}, && s \in \mathcal{T}_t^{mkt} \\ p_s^{smb} &= \left[0, \frac{1}{3}, \frac{1}{3}, \frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}, -\frac{1}{3}\right]\mathbf{p}_s^{etf}, && s \in \mathcal{T}_t^{smb} \\ p_s^{hml} &= \left[0, \frac{1}{2}, 0, -\frac{1}{2}, \frac{1}{2}, 0, -\frac{1}{2}\right]\mathbf{p}_s^{etf}, && s \in \mathcal{T}_t^{hml} \end{aligned} \tag{1.27}$$

where refresh time grids $\mathcal{T}_t^{mkt}$, $\mathcal{T}_t^{smb}$, and $\mathcal{T}_t^{hml}$ are generated by the transaction data from only those ETFs that are used for the replication of the corresponding factors. Thus, $\mathcal{T}_t^c$ from (1.26) is a subset of $\mathcal{T}_t^{mkt}$, $\mathcal{T}_t^{smb}$, and $\mathcal{T}_t^{hml}$. Consequently, each of the univariate intra-daily "log-price" series from (1.27) provide either the same or a higher number of observations than the multivariate intra-daily "log-price" from (1.26). Hence, the univariate Realized Kernel estimator of Barndorff-Nielsen et al. (2008) applied to the series from (1.27) is supposed to deliver more efficient realized measures of the factor variances than the ones in $\text{diag}\big(RC^c_{MRK,t}\big)$. If we denote by $D^c_{RK,t}$ a diagonal matrix with such univariate realized kernel estimates[14], then the vector of realized variances $\mathbf{x}_t^c$ that is used in the model consists of the diagonal elements from $D^c_{RK,t}$.

1.3.3.2 Realized Measures for Individual Asset Returns

The intra-daily realized measures of volatility of an individual asset return and correlations between a return and factors are estimated in a similar manner. At first, the univariate Realized Kernel provides a measure for the asset return variance x_t^i that is computed with intra-daily log-prices on this asset, p_i. Then, the realized correlation measures can be obtained from the multivariate Realized Kernel for the covariance between the three factors and an asset return, $RC^{c,i}_{MRK,t}$. Such 4×4 realized covariance matrix is estimated with the multivariate sequence of intra-daily factor price observations from (1.26) extended by the intra-daily asset price observations

[14]Note that realized covariance matrix $\big(D^c_{RK,t}\big)^{\frac{1}{2}} \text{diag}\big(RC^c_{MRK,t}\big)^{-\frac{1}{2}} RC^c_{MRK,t} \text{diag}\big(RC^c_{MRK,t}\big)^{-\frac{1}{2}} \big(D^c_{RK,t}\big)^{\frac{1}{2}}$ can be considered as a case of the block realized covariance estimator (Hautsch et al. (2012)), where diagonal elements are treated as blocks for which realized variances are estimated separately.

$$\begin{pmatrix} p_s^{mkt} \\ p_s^{smb} \\ p_s^{hml} \\ p_s^{i} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} & -\frac{1}{3} & -\frac{1}{3} & -\frac{1}{3} & 0 \\ 0 & \frac{1}{2} & 0 & -\frac{1}{2} & \frac{1}{2} & 0 & -\frac{1}{2} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \mathbf{p}_s^{etf} \\ p_s^{i} \end{pmatrix}, \quad s \in \mathcal{T}_t^{c,i} \tag{1.28}$$

Note that a refresh time sampling $\mathcal{T}_t^{c,i}$ is even more scarce than $\mathcal{T}_t^{c}$ in (1.26) since here we bring 7 asynchronous ETF transaction series together with the price observations from an individual asset to the unique grid.

Denote the corresponding realized correlation matrix as

$$Y_t^{c,i} = \text{diag}\big(RC_{MRK,t}^{c,i}\big)^{-\frac{1}{2}} RC_{MRK,t}^{c,i} \text{diag}\big(RC_{MRK,t}^{c,i}\big)^{-\frac{1}{2}}$$

Then the 3×1 vector of realized correlations between an asset return and the Fama-French factors, $\mathbf{y}_t^i$, is defined as the 3×1 upper right block of $Y_t^{c,i}$.

1.3.4 Core Model for the Fama-French Factors

The estimation results of the core model for the three Fama-French factors are reported in the last column in Tables 1.2-1.3. For comparative purposes and for the sake of robustness we also estimate three univariate Realized GARCH models ($k = 1$) for each of the factors separately (columns I-III) and three bivariate Realized GARCH models ($k = 2$) for the three possible pairs of the Fama-French factors (columns IV-VI). The results demonstrate that some of model parameters noticeably change with the inclusion of additional factors into the system. This might be an argument in favour of the joint estimation of the whole dynamic system instead of using marginal models to reduce the dimensionality of the optimization problem (as it is often practiced in estimation of the DCC models, for example). Note that in contrast to the variance related parameters, we can not directly associate correlation parameters to the correlations between the particular factors. It is because we model the dynamics of the off-diagonal elements of the correlation matrix logarithm, but not the correlation elements directly.

In the top of Table 1.2, we report the estimates of μ^c from the return equation. Since the Fama-French factors are essentially portfolios, the parameter μ^c indicates the mean return on a portfolio associated with a given factor. According to the estimates, only the market factor (excess return on the market portfolio) has a significant positive mean, whereas the size and the value factors (returns on the zero investment style portfolios) have insignificant unconditional expectations.

The estimates of the GARCH dynamic equations are reported further in Table 1.2. These estimates reveal the high persistence in the dynamics of the filtered conditional volatilities and correlations of the factors. As it follows from the VARMA representation of the conditional dynamics implied by the Realized GARCH model (see Hansen et al. (2012)), the persistence coefficients for the volatilities are given by $\Gamma_1^c + \Phi_1^c \Gamma_2^c$. The persistence parameters for the conditional variances of the market, size, and value factors are 0.976, 0.976, and 0.979 respectively. As for the off-diagonal elements of the transformed conditional correlation matrix, $\Lambda_1^c + \Psi_1^c \Lambda_2^c$ gives the persistence coefficients 0.973, 0.988, and 0.990.

Table 1.3 contains parameter estimates related to the measurement equations. The parameters in the measurement equations are strongly affected by the daily factors. As long as the daily factors are obtained with the Fama-French research portfolios and we use realized measures based on ETFs, we may expect a possible mismatch between the latent conditional covariance of daily factors and the ETF-based realized covariation measures. If this mismatch leads to a systematic bias, it is supposed to be captured by the parameters in the measurement equations.

Interesting to note, that all three Fama-French factors exhibit a significant leverage effect with positive T_1^c, D_1^c and negative T_2^c, D_2^c. Thus, the positive impact of negative returns on volatility, which is a stylized feature of the stock return volatility, seems also innate in the volatility dynamics of the Fama-French portfolio returns.

Table 1.4 contains indirectly estimated covariance matrix Ω^c of the measurement error vector. As it is

discussed in Hansen and Huang (2016), the measurement errors can be interpreted as shocks to the conditional dynamic variables. This is because the errors can be represented as updating terms in the Realized GARCH dynamic equations, if the measurement and the GARCH equations are combined together. As it can be seen from Table 1.4, the shocks to the conditional variances are far more volatile than the shocks to the conditional correlation components. The measurement errors related to the variance also demonstrate a significant amount of cross-correlation, which suggests that the volatilities of the Fama-French factors exhibit some common dynamics.

We plot the filtered conditional factor variances and correlations in Figure 1.1. Important to note that the variances and correlations filtered from the multivariate realized GARCH model incorporate the signals both from the realized measures and from the daily factor returns. Thus, the latent covariance dynamics of the benchmark Fama-French factors drives, to a large extent, the conditional covariance dynamics implied by the model. Also, Figure 1.1 demonstrates that the Fama-French factors exhibit a significant temporal variation in the correlation dynamics. It emphasizes the potential importance of having a time-varying specification for the whole factor covariance structure.

In Figure 1.2 we provide qq-plots of the standardized model residuals. The corresponding empirical densities do not perfectly match with the standard Normal density indicating some evidence of heavy tails. However, following Hansen et al. (2012) we rely on the quasi maximum likelihood argument that renders a consistency in the parameter estimates and we maintain the Gaussian distribution assumptions when estimate the model.

1.3.5 Conditional Models for ETFs and Stocks

In our empirical application, we estimate the individual factor models on the set of style and sectoral ETFs as well as on the large cross-section of stocks. The corresponding estimation result are reported in Tables 1.5-1.6. In addition, in Figure 1.3 we plot the cross-sectional distributions of some selected model parameters for the sample of 824 estimated individual assets. Below we comment on several regularities exhibited by the parameter estimates.

Parameters γ_1^i, γ_2^i, Λ_1^i, and Λ_2^i in the GARCH equations determine the relative contributions of the lagged values and the lagged realized measures to the current value of a given filtered variable. Intuitively, for the realized measures with a higher signal-to-noise ratio we expect to obtain higher coefficients γ_2^i and Λ_2^i. Note that although the coefficients γ_2^i related to the realized variance signals are above 0.2, the coefficients Λ_2^i related to the realized correlation signals are on average smaller. It may indicate that realized correlation measures are more noisy and less informative for the latent conditional dynamics than realized volatility measures. This finding may be a consequence of our approach to the construction of realized measures explained in the previous section. Whereas the correlation measures are estimated with the significant loss of information due to the refresh time sampling, the volatility measures are estimated from the univariate intra-daily price sequences without the efficiency losses associated with the refresh time.

The parameter ϕ_0^i captures the possible bias in the variance measurement equation and is strongly negative for almost all the considered assets. This finding is consistent with the fact that the realized volatility measures are based on the open-to-close intra-daily trading period and do not account for the overnight and "over-weekend" variation. In contrast, daily returns in return equations are close-to-close and incorporate this variation. As a result, realized measures are systematically underestimate close-to-close return variances and this is reflected in the parameters in the measurement equations. The corresponding bias correction parameters Ψ_0^i and Ψ_1^i in the correlation measurement equations have less straightforward interpretation partially because these equations operate with the non-linear Fisher transform of the correlations. As it can be seen from the results on the large cross-section of stocks, the additive bias term Ψ_0^i is on average close to 0, whereas the multiplicative bias coefficient Ψ_1^i is about 1. This may imply that the realized correlation measures are on average not too much biased for the daily close-to-close returns.

1.4 Empirical Analysis

1.4.1 Conditional Fama-French Betas

The Fama-French model in (1.25) implies that an asset return is exposed to the three common sources of a systematic risk - the market, size, and value risk factors. The factor loadings, or betas, reflect the sensitivity of a considered asset or portfolio return to the changes in the systematic risk factors and are treated as conventional indicators of risk. An availability of the precisely estimated factor loadings is important for a portfolio allocation, risk management and an empirical asset pricing analysis. Particularly, it allows to better measure and forecast the systematic risk and might lead to a more reliable risk premia estimation.

Using (1.19), we can extract model-implied factor loadings from the estimated conditional covariance structure. Figure 1.4 shows the quantile plots of the model-implied factor loadings obtained for the sample of 824 stocks. The corresponding cross-section summary statistics aggregated by years is given in Table 1.7. Interestingly, the cross-section distributions of all three Fama-French loadings exhibit a significant degree of temporal variation. This is highlighted in Figure 1.5, where the quantile plots from Figure 1.4 are zoomed to illustrate the distributions of betas in 2008 during the sharp phase of the Great Recession.

The distribution of market betas is centered around 1, which is pretty natural and implies that the average of the large cross-section of considered stocks moves closely with the market portfolio. Particularly, it indicates that our sample is representative enough for the US equity market. The distributions of loadings related to the size and value factors have means that lie between 0 and 1. We note that the distributions of the loadings associated with the book-to-market risk exhibits the highest temporal variation. This is especially noticeable around the financial crisis episode when the value betas swiftly became substantially right skewed. Particularly, it may indicate that many companies had become distressed and lost an amount of their market valuation at that period.

The factor exposures of the style and sectoral ETFs are presented on Figures 1.6-1.8. Style ETFs can be treated as portfolios sorted by the size and the book-to-market ratio. The betas on the style ETFs demonstrate substantial degrees of variation during the estimation period[15], especially for the small cap funds. The market loadings, specifically those of large funds, stay near 1 and this is pretty expected for the market representative portfolios. The relative risk exposures of the style ETFs to the size and value factors are consistent with the empirical definitions of these risk factors. Namely, the size betas are considerably high and positive for the small funds and mostly negative for the large funds, whereas the value betas are high for the value ETFs and low (negative) for the growth ETFs.

Sectoral ETFs can be considered as portfolios of stocks that belong to the specific economic industry or sector[16]. The materials, energy, and financial sectors have the most volatile dynamics of the risk exposures. Their market betas are on average larger than 1, which indicates that the returns on these funds are typically more volatile than the market. The industrials, consumer discretionaries, and technology sectoral funds demonstrate fairly stable factor loading dynamics and have market betas close to 1. Interestingly, the consumer staples, utilities, and health care sectoral funds exhibit the relatively small market risk exposures implying that these sectors act more orthogonally to the broad economy compared to the other considered funds. The size loadings of the energy, financial, consumer staples, utilities, and health care funds tend to be negative assuming the negative correlation between these ETFs and the small-minus-big factor. Particularly, this may suggest that these ETFs have an excessive exposure to large cap stocks. The loadings on the high-minus-low factor demonstrate volatile dynamics. The value betas of the energy and financials are mostly positive (except for the energy fund during the financial crisis period), whereas the technology and health care sectors demonstrate mainly negative

[15]Notice, that the factor loadings of almost all considered funds are especially volatile in the beginning of the estimation period (2005, 2006). This can be a consequence of noisy realized covariance estimates since the average number of transactions per day for the style ETFs used for intra-daily factor replication was not too high at that period.

[16]It is important to note that the ETF loadings on risk factors not necessarily indicate the risk exposures of the underlying economic sectors and may be affected by the fund managers' choice regarding a portfolio composition and style. Nonetheless, we rely on the fact that SPDR sectoral ETFs are passively managed funds and are "managed with the objective of matching the price and yield performance of their underlying sector indexes" (see the corresponding fund description on *www.sectorspdr.com*). Thus, we regard the returns on sectoral ETFs as close proxies for the returns on the industry-representative portfolios.

value exposures.

Our empirical results can shed some additional light on the findings from Andersen et al. (2005). In particular, using the reduced form model for realized CAPM betas they documented counter-cyclical behavior of the market betas with the effect is being more pronounced for the value stock portfolios. They demonstrate the substantial increase in the corresponding market betas during the 2001 recession suggesting the rise of the value premium may be associated with an increase in expected returns for value stocks in bad times. We may contribute to this discussion by analyzing the factor loadings implied by the multi-factor realized GARCH model for the style and sectoral ETFs around the Great Recession episode.

As it can be seen from Figure 1.6, the market loadings for the small and mid cap value portfolios exhibit quite moderate increase during the financial crisis, whereas we do not observe any significant change in the level of market betas for the remaining style funds at that time. A possible explanation is that, in contrast to Andersen et al. (2005), we use two additional factors that may explain a part of return variation which is solely attributed to the market factor in the one-factor CAPM model. Particularly, we observe a substantial rise of the correlation between the market and value factors in 2008 (see Figure 1.1). Therefore, in case we do not control for the value factor, the market risk loadings of value funds (which are exposed to the value factor by construction) tend to be mechanically increased during the Great Recession period. We, thus, confirm the findings in Andersen et al. (2005) in part that the counter-cyclical dynamics of market betas is the mostly feature of value stocks. Furthermore, an explicit control for the value risk substantially absorbs the effect.

The strong rise in the correlation between the market and value factors can be attributed to the growing number of distressed companies with a high book-to-market ratio across the economy in the crisis time. However, Figure 1.8 demonstrates that only the financial sector ETF exhibits a considerable increase in the book-to-market risk exposure at that period. It may imply that the most of value stocks appeared during the recession were from the financial sector and that could be the channel of the rising correlation between the market and high-minus-low risk factors.

1.4.2 Variance Decomposition

Another aspect of the Fama-French model that is convenient to analyze within the given framework relates to the decomposition of the return variance along the risk factors and the idiosyncratic return component and its evolution in time. Although it does not have direct implications for the pricing of the systematic risk, splitting the return variance can be important for a risk management and estimation of large covariance matrices.

The conditional variance of an asset (or portfolio) return r_t^i can be decomposed as follows

$$\mathrm{Var}\big(r_t^i|\mathcal{F}_{t-1}^{c,i}\big)= \big(\beta_t^i\big)'\mathrm{Var}\big(\mathbf{r}_t^c|\mathcal{F}_{t-1}^{c,i}\big)\beta_t^i + \mathrm{Var}\big(\varepsilon_t^i|\mathcal{F}_{t-1}^{c,i}\big)$$

The part of the conditional return variance h_t^i explained by the systematic risk factors can be characterized similarly to the regression R^2 and can be derived from (1.9)-(1.11)

$$V_t^{i,c} = \mathrm{Var}\big(r_t^i|\varepsilon_t^i;\mathcal{F}_{t-1}^{c,i}\big)/\mathrm{Var}\big(r_t^i|\mathcal{F}_{t-1}^{c,i}\big)= (\rho_t^i)'(R_t^c)^{-1}\rho_t^i \tag{1.29}$$

whereas the part of the conditional return variance explained by the idiosyncratic component corresponds to $1-R^2$ and reads

$$V_t^{i,\varepsilon} = 1 - V_t^{i,c} = 1 - (\rho_t^i)'(R_t^c)^{-1}\rho_t^i \tag{1.30}$$

The further decomposition of $V_t^{i,c}$ along the individual contributions of the three risk factors is not straightforward due to the implied correlation between factors. Although an exact variance decomposition requires additional identifying restrictions on the causal structure of the factors, there exist several statistical methods which may be used to assess the factors' "relative importance"[17].

[17]See, for example, a related methodological review in Gromping (2007)

We apply the approach suggested in Genizi (1993). The method rests on finding the orthogonal basis that is closest to the system of factors (regressors) with respect to a particular norm. Then, the factors together with the considered return are decomposed along this basis. The contribution of a particular factor to the total return variance is then computed indirectly as a sum of variance contributions from the orthogonal components of the decomposed factor. Although this approach to the variance decomposition is somewhat approximate, we use the corresponding results to discuss several qualitative implications.

Figure 1.9 demonstrates the decomposed variance fractions averaged across 824 individual assets considered in the analysis. Table 1.8 provides the corresponding cross-section decomposition statistics by years and factors. The idiosyncratic component on average accounts for 60-70% of the total return variation throughout most of the estimation period except for the several relatively short intervals when the variation attributed to the Fama-French factors had been significantly amplified. The timing of such periods coincides with the Great Recession (2008-09) and with the two waves of the Euro-area Sovereign Debt Crisis (2010 and the second half of 2011). Note that these periods are also flagged by the increased overall stock market volatility level.

The relative variance contribution of the Fama-French factors is not identical during the considered period. The market factor contributes on average 20-25% to the total return variance in normal times and raised up to 35-40% during the volatile episodes. The small-minus-big factor demonstrates a cyclical intensity. Thus, it explained about 7-8% of the average variation in stock returns in the first two years of the sample, then it reduced to 4-5% during the Financial Crisis, raised again to 7-8% in 2010-11 and then faded to 4-5% in the end of the sample period. Such periodic activity can hardly be attributed to the business cycle fluctuations of the U.S. economy, but may be a consequence of a mechanical effect of changing correlations between the market and the size Fama-French factors (see Figure 1.1). The value factor is reflected slightly in stock return variances and explains just about 2% for most of the period. However, the relative variance shares attributed to the value factor boosted to almost 8% in 2008 and then stayed high until the middle of 2010. This may indicate that the proportion of the value risk in the stock return variance is significantly pronounced during recessions and the aftermath periods.

Figure 1.10 contains variance decomposition results for the style and sectoral ETFs. As expected, since ETFs are actually portfolios of individual stocks, the idiosyncratic portion of variance is significantly smaller than in the previous case due to diversification. For most of sectoral ETFs the idiosyncratic component accounts for between 25% and 50% of the total variation, whereas for the style funds it rarely exceeds 15%. The 9 style ETFs represent a small cross-section of portfolios sorted by the size and the book-to-market. It is not surprising that the contribution of the market factor is higher for large cap funds (~80%), but lower for small cap funds (~60%). The size factor, conversely, brings appreciably more in the variation of the small cap funds' returns than of the large cap ones (~30% vs ~10%). As for the value factor, it contributes to the ETF return variance mostly during the period between 2008 and 2011. This supports our findings related to the value factor contribution based on the cross-section of individual assets.

An analysis of sectoral funds reveals how well the Fama-French factors explain a return variation when applied to different sectors of the U.S. equity market. We note that the variance fractions attributed to the size and the value factors are smaller than for the style funds because the composition of stocks in sectoral portfolios is sufficiently well-diversified along the market cap and the book-to-market characteristics. The utilities, health care, and consumer staples ETFs demonstrate especially low exposure to the Fama-French factors. In contrast, technology, consumer discretionary, industrials, and financials are among the most exposed funds.

Almost all considered sectoral ETFs corroborate the evidence that the intensity of the high-minus-low factor amplified during the financial crisis and then had been vanishing gradually. The financial sector ETF, however, experienced an especially pronounced impact from the value factor on the return variance (up to 35% during the Crisis phase). What is interesting, the sharp rise of its exposure to the value factor had began several months before the recession turning point identified by NBER. This fact might imply that the interconnection between the financial sector and the value risk contains an information that can be useful to detect recessions on their early stages. An enormously large exposure of the financial sectoral fund to the high-minus-low factor (relative to other sectoral funds) might indicate that the value risk is not that much "systematic" and is mostly

the feature of assets from the particular sectors of economy.

Another interesting phenomenon can be observed in the commodity related sectors - materials and energy. Shortly after the beginning of the crisis the idiosyncratic variation in the corresponding returns became strongly predominant. The rise of idiosyncratic volatility in these sectors just partially explains the phenomenon. As we can see from Figures 1.6-1.8, the risk factor loadings of the energy ETF (and of the materials ETF to a lesser extent) significantly changed during the crisis with the market beta of the energy fund decreased considerably and was near zero in late 2008. Such effect also attributed to the rising "orthogonality" between these specific sectors, the global U.S. economy, and the Fama-French systematic risk factors. Particularly, this evidence favours a widespread intuition about the strong diversification capability of commodity related securities (see Gorton and Rouwenhorst (2006), for example). After the crisis, however, the exposure of the energy and material ETFs to the common risk factors had promptly recovered.

To summarize roughly, the Fama-French three factors leave unexplained more than 60% of variation in the representative individual equity returns. Therefore, in order to make a good job in modeling large covariance matrices of individual assets the model has to be augmented by additional factors, e.g. sectoral specific, to explain more variation by the common factor dynamics. On the other hand, the three factor framework can be effectively exploited in modeling large covariances of portfolios or funds, where the amount of idiosyncratic variation is considerably lower due to the averaging effect.

1.5 Pricing of Systematic Risk

One natural application of the proposed model relates to the estimation of premia for an asset's exposure to the sources of systematic risk. In this section, we discuss how the suggested modeling framework can be useful to explain expected returns, discuss some properties of the pricing kernel implied by the model, and make an attempt to measure the risk premia for the Fama-French factors using the large cross-section of equities.

1.5.1 Risk Premia

The beta representation of conditional linear factor models (see Jagannathan and Wang (1996), Bollerslev and Zhang (2003), Lewellen and Nagel (2006), etc.) implies that $N \times 1$ vector of cross-sectional expected excess returns can be expressed as follows

$$\mathrm{E}_t \mathbf{r}_{t+1} = i_N \lambda_{0,t+1} + B_{t+1} \lambda_{1,t+1} \tag{1.31}$$

where i_N is the $N \times 1$ vector of ones and B_{t+1} is a $N \times 3$ matrix containing the conditional factor loadings of N considered assets. The vector $\lambda_{1,t+1} = \left(\lambda_{mkt,t+1}, \lambda_{smb,t+1}, \lambda_{hml,t+1}\right)'$ contains risk premia for the Fama-French factors for the period $t+1$. The scalar parameter $\lambda_{0,t+1}$ stands for the so-called zero-beta return and is intended to capture possible risk-free rate misspecifications.

A prevailing approach to a risk premia estimation in empirical asset pricing relates to the regression based methods (Black et al. (1972), Fama and MacBeth (1973)). Usually, the implementation proceeds in two steps. The first stage consists of (rolling) time series regressions of excess returns on the factors and is aimed to estimate factor loadings of the considered assets. The second stage implies the use of cross-section regressions of excess returns on the estimated betas in order to identify premia for an exposure to the systematic risk factors. Alternatively, both factor loadings and risk premia can be estimated jointly in the GMM framework (see Cochrane (2001)). For the conditional models, where both loadings and risk premia are allowed to be time-varying, an ability of static regression methods to identify the parameters of interest is limited.

Our model suggests a natural alternative to the first-stage time series regression in Fama and MacBeth (1973). The time-varying conditional betas filtered from the model provide us with the sequence of matrices B_{t+1}, so the risk premia $\lambda_{1,t+1}$ in (1.31) can be directly estimated by means of cross-section regressions or GMM.

The standard practice in estimation of asset pricing models rests on using portfolios formed by various stocks' characteristics as base assets. Grouping stocks into portfolios is supposed to remove a part of the idiosyncratic return variation and helps to more accurately estimate betas in the first stage regressions (Blume (1970), Fama and French (1993)). This approach, however, leads to considerable losses of information about the broad cross-section distribution of the factor loadings and may increase the estimation errors of risk premia, as a consequence (Ang et al. (2010)). In addition, the principal component structure of the sample of portfolios sorted according to the factor related characteristics (e.g. size or book-to-market) may cause a spurious assessment of the asset pricing model implications (Lewellen et al. (2010)). Since we suggest much more careful way for filtering the factor loadings, we expect a sufficient degree of accuracy for model-implied betas. Thus, for the risk premia estimation we are able to exploit more disperse and granular distribution of factor loadings for the large cross-section of individual assets rather than use portfolio betas estimated by rolling linear regressions.

Since our model is set at a daily frequency, potentially we can estimate conditional risk premia on a daily basis using the cross-sections of daily excess returns. In Appendix A.7, we provide the plots with estimated risk premia at a daily and a monthly frequency based on our cross-section of 824 individual assets. The plots reveal a strong time variation of the estimated risk premia providing an additional evidence in favour of taking into account the dynamics of factor loadings when pricing systematic risk. Note that the current study was limited by the time span (only 9 years are considered) and, more importantly, by the asset sample. Despite we use more than 800 stocks in our cross-section analysis, the broad market representativeness of the sample is still a concern and a possible insufficient cross-section variation along the size and book-to-market characteristics could deteriorate the estimation results.

Furthermore, we expect that the error-in-variables is still a substantial problem. Although the daily model-implied factor loadings are supposed to be more consistent with factor returns and less noisy than the plain realized betas, the measurement error may still be appreciable and the risk premia estimates will likely be considerably attenuated. Moreover, daily stock returns are relatively noisy which leads to the presence of outlying observations and additional distortions of the risk premia estimates. Therefore, we may need a coherent econometric design in order to filter the conditional risk premia (consistent with our model-implied betas) possibly at a daily frequency and from presumably very noisy signals. The work on these issues is the subject of the ongoing research.

1.5.2 Implied Pricing Kernel

In order to further illustrate the potential asset pricing implications from the use of model-implied variables we refer to the alternative SDF (stochastic discount factor) representation for pricing models (see Cochrane (1996), Hansen and Jagannathan (1997), Hodrick and Zhang (2001), etc.). The SDF, or pricing kernel, implied by the Fama-French model is affine in risk factors and reads

$$m_t = 1 + b_t'(\mathbf{r}_t^c - \mathrm{E}[\mathbf{r}_t^c]) \tag{1.32}$$

where the 3×1 vector b_t consists of conditional weights attached to demeaned factors in the pricing kernel specification. Expression (1.32) also implies that m_t is normalized, such that $\mathrm{E}[m_t] = 1$. In an absence of arbitrage, the following fundamental pricing identity holds

$$\mathrm{E}_t[m_{t+1}\mathbf{r}_{t+1}] = \mathbf{0}_N \tag{1.33}$$

Equation (1.33) can be associated with the Euler equation and the pricing kernel m_t in this case is interpreted as a state-dependent intertemporal marginal rate of substitution where states are determined by the risk factors. Therefore, coefficients b_t directly reflect the importance of a given factor in the investors' attitudes towards the systematic risk.

By substituting (1.32) in (1.33) we can establish the one-to-one correspondence between the risk premia parameters $\lambda_{1,t}$ and the pricing kernel coefficients b_t,

$$b_t = -\text{Var}_{\text{t-1}}(\mathbf{r}_t^c)^{-1}\lambda_{1,t} \tag{1.34}$$

Thus, the conditional covariance process $\text{Var}_{\text{t-1}}(\mathbf{r}_t^c)$ appears as one of the key elements of the pricing kernel m_t. In case we assume that $\text{Var}_{\text{t-1}}(\mathbf{r}_t^c) = \text{Var}(\mathbf{r}_t^c|\mathcal{F}_{t-1}^c)$, it can be filtered directly from the core model, $\text{Var}_{\text{t-1}}(\mathbf{r}_t^c) = H_t^c R_t^c H_t^c$. Then, conditionally on model-implied $\text{Var}_{\text{t-1}}(\mathbf{r}_t^c)$, the risk premia parameters $\lambda_{1,t}$ can be estimated from equations (1.32)-(1.34), for instance, by means of the GMM-based methods (see Hansen (1982), Jagannathan et al. (2002), etc.). Importantly, when we estimate risk premia using the SDF approach, we do not need to extract factor loadings. In other words, we only have to filter the conditional factor covariances from the core model and do not need to estimate N conditional factor models for individual assets or portfolios[18].

1.6 Conclusion and Outline

In this paper, we suggest a novel framework for modeling asset returns with an underlying factor structure. The model carefully describes the covariance dynamics between an asset and possibly correlated factors, incorporates information from low frequency returns, and exploits realized volatility measures extracted from high frequency data. The modeling framework essentially consists of two parts. The first part represents a multivariate generalization of the realized GARCH model (Hansen et al. (2012)). It jointly models conditional covariance dynamics for a vector of returns, or factors, and the corresponding realized covariance matrix preserving positive definiteness at each step. The second part is a linear factor model for an asset return that is formulated conditionally on the covariance structure of factors. This part extends the single-factor realized Beta GARCH model of Hansen et al. (2014b) to the case with multiple factors.

We apply our framework to the three-factor Fama-French model. We extract realized measures for the Fama-French risk factors from high frequency intra-daily data on the panel of liquid ETFs that reflect different investment styles. Then, we estimate the model at a daily frequency using the large cross-section of U.S. stocks as well as the style and sectoral ETFs as test assets. We document a significant extent of temporal variation in the filtered factor loadings and in their cross-section distributions. Furthermore, we provide a day-by-day decomposition of asset return variances along the factors of systematic risk that illustrates periodically exciting nature of the Fama-French factors. Importantly, we detect a pronounced heterogeneity in the decomposition of the explained variance across different sectors of economy. Finally, we investigate the asset pricing implications of our model by cross-section regressions of excess equity returns on the model-implied factor loadings. We find an appreciable temporal variation in the premia estimates for all three risk factors during our sample period. The empirical evidence suggested in our study additionally emphasizes the role of dynamics in the conditional covariance of factors and in the corresponding betas for statistical modeling of multivariate returns and an asset pricing analysis.

1.6.1 Outline

We now briefly outline several extensions and promising directions for the future research, which either directly complement to this study or naturally stem from it.

1. Although we focus exclusively on the Fama-French three factor model in our empirical application, the presented framework can be applied to any linear factor model as long as appropriate factor series and realized measures can be obtained. For instance, from the large cross-section of stocks it is possible to form the portfolios of short-term "winners" and "loosers" (portfolios of stocks with the best and worst recent performance) and the intra-daily data can then be used to construct the corresponding realized measures. Then the momentum risk factor (a long portfolio of "winners" and a short portfolio of "loosers") can be

[18] Besides the standard methods, equations (1.32)-(1.34) can be used to extract the risk premia $\lambda_{1,t}$ assuming it is a time-varying latent parameter. It can be made, for example, with the observation-driven approach described in Creal et al. (2015), where the dynamics of the latent parameter is updated in a way to improve (locally) the GMM criterion.

incorporated into the model as an additional factor of a systematic risk (Carhart (1997)). Similarly, many other factors can be investigated within the suggested framework, including "quality-minus-junk" (Asness et al. (2014)), profitability and investment (Fama and French (2015)), etc.

2. A natural and typical application of conditional asset pricing models is to explain asset returns out-of-sample. Our modeling framework allows to readily produce one (or several) periods ahead forecasts of all relevant covariance variables and factor loadings which can then be used to form sorted or mean-variance efficient portfolios (e.g., Bollerslev and Zhang (2003), Bali et al. (2013)). The out-of-sample performance of such portfolios can serve both for evaluation of a given asset pricing model and to assess the economic significance of using dynamic model-implied betas and high frequency data.

3. The factor structure of the presented framework facilitates the feasible modeling of mid- and large-scale covariance matrices. The model-implied covariance matrix that incorporates information from high frequency observations and is consistent with the daily (low frequency) returns allows to expect a promising degree of precision. The part of covariance structure that is not explained by factors can also be incorporated into the model dynamics after some additional regularization (as in Engle et al. (2016), for instance).

4. Throughout the paper, we considered asset and factor returns mainly at a daily frequency and constructed realized measures using intra-daily data. Nonetheless, the model can be set at other frequencies as well. For example, we may use monthly returns on assets, portfolios and factors and construct realized variance and covariance measures using daily returns observed within a given month. In result, we may estimate and test asset pricing factor models for much longer sample periods (e.g., using long daily and monthly series from the Kenneth French data library). Furthermore, at a monthly frequency our results would become better comparable with the most of empirical asset pricing results obtained for monthly returns.

5. The multivariate realized GARCH framework incorporates many stylized aspects of financial returns such as leverage effects and cross-asset spillovers of volatility, can accommodate the heavy-tailed features, and, at the same time, the estimation and forecasting from the model are remarkably simple. Moreover, the risk factors that are introduced into the model impose an additional structure on the specification of return distributions. An interesting extension would be to produce multi-horizon forecasts of a return density (univariate and multivariate) from the model (that is set at a daily or monthly frequency). Then, the role of the systematic risk factors and their covariance dynamics could be assessed from the viewpoint of density forecasting of asset returns.

6. The model generates useful variables that not only are important for modeling asset returns, but also may contain relevant information about the latent state of the macroeconomic activity. Given a voluminous theoretical research that relates the systematic factors of equity risk to the macroeconomic fundamentals, the filtered factor loadings, for example, might be used to model and forecast indicators of the macroeconomic activity (Andersen et al. (2005)). The model's ability to timely detect the temporal changes in the filtered series makes it possible to efficiently incorporate such potentially informative variables into the real time nowcasting frameworks (Aruoba et al. (2009)).

7. The cross-section distributions of the risk factor loadings extracted from the model on a daily basis can provide a useful information about the local relevance of factors. The presence of irrelevant factors (those which just weakly correlate with asset returns) in the model may lead to identification failures and spurious asset pricing inference (see Kan and Zhang (1999), Kleibergen (2009), Gospodinov et al. (2014)). An availability of granular distributions of the estimated loadings motivates the development of formal testing procedures that would be able to verify the broad market or sectoral specific relevance of a particular risk factor day-by-day. This would make it possible to accurately identify the periods of "strength" and "weakness" for the considered factors.

		Univariate			Bivariate			Full
Category	Parameters	I	II	III	IV	V	VI	VII
Return trend	μ_c	0.031			0.037	0.030		0.033
(*market*)		(*0.017*)			(*0.016*)	(*0.016*)		(*0.015*)
Return trend	μ_c		0.002		0.002		0.004	0.002
(*size*)			(*0.010*)		(*0.010*)		(*0.010*)	(*0.010*)
Return trend	μ_c			0.007		0.010	0.006	0.009
(*value*)				(*0.008*)		(*0.008*)	(*0.008*)	(*0.008*)
Variance	Γ_0^c	0.152			0.132	0.119		0.103
(*market*)		(*0.014*)			(*0.012*)	(*0.012*)		(*0.010*)
	Γ_1^c	0.707			0.745	0.742		0.774
		(*0.019*)			(*0.016*)	(*0.017*)		(*0.014*)
	Γ_2^c	0.265			0.228	0.233		0.202
		(*0.018*)			(*0.015*)	(*0.016*)		(*0.013*)
	T_1^c	-0.130			-0.099	-0.108		-0.087
		(*0.008*)			(*0.007*)	(*0.007*)		(*0.006*)
	T_2^c	0.017			0.015	0.012		0.012
		(*0.005*)			(*0.004*)	(*0.004*)		(*0.003*)
Variance	Γ_0^c		-0.033		-0.031		-0.021	-0.025
(*size*)			(*0.011*)		(*0.011*)		(*0.010*)	(*0.010*)
	Γ_1^c		0.721		0.718		0.756	0.743
			(*0.021*)		(*0.017*)		(*0.018*)	(*0.016*)
	Γ_2^c		0.252		0.255		0.224	0.233
			(*0.018*)		(*0.015*)		(*0.016*)	(*0.014*)
	T_1^c		-0.058		-0.041		-0.044	-0.033
			(*0.008*)		(*0.006*)		(*0.007*)	(*0.006*)
	T_2^c		0.019		0.012		0.019	0.012
			(*0.005*)		(*0.004*)		(*0.005*)	(*0.004*)
Variance	Γ_0^c			0.205		0.222	0.187	0.197
(*value*)				(*0.020*)		(*0.020*)	(*0.018*)	(*0.018*)
	Γ_1^c			0.768		0.753	0.784	0.775
				(*0.018*)		(*0.017*)	(*0.016*)	(*0.016*)
	Γ_2^c			0.214		0.226	0.197	0.204
				(*0.017*)		(*0.016*)	(*0.015*)	(*0.014*)
	T_1^c			-0.007		-0.006	-0.004	-0.003
				(*0.007*)		(*0.006*)	(*0.006*)	(*0.006*)
	T_2^c			0.027		0.025	0.024	0.023
				(*0.004*)		(*0.004*)	(*0.004*)	(*0.003*)
Correlation	Λ_0^c				-0.028			-0.023
(*market-size*)					(*0.011*)			(*0.011*)
	Λ_1^c				0.811			0.808
					(*0.017*)			(*0.017*)
	Λ_2^c				0.232			0.228
					(*0.029*)			(*0.029*)
Correlation	Λ_0^c					0.038		0.038
(*market-value*)						(*0.005*)		(*0.005*)
	Λ_1^c					0.829		0.837
						(*0.014*)		(*0.014*)
	Λ_2^c					0.181		0.169
						(*0.018*)		(*0.018*)
Correlation	Λ_0^c						0.001	-0.001
(*size-value*)							(*0.003*)	(*0.002*)
	Λ_1^c						0.886	0.918
							(*0.012*)	(*0.010*)
	Λ_2^c						0.116	0.092
							(*0.017*)	(*0.016*)

Table 1.2: Parameter estimates of the core model for the three Fama-French factors. Here listed those parameters that appear in the return and the GARCH equations of the model. Columns I-III provide the estimates of the univariate Realized GARCH (core) models for each of the three Fama-French factors. Columns IV-VI provide the estimates of the bivariate Realized GARCH (core) models for each pair composed from the three Fama-French factors. Column VII provides the estimates of the three-variate Realized GARCH (core) model that is based on all three factors. Standard errors are obtained using numerical gradient and reported in parenthesis. Estimation period is January 3, 2005 - December 31, 2013.

		Univariate			Bivariate			Full
Category	Parameters	I	II	III	IV	V	VI	VII
Variance	Φ_0^c	-0.588			-0.593	-0.533		-0.537
(*market*)		(*0.030*)			(*0.030*)	(*0.030*)		(*0.029*)
	D_1^c	-0.174			-0.139	-0.149		-0.125
		(*0.011*)			(*0.010*)	(*0.010*)		(*0.009*)
	D_2^c	0.057			0.038	0.045		0.031
		(*0.006*)			(*0.005*)	(*0.005*)		(*0.005*)
Variance	Φ_0^c		-0.007		-0.015		-0.018	-0.021
(*size*)			(*0.029*)		(*0.029*)		(*0.029*)	(*0.030*)
	D_1^c		-0.060		-0.031		-0.046	-0.026
			(*0.011*)		(*0.009*)		(*0.010*)	(*0.009*)
	D_2^c		0.095		0.078		0.089	0.075
			(*0.007*)		(*0.006*)		(*0.007*)	(*0.006*)
Variance	Φ_0^c			-1.107		-1.133	-1.108	-1.132
(*value*)				(*0.030*)		(*0.030*)	(*0.030*)	(*0.030*)
	D_1^c			-0.031		-0.025	-0.019	-0.018
				(*0.011*)		(*0.010*)	(*0.010*)	(*0.010*)
	D_2^c			0.062		0.056	0.053	0.050
				(*0.006*)		(*0.006*)	(*0.006*)	(*0.006*)
Correlation	Ψ_0^c				0.168			0.154
(*market-size*)					(*0.034*)			(*0.035*)
	Ψ_1^c				0.709			0.723
					(*0.066*)			(*0.070*)
Correlation	Ψ_0^c					-0.198		-0.208
(*market-value*)						(*0.024*)		(*0.026*)
	Ψ_1^c					0.884		0.893
						(*0.054*)		(*0.060*)
Correlation	Ψ_0^c						-0.016	0.002
(*size-value*)							(*0.020*)	(*0.019*)
	Ψ_1^c						0.890	0.783
							(*0.103*)	(*0.110*)

Table 1.3: Parameter estimates of the core model for the three Fama-French factors. Here listed those parameters that appear in the measurement equations of the model. Columns I-III provide the estimates of the univariate Realized GARCH (core) models for each of the three Fama-French factors. Columns IV-VI provide the estimates of the bivariate Realized GARCH (core) models for each pair composed of the three Fama-French factors. Column VII provides the estimates of the three-variate Realized GARCH (core) model that is based on all three factors. Standard errors are obtained using numerical gradient and reported in parenthesis. Estimation period is January 3, 2005 - December 31, 2013.

$$
\begin{array}{r}
 \\
u_{mkt} \\
u_{smb} \\
u_{hml} \\
v_{mkt-smb} \\
v_{mkt-hml} \\
v_{smb-hml}
\end{array}
\begin{array}{c}
\begin{array}{cccccc}
u_{mkt} & u_{smb} & u_{hml} & v_{mkt-smb} & v_{mkt-hml} & v_{smb-hml}
\end{array} \\
\begin{pmatrix}
0.242 & 0.128 & 0.096 & 0.009 & 0.006 & -0.009 \\
0.530 & 0.240 & 0.084 & 0.037 & -0.002 & -0.007 \\
0.396 & 0.345 & 0.244 & 0.005 & -0.005 & -0.006 \\
0.075 & 0.303 & 0.044 & 0.062 & -0.005 & 0.000 \\
0.050 & -0.015 & -0.040 & -0.084 & 0.067 & 0.008 \\
-0.073 & -0.059 & -0.047 & 0.000 & 0.131 & 0.059
\end{pmatrix}
\end{array}
$$

Table 1.4: Covariances and correlations between measurement errors estimated from the multivariate realized GARCH (core) model that is based on the three Fama-French factors. Estimation period is January 3, 2005 - December 31, 2013.

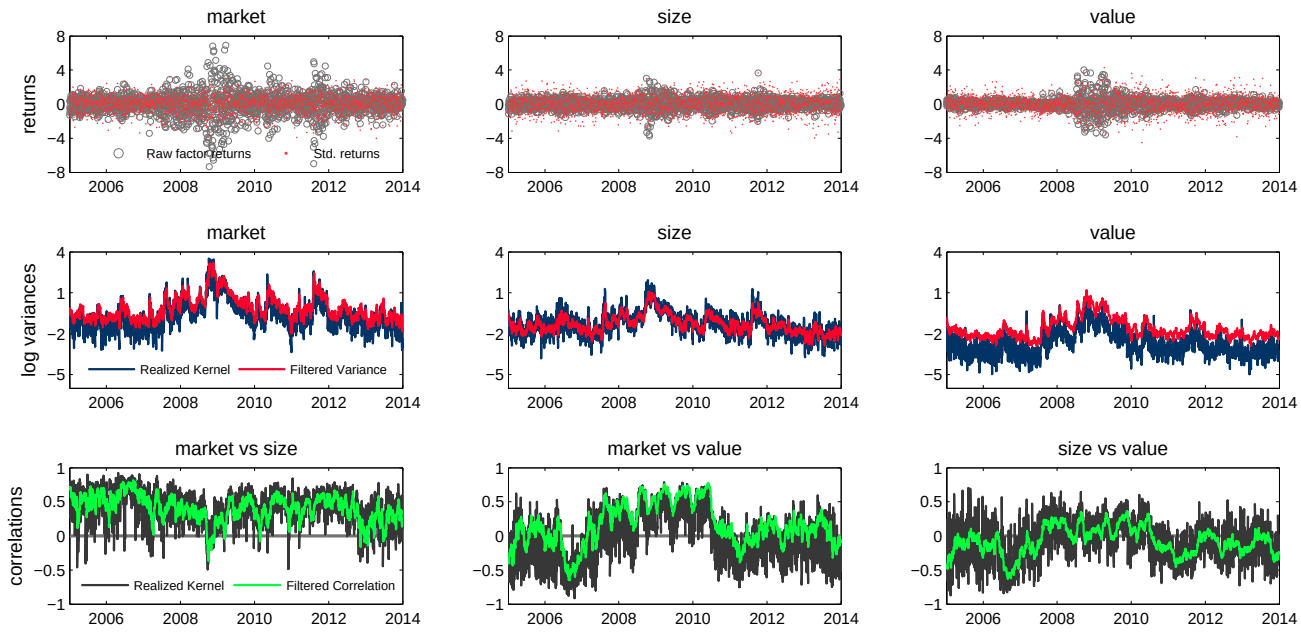

Figure 1.1: Variables filtered from the core model for the Fama-French factors (January, 2005 - December, 2013). In the top row, there are daily factor returns based on the Fama-French research portfolios and the standardized returns produced by the multivariate realized GARCH (core) model. In the middle row, there are realized kernel measures of factor variances constructed using intra-daily ETF data and the conditional variances filtered from the model. In the bottom row, there are realized kernel measures of correlations between the factors constructed using intra-daily ETF data and the conditional correlations filtered from the model.

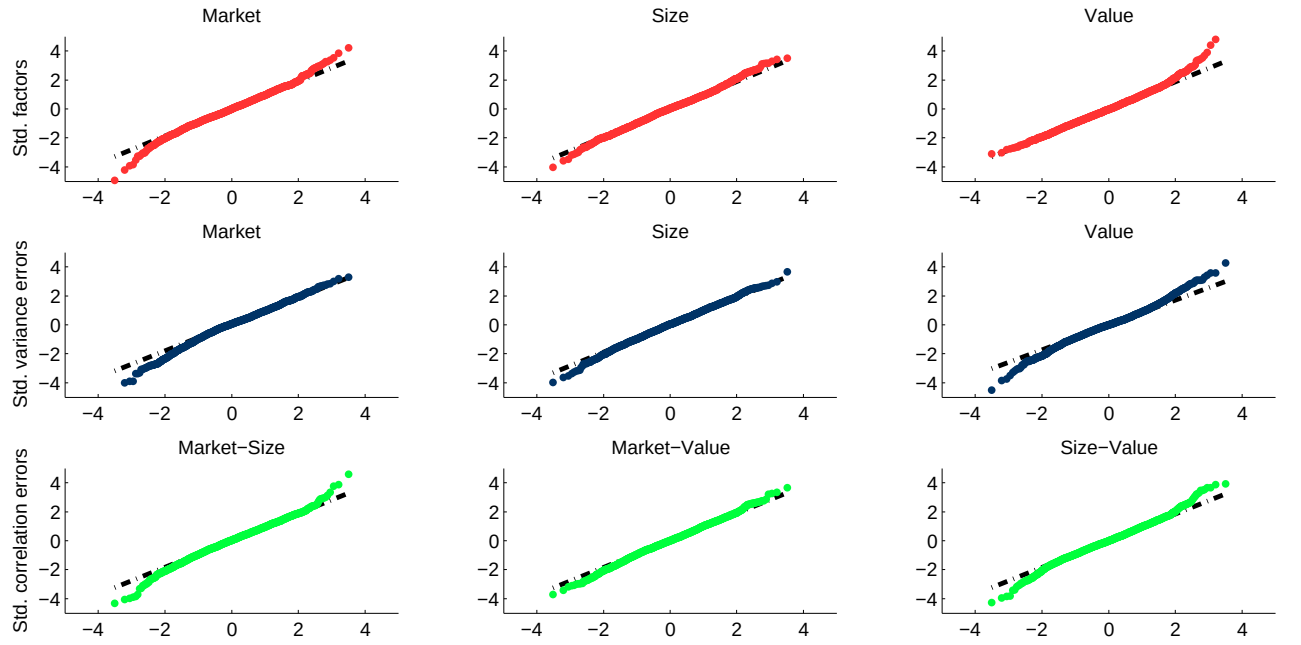

Figure 1.2: QQ-plots of residuals from the core model. In the top row, there are standardized daily factor returns, $\mathbf{z}_t$. In the middle row, there are standardized measurement errors $\mathbf{u}_t$ of the conditional variances. In the bottom row, there are standardized measurement errors $\mathbf{v}_t$ of the conditional correlations.

	GARCH equation						Measurement equation			
	μ_i	γ_0^i	γ_1^i	γ_2^i	τ_1^i	τ_2^i	ϕ_0^i	δ_1^i	δ_2^i	Ω_i^u
	Exchange traded funds									
IWD	0.026	0.118	0.770	0.205	-0.077	0.013	-0.600	-0.110	0.027	0.233
IWB	0.031	0.109	0.783	0.192	-0.089	0.012	-0.598	-0.110	0.029	0.263
IWF	0.037	0.109	0.763	0.212	-0.077	0.011	-0.542	-0.119	0.021	0.210
IWS	0.032	0.114	0.792	0.185	-0.076	0.015	-0.627	-0.100	0.029	0.261
IWR	0.036	0.111	0.789	0.187	-0.076	0.016	-0.597	-0.110	0.030	0.257
IWP	0.036	0.115	0.791	0.185	-0.078	0.015	-0.611	-0.099	0.034	0.275
IWN	0.027	0.113	0.749	0.229	-0.057	0.014	-0.455	-0.088	0.015	0.187
IWM	0.027	0.106	0.726	0.249	-0.063	0.017	-0.384	-0.099	0.006	0.148
IWO	0.036	0.117	0.744	0.230	-0.059	0.017	-0.455	-0.088	0.016	0.183
XLB	0.016	0.121	0.744	0.231	-0.058	0.021	-0.469	-0.098	0.030	0.181
XLE	0.028	0.135	0.716	0.255	-0.048	0.033	-0.440	-0.110	0.013	0.146
XLF	0.021	0.136	0.747	0.231	-0.060	0.025	-0.559	-0.075	0.034	0.206
XLI	0.033	0.107	0.769	0.207	-0.065	0.016	-0.506	-0.083	0.034	0.226
XLK	0.031	0.108	0.761	0.213	-0.055	0.017	-0.495	-0.110	0.023	0.208
XLP	0.030	0.062	0.749	0.224	-0.052	0.012	-0.359	-0.055	0.046	0.231
XLU	0.027	0.050	0.721	0.250	-0.027	0.024	-0.222	-0.059	0.059	0.218
XLV	0.035	0.087	0.736	0.237	-0.045	0.012	-0.415	-0.070	0.044	0.210
XLY	0.030	0.118	0.759	0.220	-0.063	0.013	-0.527	-0.088	0.035	0.204
	Cross-section of stocks									
mean	0.044	0.169	0.713	0.252	-0.027	0.006	-0.457	-0.020	0.065	0.278
st.dev.	0.036	0.095	0.061	0.048	0.018	0.009	0.152	0.026	0.021	0.083
skewness	-0.119	3.181	-1.912	0.884	0.938	1.227	-0.521	0.941	-0.199	1.275
kurtosis	5.382	22.003	18.859	11.444	8.058	5.726	4.867	5.995	3.595	5.598
min	-0.156	-0.009	0.084	0.058	-0.093	-0.014	-1.283	-0.105	-0.001	0.139
1	-0.046	0.037	0.520	0.137	-0.065	-0.008	-0.840	-0.075	0.010	0.160
5	-0.013	0.065	0.622	0.180	-0.052	-0.006	-0.699	-0.055	0.029	0.175
10	0.001	0.082	0.646	0.196	-0.047	-0.004	-0.649	-0.050	0.037	0.187
25	0.025	0.118	0.685	0.222	-0.038	-0.000	-0.549	-0.038	0.052	0.217
50	0.043	0.151	0.716	0.252	-0.028	0.004	-0.452	-0.024	0.067	0.265
75	0.066	0.195	0.750	0.278	-0.018	0.010	-0.360	-0.006	0.079	0.323
90	0.088	0.263	0.778	0.304	-0.007	0.018	-0.275	0.011	0.089	0.387
95	0.099	0.328	0.797	0.325	0.000	0.023	-0.226	0.025	0.095	0.420
99	0.139	0.503	0.843	0.380	0.028	0.033	-0.137	0.060	0.113	0.565
max	0.212	1.052	0.932	0.676	0.086	0.064	0.005	0.144	0.156	0.733

Table 1.5: Parameter estimates of the conditional factor models for the style and sectoral ETFs and for the cross-section of 824 stocks. Here we report parameters related to the variance dynamics of the considered assets. Estimation period is January 3, 2005 - December 31, 2013.

	Correlation with market						Correlation with size						Correlation with value					
	Λ_0^i	Λ_1^i	Λ_2^i	Ψ_0^i	Ψ_1^i	Ω_i^v	Λ_0^i	Λ_1^i	Λ_2^i	Ψ_0^i	Ψ_1^i	Ω_i^v	Λ_0^i	Λ_1^i	Λ_2^i	Ψ_0^i	Ψ_1^i	Ω_i^v
	Exchange traded funds																	
IWD	-0.001	0.766	0.234	0.273	0.867	0.092	-0.021	0.817	0.207	0.143	0.758	0.060	0.025	0.854	0.150	-0.144	0.900	0.067
IWB	0.114	0.804	0.144	-0.514	1.247	0.105	-0.018	0.814	0.204	0.141	0.772	0.061	0.035	0.847	0.151	-0.212	0.931	0.068
IWF	-0.013	0.810	0.191	0.296	0.880	0.095	-0.014	0.805	0.208	0.127	0.790	0.061	0.045	0.838	0.151	-0.293	0.979	0.071
IWS	0.006	0.788	0.226	0.264	0.785	0.096	-0.023	0.816	0.206	0.189	0.733	0.075	0.026	0.872	0.135	-0.173	0.867	0.083
IWR	-0.046	0.841	0.196	0.502	0.682	0.090	-0.025	0.825	0.197	0.202	0.739	0.075	0.031	0.863	0.142	-0.203	0.877	0.084
IWP	0.046	0.819	0.163	0.140	0.882	0.087	-0.025	0.822	0.208	0.200	0.703	0.079	0.033	0.869	0.134	-0.235	0.886	0.091
IWN	-0.085	0.789	0.298	0.460	0.598	0.055	-0.048	0.795	0.197	0.349	0.910	0.077	0.024	0.860	0.147	-0.143	0.876	0.070
IWM	-0.095	0.789	0.301	0.504	0.585	0.054	-0.061	0.802	0.205	0.400	0.838	0.077	0.030	0.867	0.141	-0.196	0.861	0.069
IWO	-0.094	0.799	0.288	0.513	0.582	0.055	-0.064	0.797	0.219	0.388	0.811	0.073	0.034	0.869	0.136	-0.246	0.872	0.070
XLB	-0.012	0.833	0.187	0.231	0.760	0.082	-0.009	0.848	0.169	0.115	0.735	0.061	0.028	0.866	0.134	-0.187	0.904	0.072
XLE	-0.005	0.806	0.210	0.103	0.844	0.081	-0.014	0.845	0.179	0.106	0.720	0.059	0.032	0.848	0.143	-0.194	0.965	0.069
XLF	0.007	0.814	0.192	0.199	0.798	0.080	-0.022	0.839	0.192	0.153	0.692	0.061	0.028	0.864	0.152	-0.154	0.823	0.077
XLI	-0.029	0.816	0.210	0.385	0.712	0.088	-0.011	0.837	0.181	0.126	0.735	0.064	0.030	0.864	0.143	-0.194	0.870	0.074
XLK	-0.089	0.819	0.235	0.570	0.628	0.075	-0.016	0.815	0.211	0.129	0.741	0.060	0.032	0.856	0.159	-0.200	0.826	0.075
XLP	0.039	0.801	0.162	0.017	0.962	0.094	-0.004	0.855	0.153	0.058	0.772	0.062	0.014	0.876	0.129	-0.098	0.846	0.072
XLU	0.034	0.826	0.143	-0.085	1.003	0.086	-0.003	0.885	0.112	0.049	0.810	0.063	0.018	0.881	0.092	-0.128	1.047	0.070
XLV	0.032	0.803	0.170	-0.000	0.984	0.098	-0.004	0.860	0.149	0.077	0.740	0.066	0.016	0.871	0.128	-0.122	0.893	0.073
XLY	-0.023	0.808	0.215	0.358	0.721	0.089	-0.010	0.825	0.195	0.115	0.747	0.065	0.030	0.859	0.154	-0.188	0.842	0.072
	Cross-section of stocks																	
mean	0.012	0.840	0.146	0.031	0.916	0.075	-0.001	0.887	0.109	0.073	0.813	0.067	0.013	0.906	0.097	-0.116	0.839	0.073
st.dev.	0.020	0.035	0.048	0.155	0.277	0.006	0.007	0.022	0.030	0.063	0.181	0.007	0.005	0.018	0.022	0.035	0.100	0.004
skewness	-0.266	0.389	0.113	-2.089	2.760	0.300	0.039	0.213	0.061	-0.681	1.336	0.656	0.368	0.380	-0.094	0.149	0.460	0.455
kurtosis	3.582	3.787	3.232	14.427	16.615	3.082	3.576	4.054	2.992	5.418	6.133	2.813	2.987	3.423	2.690	2.804	3.992	3.848
min	-0.056	0.711	0.018	-1.273	0.422	0.054	-0.027	0.811	0.027	-0.242	0.416	0.052	0.001	0.845	0.041	-0.225	0.501	0.061
1	-0.044	0.768	0.037	-0.522	0.557	0.062	-0.019	0.831	0.039	-0.130	0.530	0.055	0.003	0.866	0.047	-0.189	0.620	0.064
5	-0.020	0.787	0.068	-0.224	0.623	0.065	-0.013	0.852	0.060	-0.034	0.587	0.057	0.005	0.880	0.061	-0.172	0.696	0.066
10	-0.012	0.799	0.087	-0.143	0.661	0.067	-0.010	0.861	0.070	-0.000	0.625	0.059	0.006	0.884	0.068	-0.161	0.721	0.068
25	0.001	0.819	0.113	-0.037	0.743	0.070	-0.005	0.874	0.088	0.040	0.690	0.062	0.009	0.894	0.082	-0.140	0.773	0.070
50	0.013	0.838	0.147	0.045	0.861	0.074	-0.001	0.887	0.107	0.076	0.780	0.066	0.013	0.905	0.099	-0.117	0.835	0.073
75	0.024	0.861	0.176	0.129	1.012	0.079	0.004	0.900	0.129	0.111	0.905	0.072	0.017	0.916	0.113	-0.093	0.897	0.075
90	0.037	0.884	0.205	0.186	1.212	0.083	0.009	0.915	0.146	0.148	1.047	0.077	0.020	0.929	0.124	-0.071	0.967	0.078
95	0.044	0.900	0.225	0.231	1.376	0.086	0.011	0.924	0.157	0.174	1.176	0.081	0.023	0.938	0.130	-0.059	1.028	0.080
99	0.056	0.939	0.266	0.313	2.057	0.092	0.016	0.953	0.180	0.214	1.399	0.086	0.027	0.955	0.144	-0.031	1.109	0.086
max	0.073	0.968	0.312	0.366	3.208	0.096	0.027	0.969	0.204	0.257	1.963	0.090	0.032	0.967	0.163	-0.015	1.237	0.090

Table 1.6: Parameter estimates of the conditional factor models for the style and sectoral ETFs and for the cross-section of 824 stocks. Here we report parameters related to the dynamics of correlations between a considered asset and the three Fama-French factors. Estimation period is January 3, 2005 - December 31, 2013.

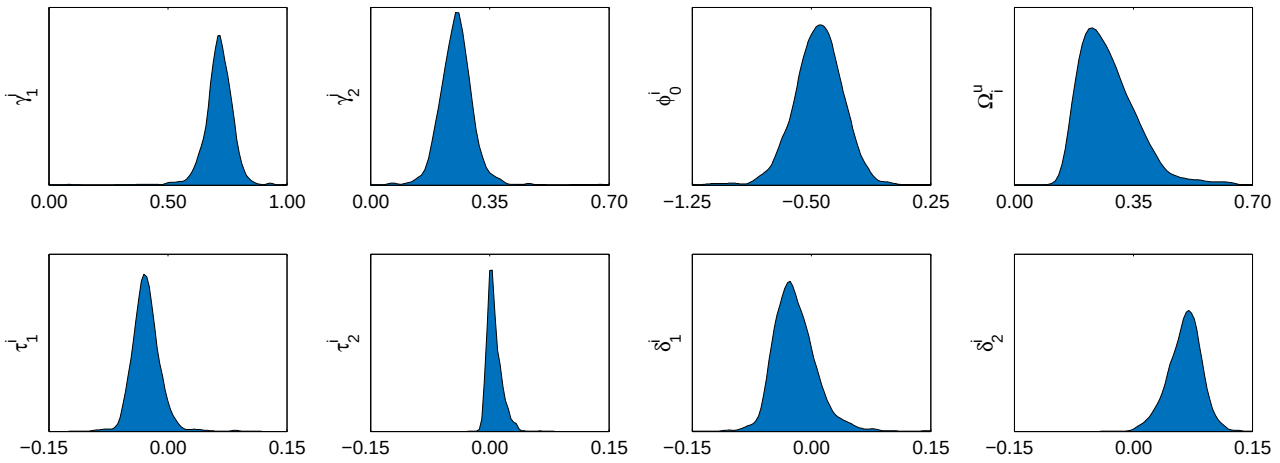

(a) Parameters related to variance dynamics

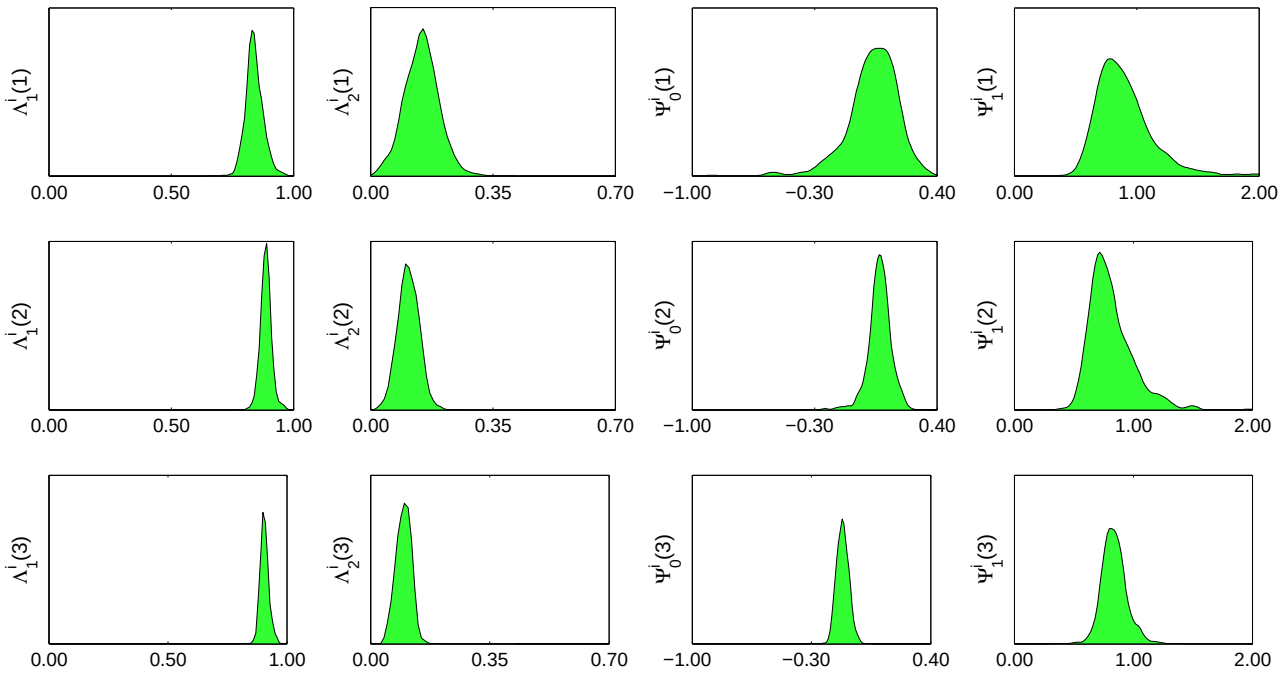

(b) Parameters related to correlation dynamics

Figure 1.3: Histograms of the estimated parameters from the conditional models for the cross-section of 824 stocks. Estimation period is January 3, 2005 - December 31, 2013.

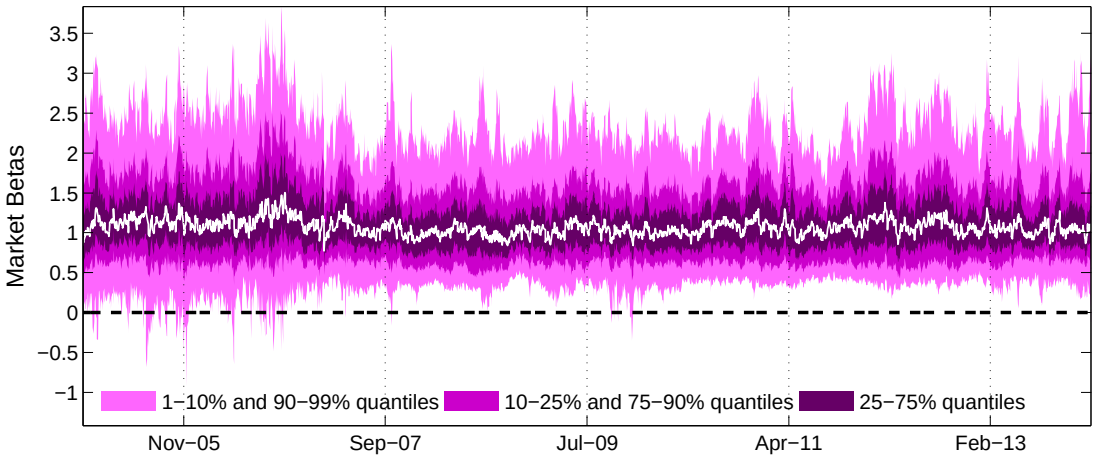

(a) Market loadings.

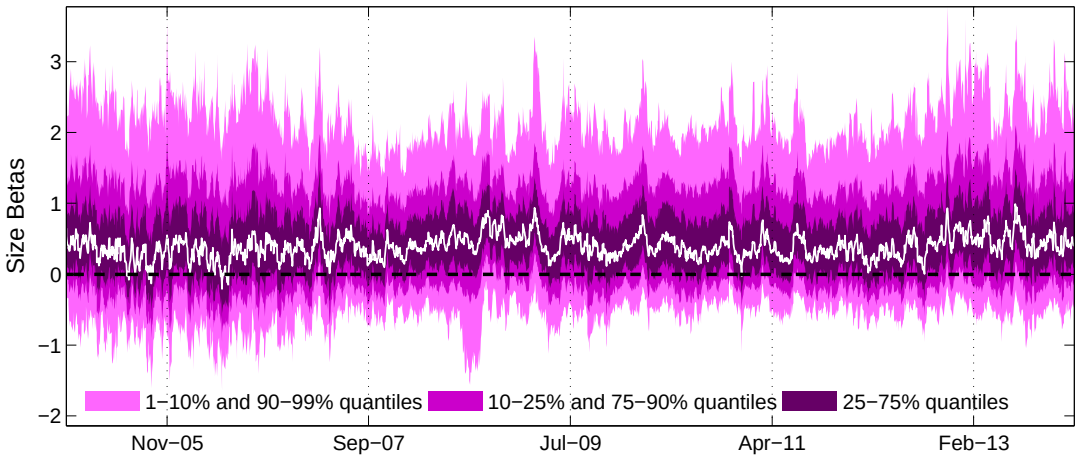

(b) Size loadings.

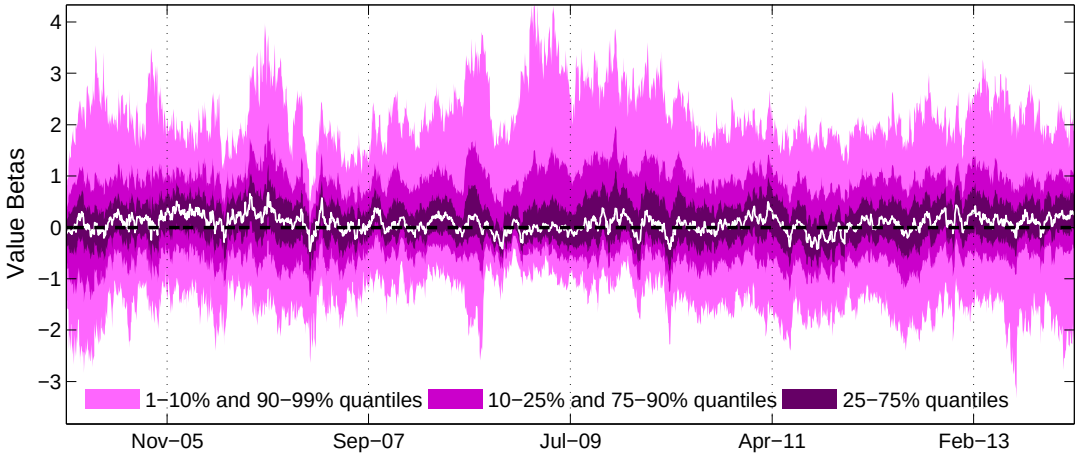

(c) Value loadings.

Figure 1.4: Quantile plots of the cross-section distributions of the model-implied factor loadings. Estimation period is January 3, 2005 - December 31, 2013.

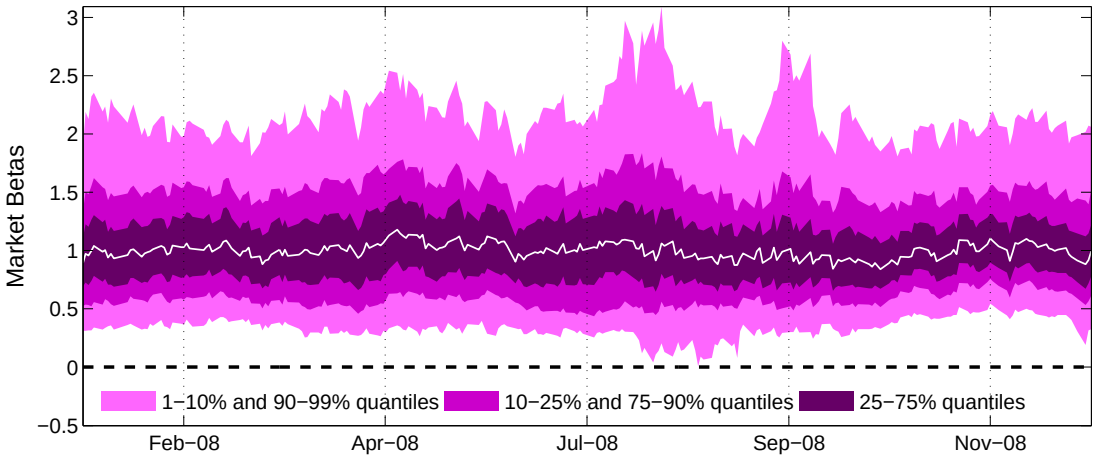

(a) Market loadings.

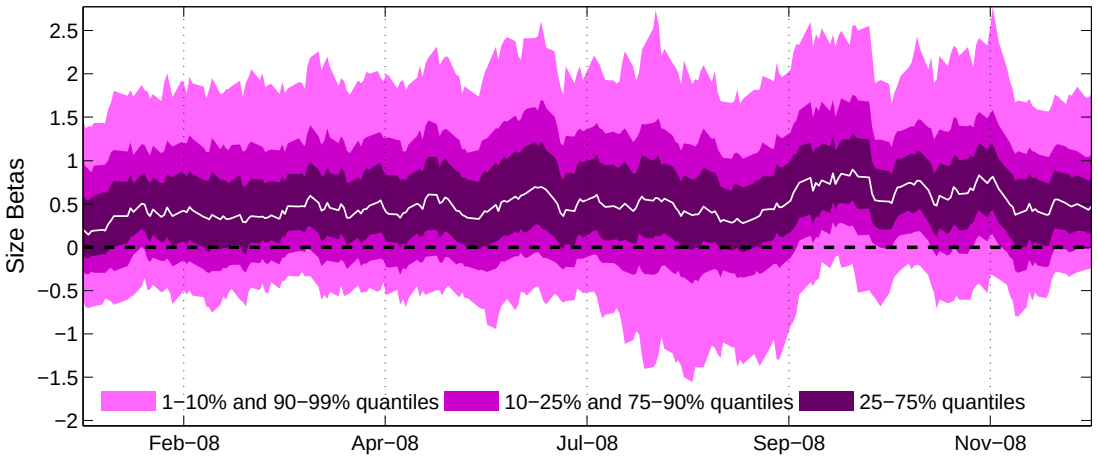

(b) Size loadings.

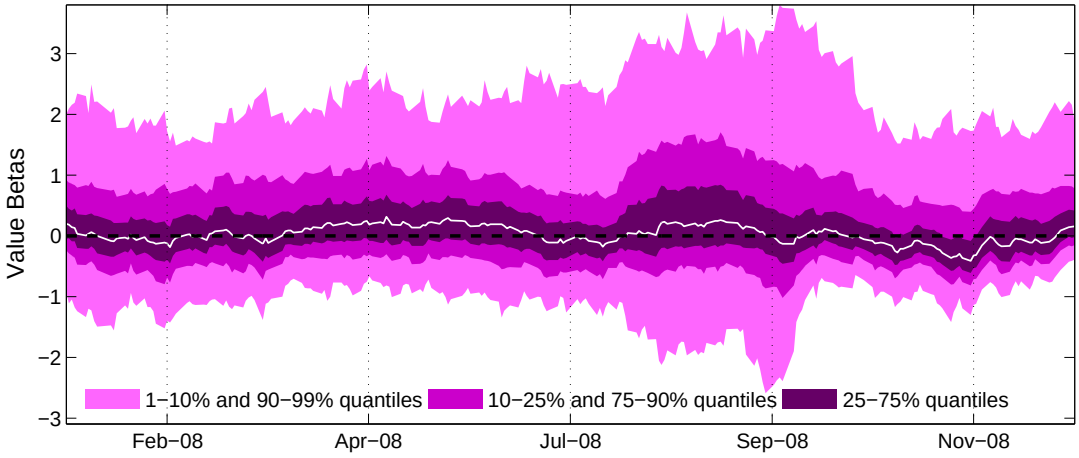

(c) Value loadings.

Figure 1.5: Quantile plots of the cross-section distributions of the model-implied factor loadings in 2008.

	2005	2006	2007	2008	2009	2010	2011	2012	2013
					Market				
mean	1.160	1.227	1.088	1.027	1.070	1.084	1.089	1.153	1.098
std.	0.292	0.360	0.287	0.308	0.301	0.309	0.341	0.394	0.311
skewness	0.360	0.541	0.473	0.470	0.449	0.486	0.370	0.639	0.948
kurtosis	3.571	2.921	3.617	3.490	4.185	3.345	2.888	3.738	4.208
min	0.100	0.426	0.361	-0.033	-0.291	0.402	0.395	0.386	0.514
1	0.503	0.587	0.514	0.442	0.522	0.500	0.452	0.462	0.579
5	0.722	0.720	0.628	0.548	0.629	0.623	0.564	0.564	0.700
10	0.823	0.814	0.752	0.643	0.719	0.696	0.636	0.652	0.745
25	0.967	0.945	0.888	0.827	0.867	0.865	0.846	0.864	0.868
50	1.140	1.182	1.066	1.009	1.051	1.071	1.068	1.125	1.052
75	1.326	1.459	1.263	1.203	1.251	1.277	1.303	1.393	1.285
90	1.543	1.714	1.456	1.432	1.445	1.473	1.539	1.651	1.501
95	1.692	1.881	1.599	1.576	1.625	1.586	1.693	1.832	1.673
99	1.905	2.154	1.881	1.836	1.951	1.987	1.989	2.283	2.063
max	2.255	2.483	2.328	2.267	2.230	2.231	2.265	2.755	2.462
					Size				
mean	0.432	0.412	0.421	0.521	0.474	0.502	0.443	0.472	0.581
std.	0.500	0.532	0.463	0.488	0.492	0.528	0.499	0.528	0.567
skewness	0.527	0.475	0.428	0.263	0.336	0.527	0.588	0.582	0.520
kurtosis	2.867	2.593	2.735	2.655	2.740	2.498	2.665	2.755	2.764
min	-0.501	-0.555	-0.502	-0.715	-0.761	-0.427	-0.387	-0.497	-0.637
1	-0.386	-0.452	-0.397	-0.459	-0.449	-0.337	-0.356	-0.370	-0.364
5	-0.262	-0.337	-0.244	-0.217	-0.267	-0.231	-0.239	-0.248	-0.213
10	-0.195	-0.248	-0.157	-0.083	-0.169	-0.137	-0.152	-0.144	-0.096
25	0.038	-0.016	0.071	0.146	0.108	0.083	0.049	0.039	0.128
50	0.378	0.346	0.365	0.485	0.445	0.410	0.344	0.389	0.512
75	0.784	0.803	0.767	0.889	0.806	0.884	0.801	0.836	0.942
90	1.135	1.142	1.050	1.162	1.124	1.282	1.148	1.214	1.395
95	1.289	1.345	1.199	1.308	1.336	1.455	1.354	1.399	1.608
99	1.750	1.751	1.574	1.695	1.675	1.827	1.726	1.946	1.966
max	2.458	2.155	2.282	2.193	2.154	2.128	2.128	2.273	2.474
					Value				
mean	0.117	0.240	0.082	0.136	0.227	0.110	0.026	0.120	0.188
std.	0.547	0.528	0.392	0.548	0.678	0.588	0.589	0.569	0.597
skewness	0.899	0.609	0.281	1.796	1.803	1.061	0.312	0.056	0.485
kurtosis	5.077	4.224	3.673	8.411	6.960	4.672	3.946	3.942	6.482
min	-1.453	-1.469	-1.167	-1.105	-0.823	-1.213	-1.898	-2.099	-2.503
1	-0.992	-0.843	-0.816	-0.706	-0.683	-0.969	-1.493	-1.449	-1.447
5	-0.742	-0.580	-0.553	-0.467	-0.459	-0.664	-0.874	-0.766	-0.622
10	-0.519	-0.407	-0.401	-0.366	-0.367	-0.522	-0.678	-0.556	-0.409
25	-0.183	-0.077	-0.156	-0.185	-0.213	-0.278	-0.308	-0.206	-0.147
50	0.093	0.219	0.064	-0.000	0.014	0.015	-0.010	0.088	0.118
75	0.347	0.494	0.315	0.332	0.433	0.380	0.314	0.465	0.517
90	0.696	0.803	0.605	0.873	1.091	0.881	0.754	0.857	0.900
95	1.237	1.332	0.744	1.212	1.695	1.262	1.200	1.117	1.145
99	1.849	1.802	1.055	2.022	2.691	2.036	1.631	1.516	2.059
max	2.179	2.238	1.713	3.818	3.667	2.729	2.099	2.107	3.438

Table 1.7: Descriptive statistics for the model-implied factor loadings reported by years.

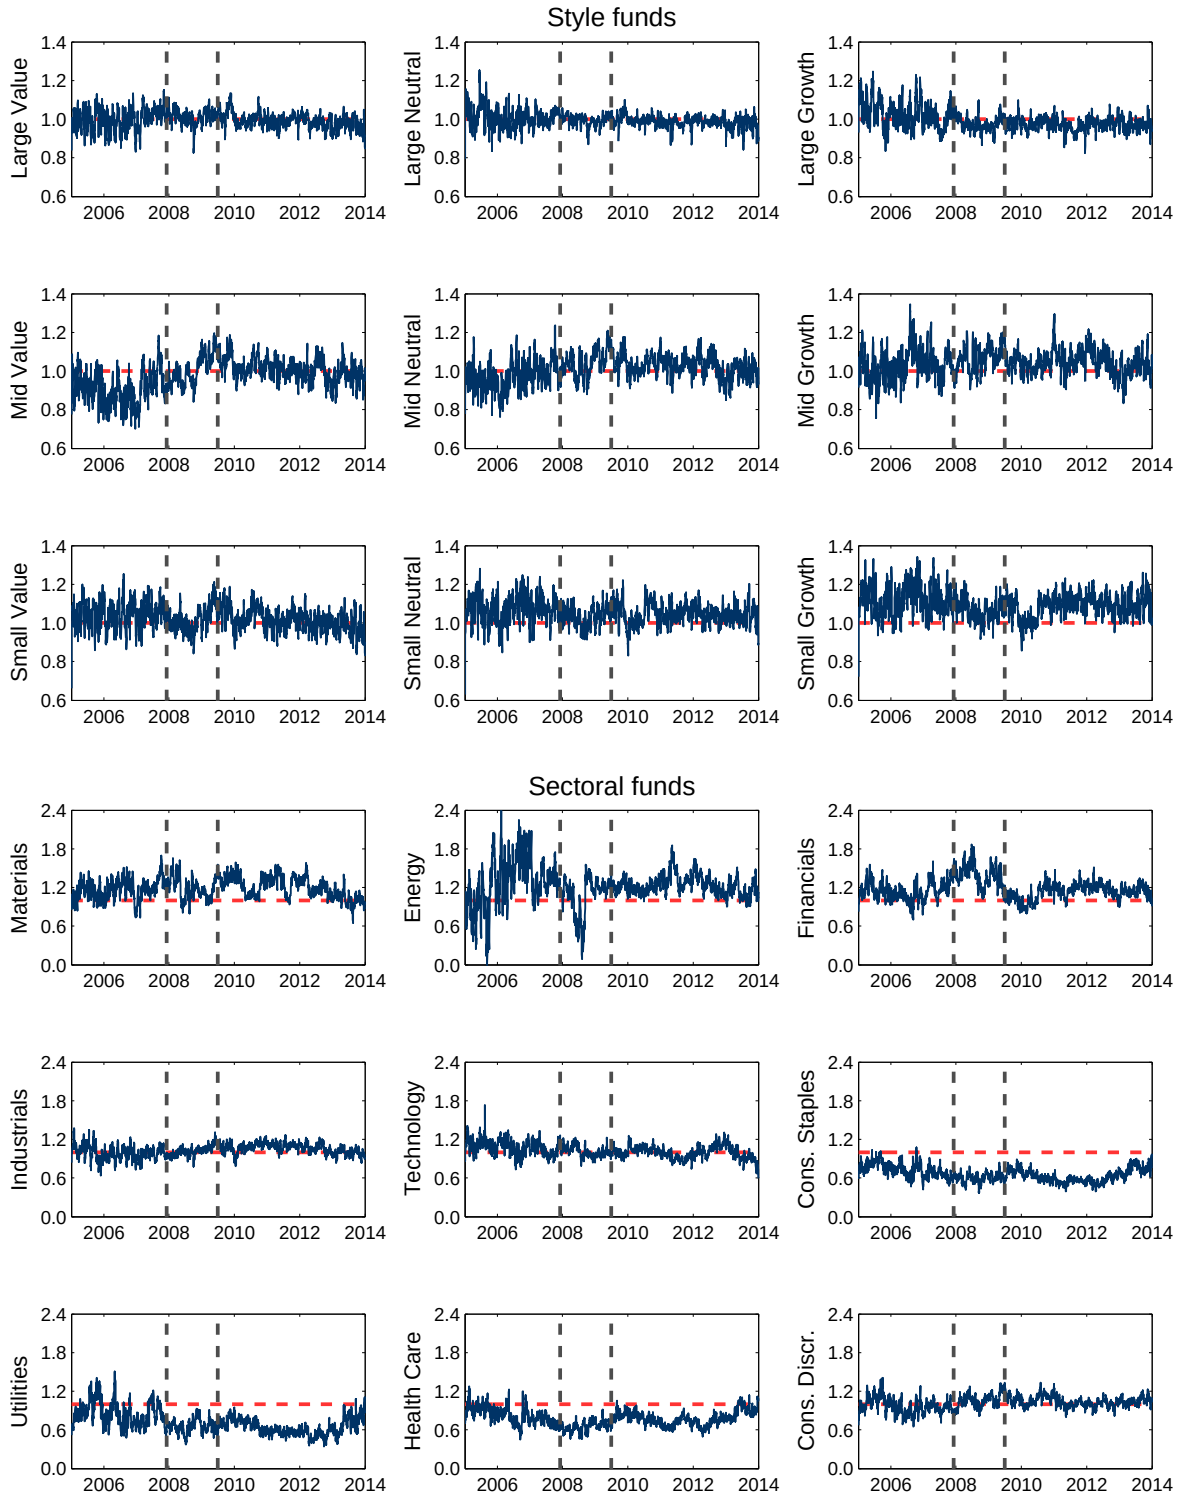

Figure 1.6: Model-implied market betas for the style and sectoral ETFs. Estimation period is January 3, 2005 - December 31, 2013.

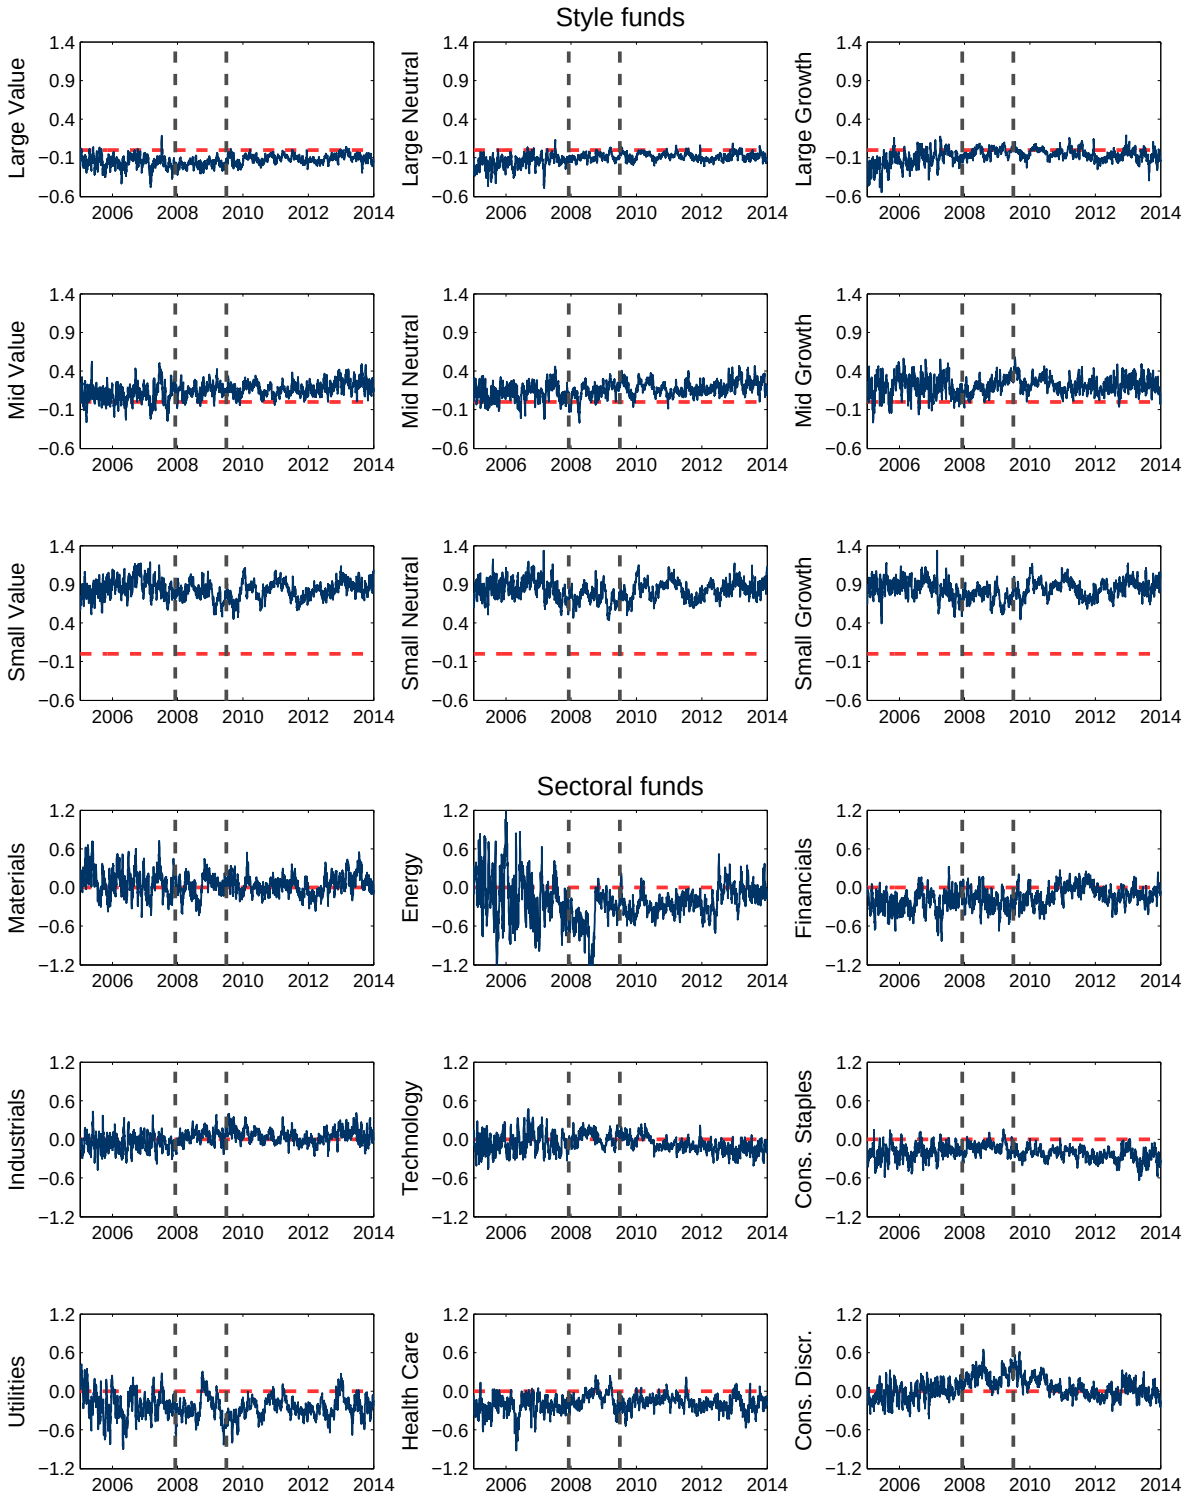

Figure 1.7: Model-implied size betas for the style and sectoral ETFs. Estimation period is January 3, 2005 - December 31, 2013.

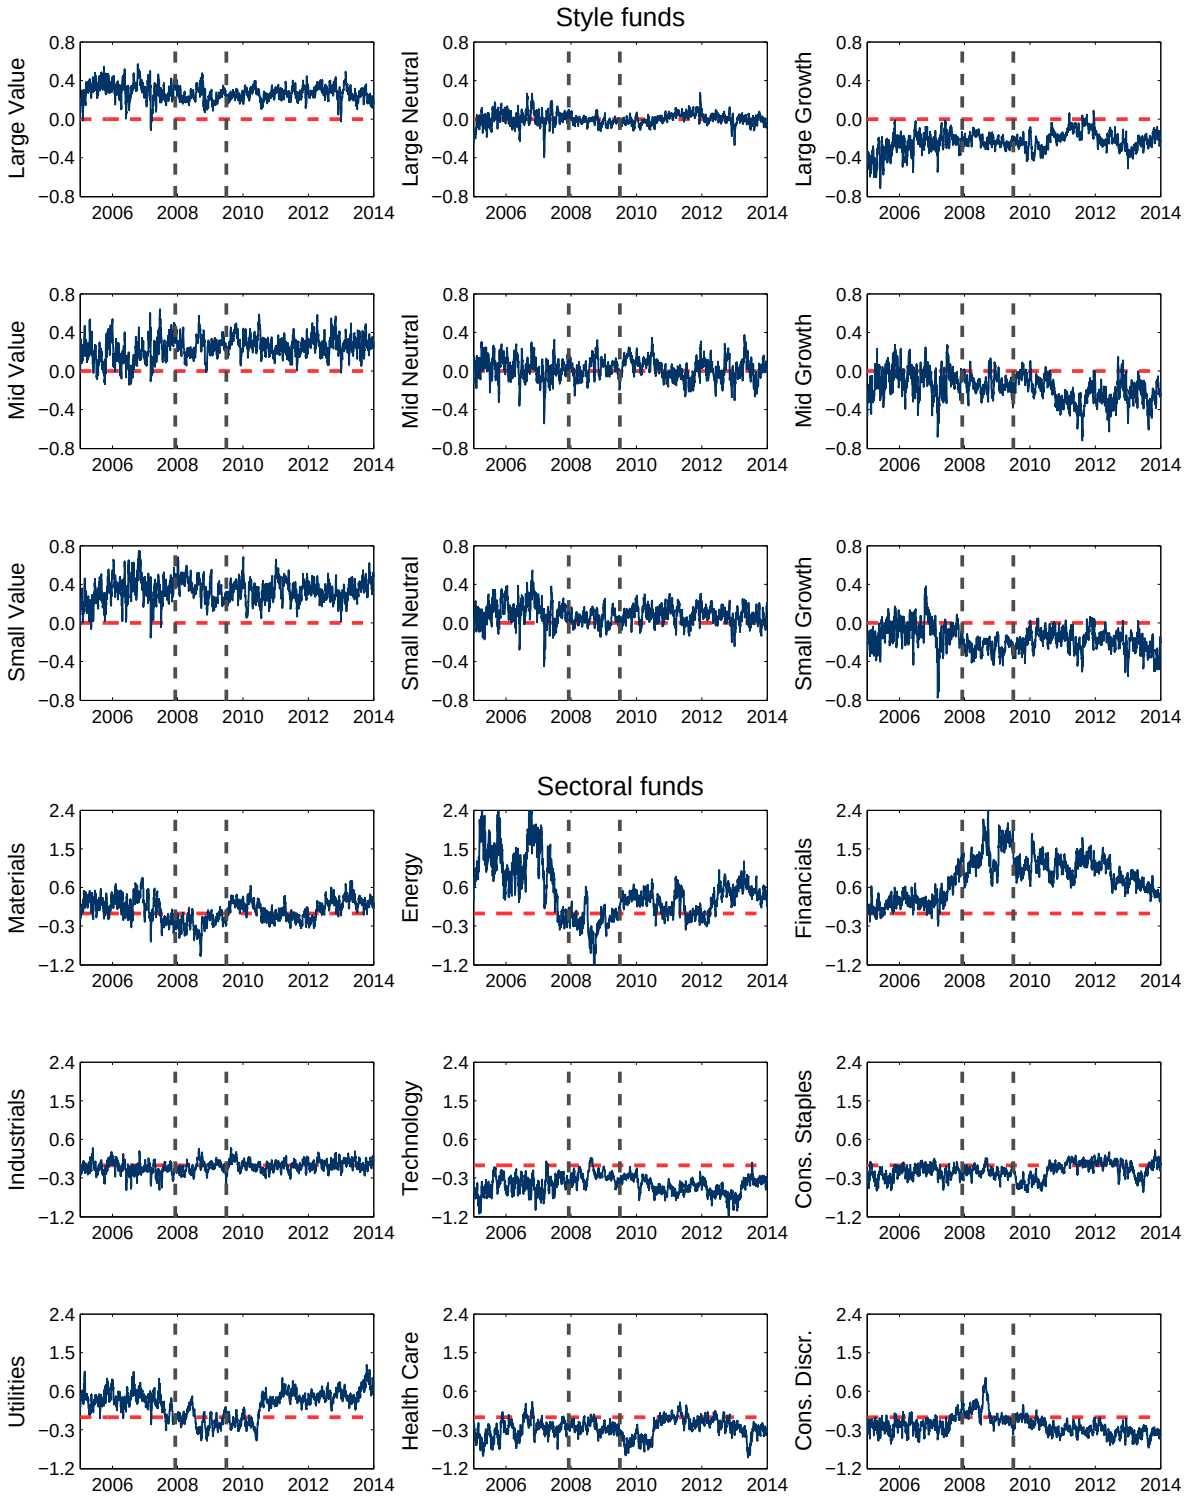

Figure 1.8: Model-implied value betas for the style and sectoral ETFs. Estimation period is January 3, 2005 - December 31, 2013.

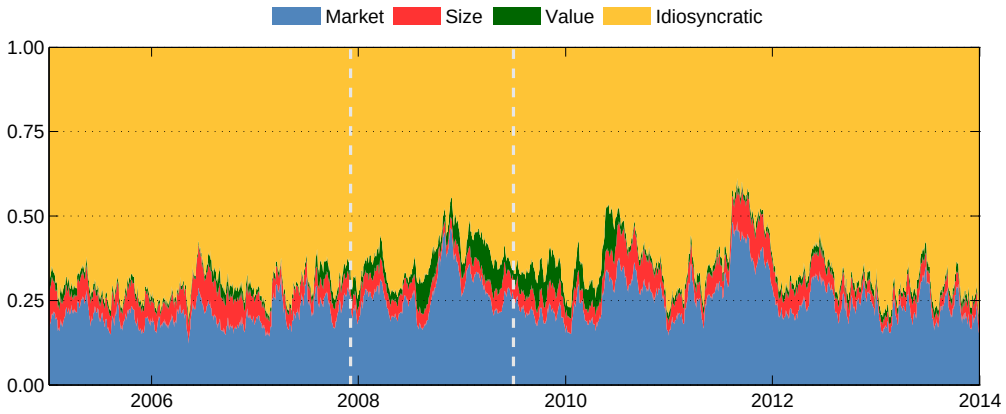

Figure 1.9: Return variance decomposition along the Fama-French risk factors and the idiosyncratic component for the cross-section of 824 stocks. The decomposition is implemented according to the method in Genizi (1993). Sample period is January 3, 2005 - December 31, 2013.

	2005	2006	2007	2008	2009	2010	2011	2012	2013
					Share of variance by the market factor				
mean	0.202	0.196	0.226	0.273	0.246	0.262	0.312	0.248	0.223
min	0.020	0.025	0.034	0.019	0.008	0.024	0.048	0.022	0.015
5	0.072	0.086	0.092	0.123	0.094	0.112	0.131	0.089	0.080
25	0.136	0.144	0.167	0.207	0.178	0.198	0.233	0.175	0.160
50	0.195	0.193	0.220	0.271	0.248	0.260	0.310	0.239	0.218
75	0.256	0.241	0.279	0.335	0.312	0.322	0.385	0.315	0.283
95	0.374	0.333	0.380	0.433	0.397	0.417	0.503	0.430	0.387
max	0.511	0.433	0.493	0.541	0.528	0.516	0.590	0.535	0.533
					Share of variance by the size factor				
mean	0.066	0.079	0.055	0.047	0.053	0.069	0.076	0.056	0.039
min	0.011	0.013	0.006	0.006	0.006	0.011	0.015	0.007	0.003
5	0.026	0.026	0.016	0.010	0.015	0.018	0.023	0.014	0.008
25	0.041	0.049	0.032	0.023	0.031	0.037	0.046	0.030	0.016
50	0.059	0.070	0.048	0.039	0.050	0.063	0.070	0.050	0.033
75	0.084	0.101	0.072	0.065	0.072	0.093	0.099	0.076	0.055
95	0.136	0.158	0.115	0.107	0.107	0.146	0.152	0.124	0.096
max	0.231	0.236	0.171	0.158	0.137	0.206	0.220	0.184	0.141
					Share of variance by the value factor				
mean	0.020	0.021	0.024	0.049	0.072	0.047	0.020	0.021	0.020
min	0.003	0.004	0.003	0.007	0.012	0.006	0.002	0.002	0.002
5	0.006	0.008	0.008	0.014	0.022	0.015	0.004	0.004	0.005
25	0.010	0.014	0.013	0.027	0.039	0.027	0.007	0.008	0.007
50	0.014	0.019	0.019	0.037	0.058	0.039	0.011	0.012	0.011
75	0.021	0.026	0.027	0.056	0.094	0.056	0.024	0.024	0.024
95	0.061	0.044	0.064	0.139	0.160	0.110	0.064	0.071	0.068
max	0.139	0.078	0.148	0.279	0.316	0.258	0.140	0.142	0.124

Table 1.8: Descriptive statistics for the return variance decomposition along the Fama-French risk factors for the cross-section of 824 stocks. Statistics is reported by years. The variance decomposition is implemented according to the method in Genizi (1993).

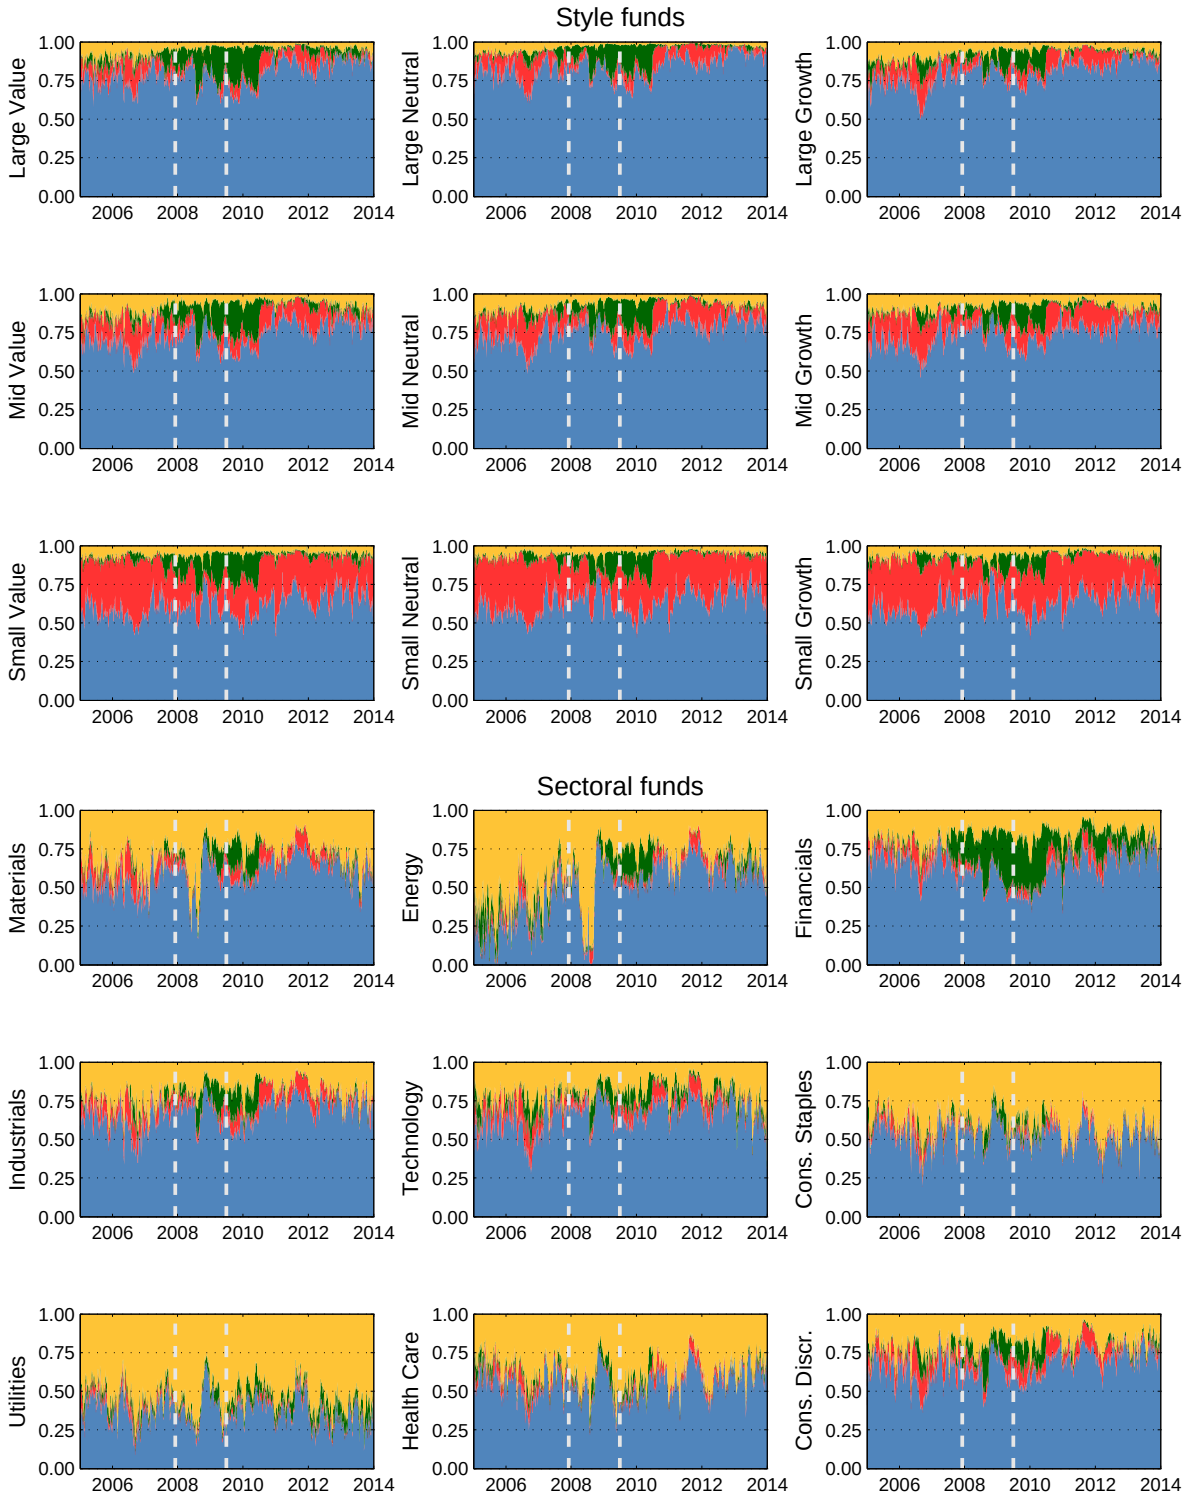

Figure 1.10: Return variance decomposition along the Fama-French risk factors and the idiosyncratic component for the style and sectoral ETFs. The decomposition is implemented according to the method in Genizi (1993). Sample period is January 3, 2005 - December 31, 2013.

Appendix A

Supplement to Chapter I

A.1 Proofs

Lemma 1. *Suppose that R is a non-singular correlation matrix. Then, the diagonal elements of logm(R) are bounded above, but the off-diagonal elements are unbounded on $\mathbb{R}$.*

Proof. Suppose that R is $k \times k$ correlation matrix, $k \geq 1$. The spectral decomposition of R reads

$$R = V\Lambda V' \tag{A.1.1}$$

where the columns of the $k \times k$ matrix V are eigenvectors of R and Λ is a diagonal matrix with the corresponding eigenvalues. In what follows, we will suppose that the columns of V are chosen to constitute an orthonormal basis. The logarithm transformation of R is a symmetric matrix which reads

$$\tilde{R} \equiv \text{logm}(R) = V \begin{pmatrix} \log\lambda_1 & 0 & \cdot & 0 \\ 0 & \log\lambda_2 & \cdot & 0 \\ \cdot & \cdot & \cdot & \cdot \\ 0 & 0 & \cdot & \log\lambda_k \end{pmatrix} V' \tag{A.1.2}$$

We denote by e_n the n-th coordinate (unit) vector. Then the n-th eigenvector can be expressed as $v^{(n)} = Ve_n$ for $n = 1, ..., k$. A generic (i, j)-element of $\tilde{R}$ can be written as follows

$$\tilde{R}_{i,j} = \sum_{n=1}^{k} v_i^{(n)} v_j^{(n)} \log \lambda_n \tag{A.1.3}$$

where $v_i^{(n)}$ is an i-th element of an eigenvector $v^{(n)}$. Note that since R is positive definite matrix, all λ_n are real and strictly positive.

We start with proving the boundedness of the diagonal elements from $\tilde{R}$, i.e. of the elements (A.1.3) with $i = j$. Consider a generic diagonal element of $\tilde{R}$. From Jensen's inequality it follows that

$$\tilde{R}_{i,i} = \sum_{n=1}^{k} (e_i' V e_n)^2 \log \lambda_n = \sum_{n=1}^{k} (e_i' v^{(n)})^2 \log \lambda_n \leq \log\big(\sum_{n=1}^{k} (e_i' v^{(n)})^2 \lambda_n\big) \tag{A.1.4}$$

Here we used the fact that $\sum_{n=1}^{k} (e_i' v^{(n)})^2 = 1$, for $i = 1, ..., k$, since the matrix V is orthogonal. Since we know that diagonal elements of a correlation matrix are equal to one, $R_{i,i} = \sum_{n=1}^{k} (e_i' v^{(n)})^2 \lambda_n = 1$, diagonal elements of $\tilde{R}$ can be bounded from above as follows

$$\tilde{R}_{i,i} \leq \log(1) = 0 \tag{A.1.5}$$

for $i = 1, ..., k$.

Consider now off-diagonal elements of $\tilde{R}$. To demonstrate that some off-diagonal element $\tilde{R}_{i,j},\ i \neq j$ is unbounded on $\mathbb{R}$, lets suppose that the correlation between i-th and j-th elements of the multivariate return process is close to perfect, i.e. $R_{i,j} \to 1$. Then, the spectrum of R contains an eigenvalue that is close to 0. Without loss of generality, we assign index h to this eigenvalue. Since eigenvalues are continuous functions of the matrix entries, it follows that

$$\lim_{R_{i,j} \to 1} \lambda_h = 0 \tag{A.1.6}$$

and hence, $\log \lambda_h \to -\infty$. The corresponding eigenvector, $v^{(h)}$, converges to the vector with two non-zero elements in entries i and j and zeroes in other entries. For an orthonormal eigensystem $\{v^{(n)}\}_{n=1}^{k}$, entries $v_i^{(h)}$ and $v_j^{(h)}$ converge to the values $\frac{1}{\sqrt{2}}$ and $-\frac{1}{\sqrt{2}}$, so we have

$$\lim_{R_{(i,j)} \to 1} v_i^{(h)} v_j^{(h)} = -\frac{1}{2} \tag{A.1.7}$$

Using (A.1.3), we can express an (i, j)-th entry (for $i \neq j$) of $\tilde{R}$ as

$$\tilde{R}_{i,j} = \sum_{n \neq h} v_i^{(n)} v_j^{(n)} \log \lambda_n + v_i^{(h)} v_j^{(h)} \log \lambda_h \tag{A.1.8}$$

As long as the remaining eigenvalues λ_n for $n \neq h$ are fixed (in other words, if all the entries of R apart from (i, j)-th are fixed), the first term in (A.1.8) is bounded. Thus, from (A.1.6) and (A.1.7) we have that

$$\lim_{R_{i,j} \to 1} \tilde{R}_{i,j} = \infty \tag{A.1.9}$$

As a result, a generic off-diagonal entry of $\tilde{R}$ is unbounded from above.

Similarly, if the correlation between i-th and j-th elements of the multivariate return process is close to be perfectly negative, i.e. $R_{i,j} \to -1$, then the spectrum of R again contains an eigenvalue which is close to 0. Denoting the index of such eigenvalue by h, we have that $\log \lambda \to -\infty$ and $v_i^{(h)} v_j^{(h)} \to \frac{1}{2}$, as $R_{i,j} \to -1$. Then, from (A.1.8) it follows that

$$\lim_{R_{i,j} \to -1} \tilde{R}_{i,j} = -\infty \tag{A.1.10}$$

so a generic off-diagonal entry of $\tilde{R}$ is unbounded from below as well. □

Lemma 2. *Suppose that R is a $(k+1) \times (k+1)$ symmetric matrix with ones on the main diagonal. Let the upper left $k \times k$ block of R consists of a non-singular correlation matrix R^c. Denote the upper right $k \times 1$ block of R by p. Then, R is a non-singular correlation matrix if and only if $p'(R^c)^{-1}p < 1$.*

Proof. Since R has ones on the main diagonal, in order to be a full rank correlation matrix, R has to be positive definite. The Schur complement of the upper left block R^c of R is given by

$$R^+ = 1 - p'\left(R^c\right)^{-1} p \tag{A.1.11}$$

According to the well known result from the linear algebra, provided that R^c is a positive definite matrix, R is positive definite if and only if $R^+ > 0$. Since matrix R^c is positive definite by assumption, R is a non-singular correlation matrix if and only if $p'\left(R^c\right)^{-1} p < 1$. □

A.2 Modeling Correlations between Asset and Factors through the Spherical Coordinate Transform

In this extension, we suggest an alternative approach to modeling a vector of conditional correlations between the individual asset returns and k factors. The $k \times 1$ vector of unobservable conditional correlations is denoted by ρ_t^i, whereas an observed vector of corresponding realized measures is denoted by $\mathbf{y}_t^i$. Equations (1.15)-(1.16) suggest the joint dynamics for $\mathbf{y}_t^i$ and ρ_t^i, but do not ensure positive definiteness of the correlation matrix $R_t^{c,i}$ which is given by

$$R_t^{c,i} = \begin{pmatrix} R_t^c & \rho_t^i \\ (\rho_t^i)' & 1 \end{pmatrix}$$

where R_t^c is a $k \times k$ correlation matrix that is obtained from the core model and is treated as exogenous.

At first, we introduce the $k \times 1$ vector $\tilde{\rho}_t^i = (R_t^c)^{-\frac{1}{2}} \rho_t^i$, where $(R_t^c)^{\frac{1}{2}}$ is a principal square root of R_t^c. Since R_t^c is an exogenous matrix process, we may model $\tilde{\rho}_t^i$ instead of ρ_t^i and then recover the latter by simply multiplying $\tilde{\rho}_t^i$ and $(R_t^c)^{\frac{1}{2}}$. Suppose a realized correlation matrix $Y_t^{c,i}$ is available and Y_t^c stands for the upper left $k \times k$ block which describes the correlation among the common factors and the upper right $k \times 1$ block is $\mathbf{y}_t^i$

$$Y_t^{c,i} = \begin{pmatrix} Y_t^c & \mathbf{y}_t^i \\ (\mathbf{y}_t^i)' & 1 \end{pmatrix}$$

Then, the proper realized measure for $\tilde{\rho}_t^i$ can be analogously obtained as $\tilde{\mathbf{y}}_t^i = \left(Y_t^c\right)^{-\frac{1}{2}} \mathbf{y}_t^i$. Note that matrix $Y_t^{c,i}$ is supposed to be a correlation matrix by construction.

From Lemma 2 it follows that given R_t^c is a correlation matrix, $(\tilde{\rho}_t^i)' \tilde{\rho}_t^i < 1$ is necessary and sufficient requirement on $\tilde{\rho}_t^i$ to guarantee that $R_t^{c,i}$ remains a correlation (positive definite) matrix. In order to satisfy this requirement we transform the correlation vector $\tilde{\rho}_t^i$ providing a mapping from the Cartesian to spherical coordinates.

In a general case of multidimensional k-sphere, a spherical coordinate system consists of a radial coordinate (radial distance) and $k-1$ angular coordinates. Imagine now that $\tilde{\rho}_t^i$ is a Cartesian coordinate vector. Then, the radial distance in spherical coordinates is defined by

$$q_t = \sqrt{(\tilde{\rho}_t^i)' \tilde{\rho}_t^i} \tag{A.2.12}$$

The radial distance can be interpreted as an aggregate measure of the correlation between an individual asset and the considered system of k factors. To preserve positive definiteness of $R_t^{c,i}$ we must require that $q_t \in (0,1)$. Therefore, for example, we may suggest to model the dynamics of the logistic transform, $\mathrm{Lgt}(q_t)$, which maps from $(0,1)$ to $\mathbb{R}$. The $k-1$ angular coordinates $\theta_{1,t}$,$\theta_{2,t}$,...,$\theta_{k-2,t}$,φ_t ($k-2$ polar angles and 1 azimuth angle) are then constructed using standard formulae for k-spherical coordinate transform. All angular coordinates are defined on bounded intervals, so the mapping onto $\mathbb{R}$ can be applied as in the case of the radial distance q_t.

Denote vector $\tilde{\rho}_t^i$ in spherical coordinates as

$$\tilde{\rho}_t^{i,S} = \left(q_t, \theta_{1,t}, \theta_{2,t}, ..., \theta_{k-2,t}, \varphi_t\right)'$$

The realized measure $\tilde{\mathbf{y}}_t^{i,S}$ corresponding to $\tilde{\rho}_t^{i,S}$ can be obtained from $\tilde{\mathbf{y}}_t^i$ analogously by means of the same spherical transformation. Then, slightly rewritten equations (1.15)-(1.16) can be used to model the dynamics of $\tilde{\rho}_t^{i,S}$ and $\tilde{\mathbf{y}}_t^{i,S}$ instead of ρ_t^i and $\mathbf{y}_t^i$, so

$$\mathrm{Lgt}(\tilde{\mathbf{y}}_t^{i,S}) = \Psi_0^i + \Psi_1^i \,\mathrm{Lgt}(\tilde{\rho}_t^{i,S}) + \mathbf{v}_t^i \tag{A.2.13}$$

$$\mathrm{Lgt}(\tilde{\rho}_t^{i,S}) = \Lambda_0^i + \Lambda_1^i \,\mathrm{Lgt}(\tilde{\rho}_{t-1}^{i,S}) + \Lambda_2^i \,\mathrm{Lgt}(\tilde{\mathbf{y}}_{t-1}^{i,S}) \tag{A.2.14}$$

where Lgt$(\cdot)$ is a logistic transform. After $\tilde{\rho}_t^{i,S}$ is filtered, it can be then easily converted to $\tilde{\rho}_t^i$ using the inverse spherical transformation. Afterwards, $\tilde{\rho}_t^i$ can be rescaled back to ρ_t^i, which is an object of interest. By construction, such ρ_t^i will satisfy the criterion from Lemma 2 because q_t is confined on $(0,1)$ and positive definiteness of $R_t^{c,i}$ will be preserved for each t.

Nonetheless, there are several reasons of why this approach is not much convenient for a practical implementation. First of all, the problem arises with the proper identification of some angular coordinates since the inverse trigonometric functions are, generally speaking, multivalued. To illustrate this issue, lets consider the case with a 3-dimensional sphere ($k = 3$). Suppose that we transform vector $x = (x_1, x_2, x_3)'$ into the spherical coordinates. The radial distance and the polar angle are expressed as $q = \sqrt{x'x}$ and $\theta = \arccos(x_1/q)$. The azimuth angle φ has to be properly defined because it has a domain $[0, 2\pi)$ and inverse trigonometric functions are not single valued on a such range. Therefore, the transform should be suitable with respect to a particular quadrant on the (x_2, x_3)-plane. Commonly, the arctangent function with two arguments, $\varphi = \text{atan2}(x_2, x_3)$, is used to define this angle. An equivalent characterization, that is more geometrically appealing, can be suggested by virtue of the complex logarithm function. Namely

$$\varphi = \text{Im}\big(\log w\big)$$

where $w = x_3 + \text{i}x_2$ and Im$(\cdot)$ stands for an imaginary part of the complex number. Since the complex logarithm is a multivalued function, the proper branches on the Riemann manifold associated with the imaginary part of the logarithm have to be selected. In the context of our modeling strategy, we treat this angle as a dynamic coordinate, $\varphi = \varphi_t$. The representation $\varphi_t = \text{Im}\big(\log w_t\big)$ implies that we should model the dynamics of φ_t as a floating point on the Riemann surface, rather than directly map $[0, 2\pi)$ onto $\mathbb{R}$ as it was the case for other angular coordinates defined on $[0, \pi)$. Thus, we are able to avoid discontinuities and jumps in the dynamics of φ_t when it crosses the branches of the multivalued function. In practice, such modeling approach requires to intelligently select the suitable branch at every time point t, so it complicates the modeling and estimation routines considerably.

Beyond that, this modeling strategy implies three sequential transformations of the conditional correlation vector, whereas the basic approach suggested in the main paper requires just a monotone Fisher transform of the conditional correlations. Therefore, the variables Lgt$(\tilde{\mathbf{y}}_t^{i,S})$ and Lgt$(\tilde{\rho}_t^{i,S})$ in the dynamic equations (A.2.13)-(A.2.14) do not have a straightforward interpretation. As a consequence, we can not exlude the possibility that the vector autoregressions embedded in (A.2.13)-(A.2.14) are not adequate model for these transformed variables.

One potential advantage of such modeling approach is the way to reduce the dimensionality of the modeling dynamic variable. For instance, we may treat only the radial distance q_t as a time-varying component, whereas all angular coordinates $\theta_1, ..., \theta_{k-2}$ and φ can be treated as constant parameters. Therefore, all components of the conditional correlation vector ρ_t^i will still be endowed with some dynamics, but only one variable instead of k are modeled in (A.2.13)-(A.2.14). This would allow to significantly diminish the number of estimated parameters and, at the same time, the positive definiteness of $R_t^{c,i}$ will be achieved for all t.

A.3 Estimation Algorithm

Here we provide a detailed description of the estimation procedure that was adopted in our empirical analysis.

A.3.1 Core Model

The log-likelihood function for the observable variables in the core model is given by

$$\mathcal{L}^c = \sum_{t=1}^{T} \log p\big(\mathbf{r}_t^c | \mathcal{F}_{t-1}^c; \theta^c\big) + \sum_{t=1}^{T} \log p\big(\mathbf{x}_t^c, Y_t^c | \mathbf{r}_t^c, \mathcal{F}_{t-1}^c; \theta^c\big)$$

Denote the vector of all measurement errors from (1.5) and (1.7) as $U_t^c = \big((\mathbf{u}_t^c)', (\mathbf{v}_t^c)'\big)'$. Then, from Gaussian assumptions it follows that

$$\mathcal{L}_z^c = \sum_{t=1}^{T} \log p\big(\mathbf{r}_t^c | \mathcal{F}_{t-1}^c; \theta^c\big) = -\frac{1}{2} \sum_{t=1}^{T} \Big(k \log(2\pi) + \log |H_t^c R_t^c H_t^c| + (\mathbf{z}_t^c)' (R_t^c)^{-1} \mathbf{z}_t^c \Big) \tag{A.3.15}$$

$$\mathcal{L}_{u|z}^c = \sum_{t=1}^{T} \log p\big(\mathbf{x}_t^c, Y_t^c | \mathbf{r}_t^c, \mathcal{F}_{t-1}^c; \theta^c\big) = -\frac{1}{2} \sum_{t=1}^{T} \Big(\frac{1}{2} k(k+1) \log(2\pi) + \log |\Omega^c| + (U_t^c)' (\Omega^c)^{-1} U_t^c \Big) \tag{A.3.16}$$

Given a parameter vector θ^c and observable variables $\{\mathcal{X}_t^c\}_{t=1}^T$, the evaluation of these likelihood functions amounts to the following steps.

1. Initialize time-varying conditional volatilities $\mathbf{h}_1^c$ and off-diagonal elements $l\rho_1^c$ of the log-tranformed conditional correlation matrix $\text{logm}(R_t^c)$. We treat initial values $\mathbf{h}_1^c$ and $l\rho_1^c$ simply as additional parameters in the vector θ^c.

2. Using the model equations (1.4)-(1.8), we extract the series of time-varying parameters $\mathbf{h}_t^c(\theta^c)$ and $l\rho_t^c(\theta^c)$ for $t = 1, ..., T$.

3. Obtain matrices $H_t^c(\theta^c)$ and $R_t^c(\theta^c)$ using $\mathbf{h}_t^c(\theta^c)$ and $l\rho_t^c(\theta^c)$ for $t = 1, ..., T$. In order to recover matrix $R_t^c(\theta^c)$ from the off-diagonal elements of $LR_t^c(\theta^c) = \text{logm}(R_t^c(\theta^c))$, we rely on the implication of Conjecture 1. Namely, for a given symmetric matrix A there exists at most one vector of diagonal elements x such that the exponential transform of a matrix A with diagonal elements x, $e^{\tilde{A}(x)}$, is a correlation matrix. Thus, for each t, we find the vector d_t which solves $\text{diag}\big(\text{expm}\big(LR_t^c(d_t; \theta^c)\big)\big) = i_k$, where i_k is the $k \times 1$ vector of ones. Then the conditional correlation matrix is given by $R_t^c(\theta^c) = \text{expm}\big(LR_t^c(d_t; \theta^c)\big)$. In section A.3.4, we provide the iterative algorithm designed for this purpose.

4. Using (1.4), (1.5), and (1.7), obtain model residuals related to the standardized returns $\hat{\mathbf{z}}_t^c(\theta^c)$ and measurement errors $\hat{U}_t^c(\theta^c)$ for $t = 1, ..., T$.

5. Using filtered sequences $H_t^c(\theta^c)$, $R_t^c(\theta^c)$, and $\hat{\mathbf{z}}_t^c(\theta^c)$, the log-likelihood contribution from (A.3.15) can be evaluated, $\mathcal{L}_z^c(\theta^c)$.

6. We simplify the evaluation of the log-likelihood from (A.3.16) by concentrating it around the covariance matrix of the measurement errors, Ω^c. At first, we define its empirical analogue as $\hat{\Omega}^c(\theta^c) = \frac{1}{T} \sum_{t=1}^{T} \hat{U}_t^c(\theta^c) \hat{U}_t^c(\theta^c)'$. Then, the log-likelihood contribution given in (A.3.16) can be approximated as follows

$$\mathcal{L}_{u|z}^c(\theta^c) \approx -\frac{T}{2} \Big(\frac{1}{2} k(k+1) \log(2\pi) + \log |\hat{\Omega}^c(\theta^c)| + \frac{1}{2} k(k+1) \Big)$$

   As a result, we do not estimate parameters from Ω^c directly as components in θ^c. For large k it may significantly reduce the dimensionality of our maximization problem. We nevertheless obtain the indirect estimate of the covariance matrix of measurement errors, $\hat{\Omega}^c$.

7. The total log-likelihood can be found for a given θ^c as a sum of log-likelihood contributions from (A.3.15) and (A.3.16), $\mathcal{L}^c(\theta^c) = \mathcal{L}_z^c(\theta^c) + \mathcal{L}_{u|z}^c(\theta^c)$.

By using the provided algorithm, the log-likelihood function of the core model can then be maximized with respect to θ^c. Then, the solution $\hat{\theta}^c$, the vector that maximizes $\mathcal{L}^c(\theta^c)$, is used further to filter the model implied variables such as $H_t^c(\hat{\theta}^c)$, $R_t^c(\hat{\theta}^c)$, $\hat{\mathbf{z}}_t^c(\hat{\theta}^c)$, $\hat{U}_t^c(\hat{\theta}^c)$, and $\hat{\Omega}^c(\hat{\theta}^c)$, which will be further employed in the estimation of conditional factor models for individual asset returns.

A.3.2 Conditional Factor Model

The log-likelihood function for the observable variables in the conditional model for asset i is given by

$$\mathcal{L}^{i|c} = \sum_{t=1}^{T} \log p(r_t^i | \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}; \theta^i) + \sum_{t=1}^{T} \log p(x_t^i, \mathbf{y}_t^i | r_t^i, \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}; \theta^i)$$

At first, (1.12) provides us with the conditional distribution of a standardized asset return z_t^i. Denote the vector of all measurement errors from (1.13) and (1.15) of the realized measures related to asset i as $U_t^i = \big((u_t^i)', (\mathbf{v}_t^i)'\big)'$. Then, from (1.17) and using the properties of the multivariate Normal distribution we may obtain the conditional distribution of measurement errors U_t^i as

$$U_t^i \Big| U_t^c \sim N\Big(\Omega^{c,i}(\Omega^c)^{-1} U_t^c,\ \Omega^i - \Omega^{c,i}(\Omega^c)^{-1}(\Omega^{c,i})' \Big)$$

From the assumption of Normality it follows that

$$\mathcal{L}_z^{i|c} = \sum_{t=1}^{T} \log p(r_t^i | \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}; \theta_i) = -\frac{1}{2} \sum_{t=1}^{T} \Big(\log(2\pi) + \log h_t^{i|c} + \big(z_t^i - z_t^{i|c}\big)^2 h_t^i (h_t^{i|c})^{-1} \Big) \tag{A.3.17}$$

$$\mathcal{L}_{u|z}^{i|c} = \sum_{t=1}^{T} \log p(x_t^i, \mathbf{y}_t^i | r_t^i, \mathcal{X}_t^c, \mathcal{F}_{t-1}^{c,i}; \theta_i) = -\frac{1}{2} \sum_{t=1}^{T} \Big((k+1)\log(2\pi) + \log\big|\Omega^{i|c}\big| + \big(U_t^i - U_t^{i|c}\big)' (\Omega^{i|c})^{-1} \big(U_t^i - U_t^{i|c}\big) \Big) \tag{A.3.18}$$

where $z_t^{i|c} = (\rho_t^i)'(R_t^c)^{-1}\mathbf{z}_t^c$, $h_t^{i|c} = h_t^i\big(1 - (\rho_t^i)'(R_t^c)^{-1}\rho_t^i\big)$, $U_t^{i|c} = \Omega^{c,i}(\Omega^c)^{-1}U_t^c$, and $\Omega^{i|c} = \Omega^i - \Omega^{c,i}(\Omega^c)^{-1}(\Omega^{c,i})'$.

Given a parameter vector θ^i, all relevant filtered variables from the core model, and observable variables related to the individual asset $\{\mathcal{X}_t^i\}_{t=1}^T$, the evaluation of the provided log-likelihood functions proceeds as follows.

1. Initialize time-varying conditional volatility h_1^i and conditional correlations ρ_1^i. In contrast to the estimation of the core model, here we do not treat h_1^i and ρ_1^i as static parameters from the vector θ^i. Instead, we estimate univariate exponential smoothing dynamic models for the corresponding realized measures, where the initial values of the smoothed variables are estimated parameters. Then, we use these estimated values to initialize h_1^i and ρ_1^i. Though such approach is less appealing, it allows to avoid an increase in the number of estimated parameters in θ^i and, thus, reduces the estimation time. This gain is especially important when we have to estimate large number N of individual assets.

2. Using the model equations (1.9)-(1.16), we extract the series of time-varying parameters $h_t^i(\theta^i)$ and $\rho_t^i(\theta^i)$ for $t = 1, ..., T$.

3. For each day $t = 1, ..., T$ we check the positive definiteness of $R_t^{c,i}$ determined by the condition $(\rho_t^i)'(R_t^c)^{-1}\rho_t^i < 1$ in Lemma 2. If this condition is violated for at least one day, the given parameter vector θ^i generates the variables which are not admissible and such θ^i can not be a solution. Whereas the rigour approach would require a likelihood maximization subject to the dynamic constraint $(\rho_t^i)'(R_t^c)^{-1}\rho_t^i < 1$, we implement a simplified method. As long as the condition is violated for at least one trading day, we slightly perturb initial values h_1^i and ρ_1^i and look for another vector θ^i, so that the condition is getting satisfied.

4. Using (1.9), (1.13), and (1.15), we obtain model residuals related to the standardized returns $\hat{z}_t^i(\theta^i)$, $\hat{z}_t^{i|c}(\theta^i)$, and measurement errors $\hat{U}_t^i(\theta^i)$, $\hat{U}_t^{i|c}(\theta^i)$ for $t = 1, ..., T$.

5. Using filtered sequences $h_t^i(\theta^i)$, $\hat{z}_t^{i|c}(\theta^i)$, $h_t^{i|c}(\theta^i)$, and $\hat{z}_t^i(\theta^i)$, the log-likelihood contribution from (A.3.17) can be evaluated, $\mathcal{L}_z^{i|c}(\theta^i)$.

6. As before, we simplify the evaluation of the log-likelihood from (A.3.18) by concentrating it around the conditional covariance matrix of the measurement errors, $\Omega^{i|c}$. We define its empirical analogue as $\hat{\Omega}^{i|c}(\theta^i) = \frac{1}{T}\sum_{t=1}^{T}\big(\hat{U}_t^i(\theta^i) - \hat{U}_t^{i|c}(\theta^i)\big)\big(\hat{U}_t^i(\theta^i) - \hat{U}_t^{i|c}(\theta^i)\big)'$. Then, the log-likelihood contribution given in (A.3.18) can be approximated via

$$\mathcal{L}_{u|z}^{i|c}(\theta_i) \approx -\frac{T}{2}\bigg((k+1)\log(2\pi) + \log\big|\Omega^{i|c}(\theta^i)\big| + k + 1\bigg)$$

7. The total log-likelihood can be found for a given θ^i as a sum of log-likelihood contributions from (A.3.17) and (A.3.18), $\mathcal{L}^{i|c}(\theta^i) = \mathcal{L}_z^{i|c}(\theta^i) + \mathcal{L}_{u|z}^{i|c}(\theta^i)$.

Thus, the log-likelihood function of the conditional model for an individual asset i can be maximized with respect to θ^i. The solution $\hat{\theta}^i$ is used then to filter the model implied variables such as $h_t^i(\hat{\theta}^i)$, $\rho_t^i(\hat{\theta}^i)$, and $\hat{z}_t^i(\hat{\theta}^c)$, which can be further employed to obtain model-implied factor loadings, risk premia, and other relevant variables.

A.3.3 Large Systems

The hierarchical structure of the modeling framework suggests a following simple way of estimating large systems of asset returns. Suppose we have N assets and want to discover their cross-sectional properties.

1. The core model is estimated as it is explained in Section A.3.1. Variables $H_t^c(\hat{\theta}^c)$, $R_t^c(\hat{\theta}^c)$, $\hat{\mathbf{z}}_t^c(\hat{\theta}^c)$, $\hat{U}_t^c(\hat{\theta}^c)$, and $\hat{\Omega}^c(\hat{\theta}^c)$ related to the dynamics of the multivariate conditional distribution of the common risk factors are obtained.

2. For each asset $i = 1, ..., N$ we estimate a conditional factor model conditional on the same output variables filtered from the estimated core model as it is explained in Section A.3.2. Variables $\{h_t^i(\hat{\theta}^i)\}_{i=1}^N$, $\{\rho_t^i(\hat{\theta}^i)\}_{i=1}^N$, and $\{\hat{z}_t^i(\hat{\theta}^c)\}_{i=1}^N$ related to the factor dynamics of the conditional distributions of asset returns are obtained. These N conditional models can be estimated independently from each other, so the simultaneous parallel estimation is allowed and the estimation time can be considerably reduced.

Afterwards, according to equation (1.22) the common component of a conditional variance of the large dimensional return vector can be estimated. In addition, the cross-sections of the factor loadings (betas) can be computed to analyze asset pricing implications of the considered linear factor model.

A.3.4 Reconstruction of a Correlation Matrix

We now provide the algorithm to recover the correlation matrix from its logarithmic transformation when the diagonal elements of such transformation are unknown.

Suppose that $A \in \mathbb{R}^{n\times n}$ is a symmetric matrix with the off-diagonal elements obtained from the logarithmic transformation of some correlation matrix and arbitrarily specified diagonal elements. Our aim is to change the diagonal of A in such a way that the exponential transform of the new matrix $\tilde{A}$ produces a correlation matrix. We suggest the following simple iterative procedure that ensures a convergence of the initial matrix A to the desired matrix $\tilde{A}$:

$$A_{k+1} = A_k - \operatorname{logm}\big(\operatorname{Diag}(e^{A_k})\big), \qquad k = 0, 1, 2, ...$$

where $A_0 = A$ and $\operatorname{Diag}(A)$ represents a diagonal matrix with diagonal elements from A. Since the updating term is a diagonal matrix, iterations affect only diagonal elements of the matrix A and leaves the off-diagonal elements unchanged. With a reasonably chosen tolerance level, the algorithm converges to $\tilde{A}$ within just a few iterations.

We relate the detailed analysis of the suggested algorithm and its convergence properties to the separate companion note.

A.4 Descriptive Statistics

A.4.1 Sample description

Sector	#	Group	#
Energy	49	Energy	49
Materials	61	Materials	61
Industrials	121	Capital Goods	80
		Commercial Services and Supplies	20
		Transportation	21
Consumer Discretionary	153	Automobiles and Components	12
		Consumer Durables and Apparel	36
		Consumer Services	29
		Media	19
		Retailing	57
Consumer Staples	40	Food and Staples Retailing	11
		Food, Beverage and Tobacco	19
		Household and Personal Products	10
Health Care	98	Health Care Equipment and Services	51
		Pharmaceuticals, Biotechnology and Life Sciences	47
Financials	93	Banks	39
		Diversified Financials	23
		Insurance	25
		Real Estate	6
Information Technology	152	Software and Services	50
		Technology Hardware and Equipment	48
		Semiconductors and Semiconductor Equipment	54
Telecommunication Services	9	Telecommunication Services	9
Utilities	48	Utilities	48

Table A.4.1: Table contains the breakdown of sample assets used in the empirical analysis according to the Global Industry Classification Standard (GICS).

Fund's type	Ticker	Description	Fund's type	Ticker	Description
Style funds	IWD	large value	Sectoral funds	XLB	Materials
BlackRock, iShares	IWB	large neutral	*SSGA, SPDR*	XLE	Energy
	IWF	large growth		XLF	Financials
	IWS	mid value		XLI	Industrials
	IWR	mid neutral		XLK	Technology
	IWP	mid growth		XLP	Consumer staples
	IWN	small value		XLU	Utilities
	IWM	small neutral		XLV	Health care
	IWO	small growth		XLY	Consumer discretionary

Table A.4.2: List of ETFs used in the empirical analysis.

A.4.2 Summary Statistics

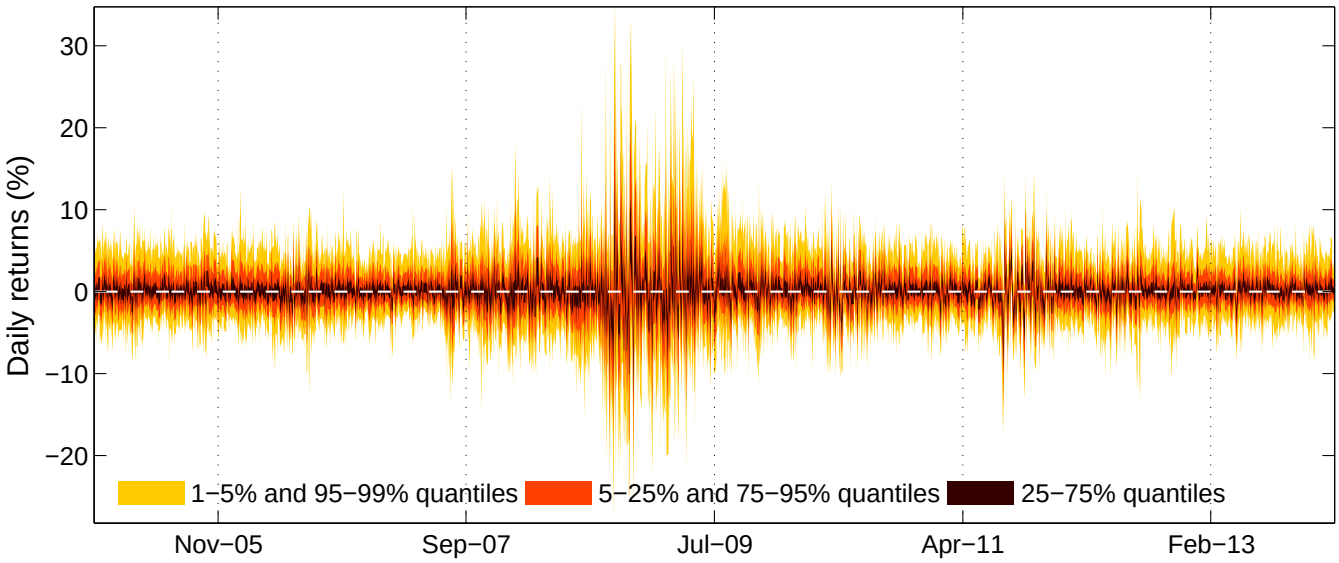

Figure A.4.1: Quantile plot of daily returns from the cross-section of 824 individual assets used in the empirical analysis.

Quantile	0.01	0.05	0.10	0.25	0.40	0.50	0.60	0.75	0.90	0.95	0.99
Mean	-0.016	0.017	0.026	0.042	0.052	0.061	0.068	0.085	0.110	0.130	0.183
Std.	1.141	1.378	1.565	2.000	2.325	2.510	2.726	3.155	3.832	4.443	5.884
Skewness	-1.307	-0.598	-0.322	-0.041	0.123	0.230	0.344	0.546	0.998	1.707	4.525
Kurtosis	5.850	7.109	7.698	9.136	10.723	11.912	13.425	16.726	26.656	36.574	88.837
Min	-61.105	-39.567	-32.948	-24.239	-19.228	-17.707	-15.952	-13.264	-10.397	-8.875	-7.870
Median	-0.139	-0.059	-0.002	0.000	0.021	0.035	0.043	0.061	0.090	0.104	0.148
Max	8.603	10.202	11.722	14.966	17.726	19.397	21.761	26.956	38.282	50.556	103.458
ACF(1)	-0.176	-0.123	-0.102	-0.065	-0.045	-0.033	-0.022	-0.002	0.026	0.042	0.077
ACF(5)	-0.103	-0.075	-0.064	-0.046	-0.031	-0.024	-0.017	-0.005	0.012	0.025	0.062
ACF(20)	-0.042	-0.023	-0.014	0.002	0.014	0.021	0.027	0.039	0.056	0.069	0.106

Table A.4.3: Summary statistics of daily returns.

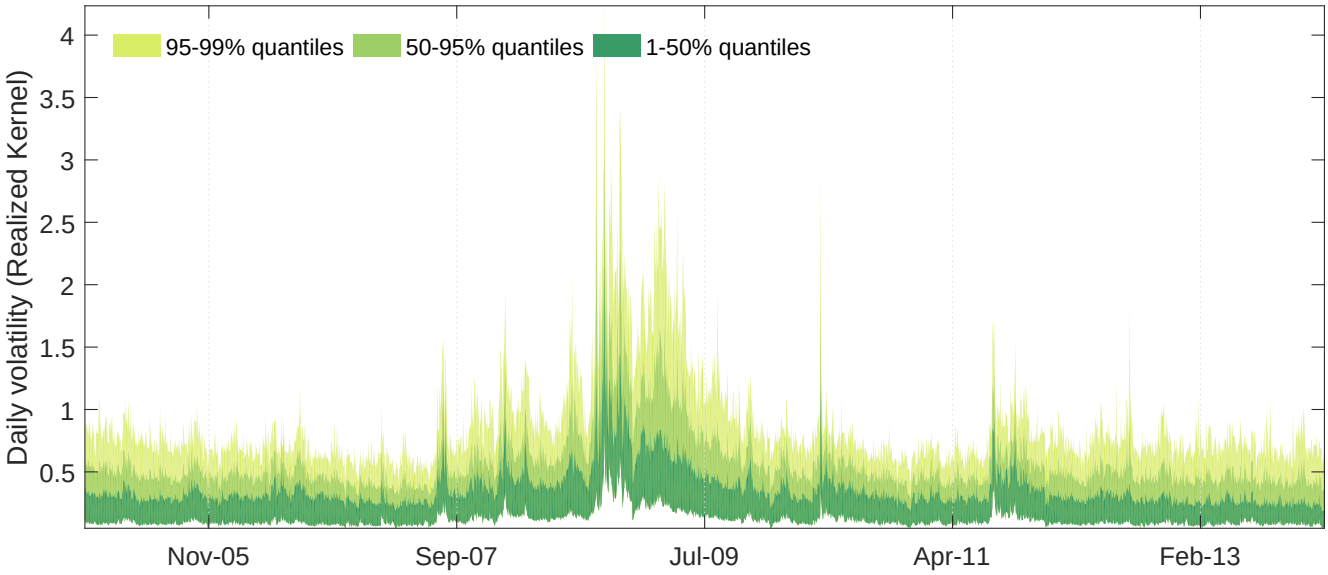

Figure A.4.2: Quantile plot of daily annualized volatilities from the cross-section of 824 individual assets used in the empirical analysis. Variances are computed using the Realized kernel estimator.

Quantile	0.01	0.05	0.10	0.25	0.40	0.50	0.60	0.75	0.90	0.95	0.99
Mean	0.155	0.180	0.197	0.239	0.277	0.296	0.320	0.369	0.438	0.499	0.612
Std.	0.083	0.103	0.113	0.133	0.154	0.167	0.180	0.211	0.284	0.344	0.483
Skewness	1.751	1.928	2.110	2.453	2.682	2.888	3.081	3.491	4.258	5.075	7.826
Kurtosis	7.707	9.247	10.463	13.125	15.381	17.260	19.695	24.277	38.099	59.925	149.461
Min	0.016	0.030	0.036	0.046	0.054	0.058	0.063	0.072	0.095	0.108	0.133
Median	0.133	0.152	0.167	0.200	0.231	0.248	0.271	0.314	0.374	0.419	0.506
Max	1.025	1.220	1.354	1.602	1.800	1.992	2.200	2.631	3.724	4.522	7.238
ACF(1)	0.458	0.537	0.581	0.663	0.718	0.747	0.775	0.807	0.840	0.859	0.898
ACF(5)	0.262	0.378	0.431	0.532	0.599	0.638	0.664	0.706	0.748	0.772	0.816
ACF(20)	0.170	0.274	0.332	0.427	0.493	0.522	0.545	0.585	0.631	0.655	0.715

Table A.4.4: Summary statistics of daily volatilities.

	IWD	IWB	IWF	IWS	IWR	IWP	IWN	IWM	IWO	XLB	XLE	XLF	XLI	XLK	XLP	XLU	XLV	XLY
	Daily returns (%)																	
Mean	0.025	0.028	0.033	0.034	0.037	0.039	0.028	0.035	0.041	0.034	0.059	0.003	0.034	0.032	0.033	0.023	0.035	0.035
Std.	1.407	1.303	1.250	1.474	1.436	1.451	1.725	1.671	1.644	1.696	1.981	2.316	1.442	1.354	0.878	1.178	1.042	1.460
Skewness	0.029	-0.119	0.018	-0.385	-0.178	-0.309	-0.228	-0.223	-0.142	-0.085	-0.202	0.343	-0.059	0.306	-0.218	0.657	0.013	-0.196
Kurtosis	13.233	12.763	13.706	9.576	10.501	9.877	8.530	7.828	7.449	9.742	12.740	15.067	9.264	13.769	9.111	16.818	18.114	9.610
Min	-9.171	-9.375	-8.630	-9.828	-9.413	-10.526	-12.641	-11.238	-10.225	-12.412	-14.444	-16.667	-9.405	-8.654	-6.024	-7.440	-9.783	-11.625
Median	0.075	0.080	0.082	0.103	0.098	0.106	0.093	0.094	0.091	0.104	0.119	0.030	0.086	0.090	0.076	0.076	0.040	0.061
Max	12.831	11.372	12.309	8.810	10.995	10.182	9.485	8.637	10.944	14.058	16.475	16.462	10.705	13.898	6.886	12.073	12.055	9.775
ACF(1)	-0.119	-0.102	-0.083	-0.076	-0.056	-0.043	-0.118	-0.090	-0.056	-0.050	-0.116	-0.128	-0.056	-0.101	-0.120	-0.125	-0.091	-0.040
ACF(5)	-0.094	-0.092	-0.077	-0.082	-0.081	-0.061	-0.066	-0.065	-0.075	-0.041	-0.087	-0.080	-0.044	-0.046	-0.101	-0.128	-0.087	-0.053
ACF(20)	0.028	0.002	0.001	0.027	0.013	0.037	0.022	0.028	0.023	0.006	-0.011	0.056	0.022	0.028	-0.039	0.006	-0.004	-0.009
	Daily volatility (Realized Kernel)																	
Mean	0.134	0.128	0.130	0.141	0.142	0.146	0.183	0.186	0.184	0.191	0.214	0.201	0.154	0.148	0.107	0.148	0.120	0.154
Std.	0.100	0.093	0.088	0.102	0.101	0.100	0.110	0.106	0.101	0.120	0.125	0.172	0.104	0.092	0.069	0.095	0.072	0.105
Skewness	2.968	3.061	3.209	2.596	2.774	2.805	2.423	2.605	2.589	2.806	3.105	2.583	2.865	3.159	3.898	3.563	3.249	2.634
Kurtosis	15.318	16.353	17.607	12.224	13.628	14.010	11.435	12.897	12.994	15.224	17.743	11.258	14.668	17.481	26.358	21.469	18.890	12.806
Min	0.025	0.020	0.031	0.022	0.022	0.032	0.044	0.053	0.050	0.044	0.056	0.034	0.034	0.030	0.027	0.036	0.027	0.031
Median	0.103	0.101	0.105	0.109	0.112	0.117	0.151	0.155	0.156	0.155	0.186	0.138	0.124	0.122	0.089	0.122	0.100	0.118
Max	0.929	0.894	0.819	0.843	0.887	0.883	0.987	0.965	0.921	1.214	1.261	1.445	0.979	0.899	0.880	1.001	0.777	0.937
ACF(1)	0.891	0.878	0.877	0.887	0.876	0.871	0.868	0.877	0.845	0.891	0.891	0.915	0.890	0.850	0.837	0.875	0.856	0.899
ACF(5)	0.789	0.772	0.771	0.801	0.783	0.773	0.766	0.767	0.726	0.785	0.796	0.822	0.791	0.738	0.737	0.776	0.737	0.804
ACF(20)	0.649	0.628	0.621	0.675	0.654	0.636	0.635	0.625	0.581	0.634	0.620	0.717	0.659	0.598	0.559	0.597	0.598	0.691

Table A.4.5: Summary statistics of daily returns and volatilities for ETFs used in the empirical analysis.

A.5 Fama-French factors from ETF data

Here we describe the procedure for replication of the three Fama-French risk factors using the data from style and broad market exchange traded funds (ETFs). We begin with a general empirical formulation of these risk factors, outline the replication strategy and introduce a descriptive list of all ETFs used in the analysis. Then, we describe the scheme that we use to construct ETF-based risk factors on a daily and an intra-daily frequencies.

A.5.1 Three Fama-French Factors

The three risk factors used in the Fama-French model are essentially unobservable. According to the original approach (Fama and French (1993)), the factors can be replicated empirically as follows:

$$\begin{aligned} r_{mkt} &= r_m - r_{rf} \\ r_{smb} &= \frac{1}{3}\big(r_{sv} + r_{sn} + r_{sg}\big) - \frac{1}{3}\big(r_{bv} + r_{bn} + r_{bg}\big) \\ r_{hml} &= \frac{1}{2}\big(r_{sv} + r_{bv}\big) - \frac{1}{2}\big(r_{sg} + r_{bg}\big) \end{aligned} \tag{A.5.19}$$

where r_{mkt} denotes the market risk factor (MKT), r_{smb} is the size factor which reflects a market capitalization of stocks (small-minus-big, SMB), and r_{hml} is value factor that captures a book-to-market ratio (high-minus-low, HML). Factors are constructed using the following returns on specific portfolios:

return	associated portfolio
r_{rf}	riskless portfolio (usually Treasury Bill rate is used)
r_m	market representative
r_{sv}	stocks with small market capitalization and high book-to-market ratio
r_{sn}	stocks with small market capitalization and medium book-to-market ratio
r_{sg}	stocks with small market capitalization and low book-to-market ratio
r_{bv}	stocks with big market capitalization and high book-to-market ratio
r_{bn}	stocks with big market capitalization and medium book-to-market ratio
r_{bg}	stocks with big market capitalization and low book-to-market ratio

As we can see, the market factor represents a return on a market portfolio in excess of the risk-free rate. As long as the riskless portfolio is, generally speaking, an infeasible object, the corresponding risk-free return is usually approximated by the one-month Treasury Bill rate. The size factor is a return on zero investment portfolio formed by taking a long position on small cap stocks and short on large cap stocks. Finally, the value factor is a return on zero investment portfolio formed by taking a long position on stocks with a high book-to-market ratio and a short position on stocks with a low book-to-market ratio.

Usually, in empirical research people refer to specifically constructed research portfolios that are based on the data from many individual stocks and are representative for a particular equity market. As an example, Kenneth French Data Library provides the data on the Fama-French factors for the U.S. equity market based on the regularly updated research portfolios started from 1927.

A.5.2 Replication Strategy

Our ultimate purpose in this section is to replicate returns on the Fama-French factors on a daily and intra-daily frequencies. The high frequency intra-daily factors allow then to obtain the corresponding realized measures of volatilities and correlations for the daily factors. Afterwards, the daily "close-to-close" factors and realized measures are supposed to be input variables for the multi-factor Realized GARCH model developed in this paper.

As for the daily factors, they can be directly obtained from the Kenneth French Data Library. They are constructed with market representative portfolios using formula (A.5.19) and are considered as a benchmark

in the empirical asset pricing literature being used intensively in the related research. The daily frequency is the highest available frequency of the factor series at the Kenneth French website, so the replication of the Fama-French factors intra-daily is the main challenge.

One straightforward possibility would be to use high frequency transaction observations of the assets from the same research portfolios which were used to construct the daily factors. In that case, intra-daily portfolio returns can be obtained and the corresponding intra-daily factors can be then directly computed using formula (A.5.19). Although the intra-daily factors obtained in such a way would be fairly consistent with the daily factors from the Kenneth French Library, this approach requires to aggregate large amounts of high frequency data from the individual stocks as well as to maintain the exact compositions of the research portfolios for each day. We, therefore, resort to another possibility of getting the high frequency factors that is much easier to implement and replicate.

Data from market and style ETFs suggests an alternative approach to construct the risk factors. Since an ETF represents a portfolio of stocks traded on a market as a single security, the latent Fama-French factors can be replicated by means of proper combinations of ETF returns. Using relations from (A.5.19), where returns on corresponding market and style ETFs are used instead of the returns on research portfolios, we may obtain the ETF-based Fama-French factors. Because intra-daily ETF transactions are directly observable, we may use high frequency ETF data to replicate the latent risk factors at an intra-daily level.

Since the compositions of assets in the ETFs and the Fama-French portfolios are not the same, the realized measures obtained using ETF data are not fully consistent with the daily factors which are based on the research portfolios. To reduce the potential mismatch, we look across several alternative ETF families and gauge their abilities to replicate the benchmark Fama-French factors. Since the benchmark factors are available just at a daily frequency, but not intra-daily, we perform our comparison at a daily level. We conjecture, that if some given ETF family replicates the daily risk factors, i.e. better correlates with the benchmark (daily factors from the Kenneth French Data Library), then on an intra-daily frequency this ETF family would also better replicate the latent benchmark (intra-daily factors based on the Fama-French portfolios).

A.5.3 Exchange Traded Funds Used for Factor Replication

In order to replicate the Fama-French factors we consider data on style and broad market equity ETFs from the three main U.S. investment management corporations: BlackRock, State Street Global Advisors (SSgA), and Vanguard Group. Table (A.5.6) provides a brief description of the corresponding funds. In addition, there are several issues that we point out here in more details.

We perform our comparison across 4 fund families. We selected these families based on several general requirements. First of all, we consider the funds with sufficiently long trading history. This requirement is natural since an accurate estimation of the dynamic model formulated in this paper needs sufficiently long data span. Secondly, it is highly desirable that a considered ETF is traded intensively. In this case, the number of intra-daily transactions is supposed to be higher, the price responds to the arrival of relevant information more timely and, thus, more precise realized measures of variances and correlations can be obtained. Finally, we want to use funds that can be treated as representative portfolios for the U.S. equity market. Thus, we focus on those ETFs which track representative stock market indexes, but not those which are related to a particular sector or industry.[1]

As it can be seen from Table (A.5.6), iShares (Russell 3000) funds are the most appealing from the perspective of the aforementioned requirements. They all are incepted in 2000, are comparatively liquid, and are based on the constituents of the Russell 3000 index which represents about 98% of the U.S. equity market. In contrast, SPDR family which includes well known large blend SPY, the most liquid ETF with the longest trading history,

[1]We notice that these selection criteria do not support the use of momentum ETFs to proxy a momentum factor which could be added then as an additional risk factor (Carhart (1997)) in our empirical pricing model. In particular, although DWAQ (PowerShares, Invesco) is traded since 2003, it tracks momentum in NASDAQ companies and is not representative for the broad equity market. Other popular momentum funds include PDP (PowerShares, Invesco, since 2007), MTUM (iShares, BlackRock, since 2013), MMTM (SSgA, since 2012), MOM (QuantShares, since 2011), etc. Such funds have insufficiently long trading history, low trading volumes, and, as a result, are inappropriate for our analysis.

is traded less intensively than alternatives and is less representative for the broad U.S. stock market (compared to the Russell 3000 and CRSP Total) since the funds are based on the components of the S&P 1500 index.

All the considered funds are passively managed. It means that the underlying security composition is supposed to replicate the performance of the corresponding index or its specific segment characterized by a particular investment style (value, neutral, or growth stocks; large or small-cap stocks). Nonetheless, it does not mean that the composition of securities in ETFs is fully representative for the whole U.S. equity market (unlike the research portfolios from the Kenneth French Library). Moreover, it is pretty unlikely that the ETFs perfectly replicate the corresponding indices and style segments at least due to the presence of tracking errors (see Hassine and Roncalli (2013)). In addition, a composition of an ETF may be slightly tilted and differ from the index representative portfolio because of the fund-specific weighting schemes. Thus, as we have already pointed out, we can not expect that the ETF-based Fama-French factors perfectly match the Fama-French risk factors obtained with the highly representative research portfolios.

In what follows, we will compare how well the factors constructed with different fund families match the benchmark factors at a daily frequency. Based on this comparison, we select the funds for the intra-daily replication of the Fama-French factors.

A.5.4 Matching ETF-based Factors with the Benchmark

As a benchmark, we use daily factors from the Kenneth French Data Library. Daily ETF returns are obtained from the Yahoo Finance website and the one-month Treasury bill rate (from Ibbotson Associates) is treated as a risk-free rate r_{rf}.

A.5.4.1 Market Factor

Several ETFs from those listed in Table (A.5.6) might be used to approximate the market factor. Namely, we consider daily returns on 4 large-cap neutral style funds and 2 broad market funds. These funds are supposed to closely follow the stock market indexes associated with the corresponding ETF families. Whereas the large neutral funds track only the companies with a large capitalization, broad market funds do not relate to a particular style segment and can be treated as more representative market portfolios.

We define r_m as a daily return on one of these ETFs and then obtain the corresponding ETF-based market factor using the definition from (A.5.19). Then, for the period 2005-2013 we compute sample correlations between the ETF-based market factors and the benchmark market factor from the Kenneth French Library.

ETF for market factor	IWB	IVV	VV	SPY	IWV	VTI
Correlation with the benchmark	0.993	0.991	0.987	0.985	0.994	0.992

As we may see, all 6 funds match the benchmark almost perfectly and either can be used as a market factor. Nonetheless, we select VTI as a replicating ETF for the market risk factor since this fund is representative for the broad market (but not just for large-cap stocks) and is traded more intensively than the other broad market fund IWV.

A.5.4.2 Size Factor

We construct the size factor according to the definition in (A.5.19) by using returns on the style ETFs. The sample correlations between SMB factors replicated with ETFs from 4 different fund families and the benchmark SMB are

ETF family for size factor	iShares, Russell	iShares, S&P	Vanguard	SSgA
Correlation with the benchmark	0.887	0.899	0.865	0.695

We can see that both iShares families as well as Vanguard funds replicate the benchmark Fama-French size factor quite well. In contrast, SSgA funds perform relatively worse. This may be explained by the fact that SSgA style ETFs are on average much less liquid than the alternatives. Since the iShares Russell funds

are traded more intensively than Vanguard and iShares S&P funds, we choose this family to replicate the size factor.

A.5.4.3 Value Factor

We obtain the daily HML factors with daily returns on the style ETFs using (A.5.19) and the comparison with the benchmark value factor gives

ETF family for value factor	iShares, Russell	iShares, S&P	Vanguard	SSgA
Correlation with the benchmark	0.783	0.723	0.723	0.497

Similar to the previous case, iShares and Vanguard funds replicate the benchmark value factor decently (but significantly worse than in case of the size factor). HML factor based on the SSgA funds exhibits very low correlation with the benchmark factor. We select the iShares Russell ETF family for the replication since the value factor constructed with these funds has the highest correlation with the benchmark.

Figure A.5.3 contains cumulative series of the benchmark research factors from the Kenneth French Data Library and the ETF-based factors constructed with the funds selected above. Particularly, these plots illustrate that although the correlations between the ETF-based factors and the benchmark factors vary across different sample years, a significant co-movement is always observable.[2]

Finally, we acknowledge that despite our ETF-based factors correlate pretty well with the benchmark at a daily frequency, the realized measures constructed with the ETF-based factors at an intra-daily frequency still may be inconsistent measures of variances and correlations for the daily benchmark Fama-French factors. Importantly, we suppose that such realized measures capture the dynamics of these latent variances and correlations good enough and can be treated as valid measures (instruments) of the latter. At the same time, a potential systematic bias in the realized measures is supposed to be controlled by the measurement equation parameters in the model.

[2]We note that it is also possible to exploit the infrormation from all considered funds and obtain combined ETF-based factors. For example, such a combined ETF-based factor can be built as a linear combination of the ETF-based factors obtained from different fund families. To determine the weights in such a linear combination, a simple linear regression can be used

$$r_{x,t}^{bench} = \bar{r}_{x,t}\beta_x + \varepsilon_t$$

where $r_{x,t}^{bench}$ is the (daily) benchmark factor x (market, size, or value) and $\bar{r}_{x,t}$ is a vector of the ETF-based factors x obtained with different fund families. A combined ETF-based factor is, therefore, defined by $\bar{r}_{x,t}\hat{\beta}_x$, where $\hat{\beta}_x$ is a vector of estimated coefficients. In that case, an increased representativeness of an underlying stock composition in the combined ETF-based factors may lead to an improved matching with the Fama-French factors constructed by means of the representative researcher portfolios. However, to avoid overfitting issues and potential adverse effects related to the intra-daily aggregation of many price series, we do not use such combined ETF-based factors in our empirical analysis.

Fund family and Index	Ticker	Style	Inception date	Average daily volume	Stock splits (since 2004)	
					Proportion	Date
BlackRock, iShares	IWV	broad market	22.05.2000	$4.44 \cdot 10^5$		
Russell 3000	IWD	large value	22.05.2000	$1.73 \cdot 10^6$		
	IWB	large neutral	15.05.2000	$1.29 \cdot 10^6$		
	IWF	large growth	22.05.2000	$2.33 \cdot 10^6$		
	IWN	small value	24.07.2000	$1.66 \cdot 10^6$	1:3	09.06.2005
	IWM	small neutral	22.05.2000	$5.34 \cdot 10^7$	1:2	09.06.2005
	IWO	small growth	24.07.2000	$1.85 \cdot 10^6$		
BlackRock, iShares	IVE	large value	22.05.2000	$6.04 \cdot 10^5$		
S&P 1500	IVV	large neutral	15.05.2000	$3.03 \cdot 10^6$		
	IVW	large growth	22.05.2000	$7.54 \cdot 10^5$		
	IJS	small value	24.07.2000	$2.04 \cdot 10^5$	1:2	09.06.2005
	IJR	small neutral	22.05.2000	$1.46 \cdot 10^6$	1:3	09.06.2005
	IJT	small growth	24.07.2000	$2.23 \cdot 10^5$	1:2	24.07.2008
Vanguard	VTI	broad market	24.05.2001	$1.80 \cdot 10^6$	1:2	18.06.2008
CRSP US Total	VTV	large value	26.01.2004	$3.60 \cdot 10^5$		
	VV	large neutral	27.01.2004	$2.20 \cdot 10^5$		
	VUG	large growth	26.01.2004	$3.79 \cdot 10^5$		
	VBR	small value	26.01.2004	$1.24 \cdot 10^5$		
	VB	small neutral	26.01.2004	$2.68 \cdot 10^5$		
	VBK	small growth	26.01.2004	$1.27 \cdot 10^5$		
SSgA, SPDR	SPYV	large value	25.09.2000	$1.31 \cdot 10^4$	1:2	22.09.2005
S&P 1500	SPY	large neutral	22.01.1993	$1.57 \cdot 10^8$		
	SPYG	large growth	25.09.2000	$2.11 \cdot 10^4$		
	SLYV	small value	25.09.2000	$1.33 \cdot 10^4$	1:3	22.09.2005
	SLY	small neutral	08.11.2005	$1.08 \cdot 10^4$		
	SLYG	small growth	25.09.2000	$1.21 \cdot 10^4$		

Table A.5.6: Descriptive information on the passively managed equity exchange traded funds which are considered and used in the replication of the Fama-French risk factors. The first column contains the name of investment management company, the name of the fund family, and the stock market index related to the particular fund family. The second column includes the ticker symbols, while the investment styles are in the third column. The fourth column holds the inception dates of the considered ETFs. The fifth column contains the average daily traded volumes which are computed for the trading days within the period 2004 - 2013 inclusively. The last two columns provide an information about the dates and proportions of the ETF stock splits (if happened) since 2004.

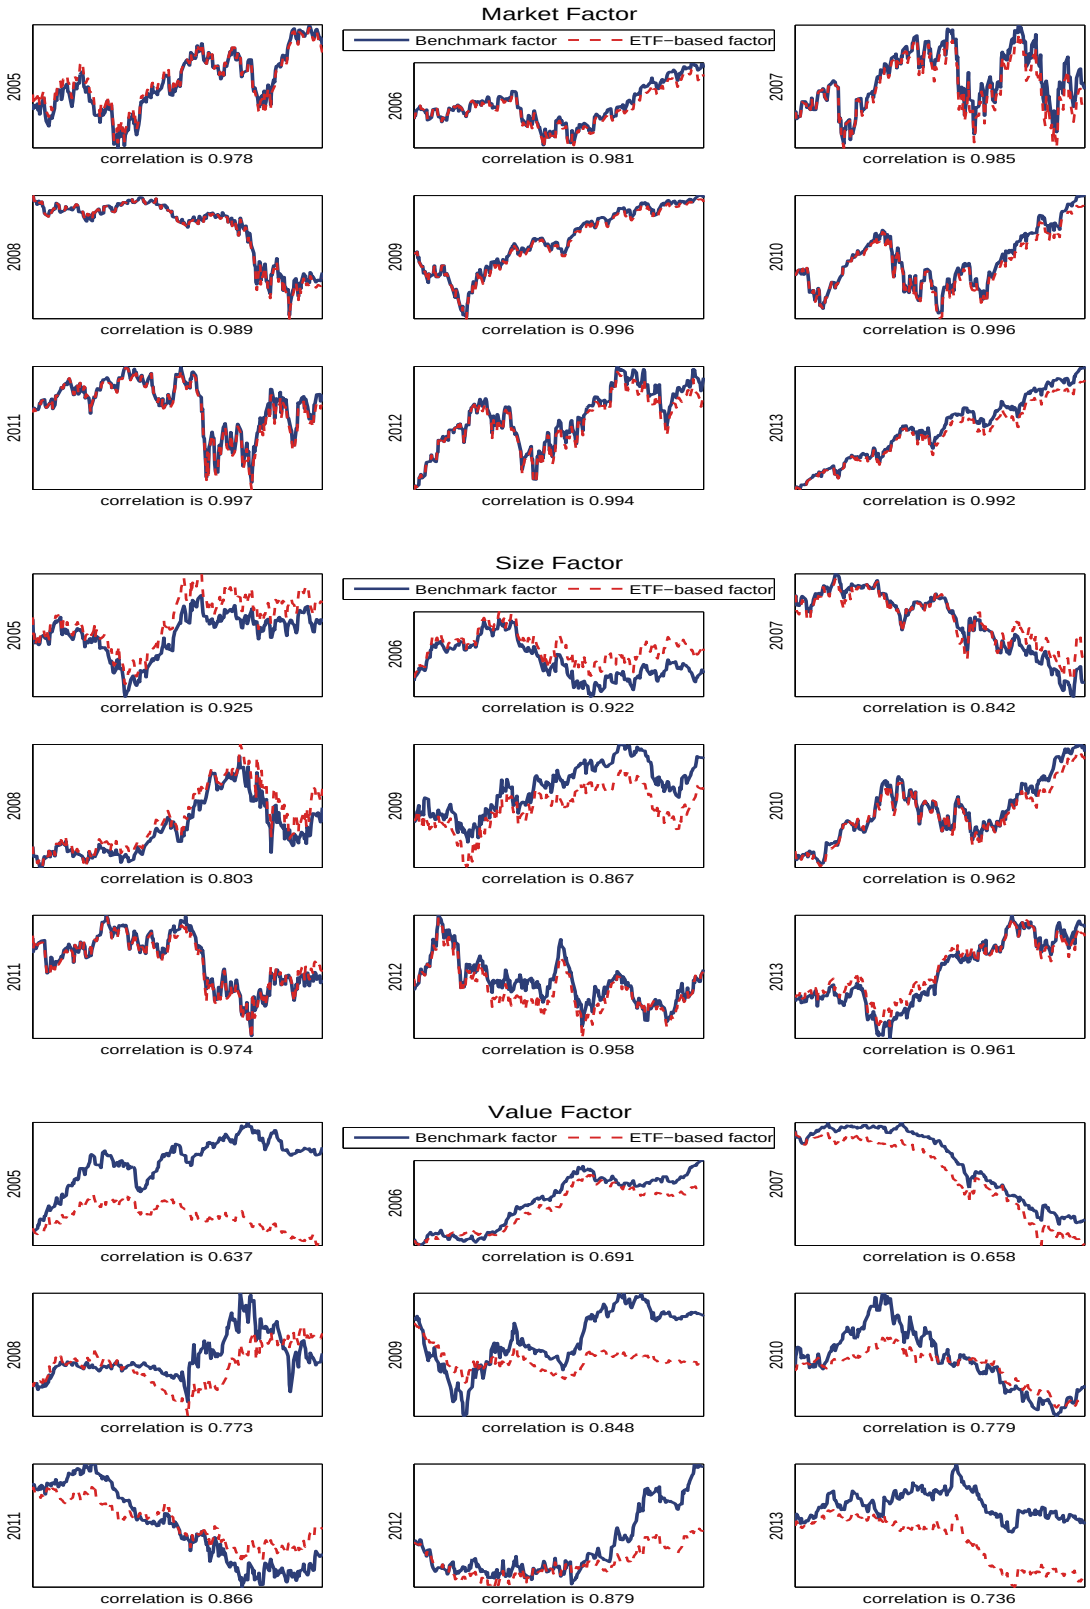

Figure A.5.3: Cumulative series of the benchmark daily risk factors from the Kenneth French Data Library (solid lines) and the ETF-based daily risk factors (dashed lines). Plots are provided for each calendar year from the sample period between 2005 and 2013 inclusively and the corresponding correlations between factors based on a one year period are indicated beneath the plots. The ETF-based risk factors are constructed as described in Section A.5.4.

A.6 Model Evaluation

We provide a statistical evaluation of the suggested modeling framework. At first, we formulate the set of alternative models that will be used for comparative purposes. Then, we discuss the results of in-sample and out-of-sample statisical analysis.

A.6.1 Alternative models

The model described in the paper allows for a dynamic correlation structure, exploits high frequency (intra-daily) data, and also incorporates the signals from low frequency (daily) returns. To evaluate the potential benefits suggested by these features, we specify the pool of alternative models. Each model from this pool isolates one of the key aspects from our baseline framework.

1. ***A Realized CCC Model***. In this model, only variances are treated as time-varying, whereas correlations between asset returns are assumed to be constant. Therefore, we rule out the correlation dynamics from the framework. Such model can be considered as a realized version of the Multivariate CCC (Constant Conditional Correlations) GARCH of Bollerslev (1990). Note that this model is nested in our baseline framework.

2. ***A Realized Reduced Form Model for Variances and Correlations***. Here we discard the information from the low frequency daily returns, so the model can be treated as a simple filter for realized measures. Without this information, the filtered covariance process is no more an unbiased covariance measure of asset returns, but rather a smoothed version of realized measures. In addition, such aspects of the baseline modeling framework as a "leverage effect" are left out.

3. ***A DCC Model with Log-matrix Transform***. In this case, we exclude realized measures based on high frequency data. Instead, we provide a close counterpart of multivariate GARCH frameworks of Engle (2002) and Tse and Tsui (2002), where the correlation dynamics is modeled using the log-matrix transformation, as it is developed in our baseline framework. Although such model is not nested in our main model, it can be treated as a closely related low frequency analogue.

A.6.2 Statistical Evaluation

We evaluate both the core and the factor models from our baseline framework. As for the core model, we compare partial and predictive log-likelihoods for the system of the three Fama-French risk factors. A comparison of results across the pool of alternative models will allow us to assess the relative benefits in a statistical fit that are suggested by different aspects of our baseline framework. As for the factor models, we implement a beta hedging tracking exercise (in- and out-of-sample) using style and sectoral ETFs as individual assets.

Table A.6.7 contains log predictive likelihoods of the three-variate Fama-French return for 4 models. It is quite naturally that the DCC model performs in-sample better than the alternatives since the likelihood estimation objective coincides with the evaluation criterion. This is not the case for other three models, where the likelihood objectives have additional components related to the observable realized measures. In contrast, in the out-of-sample comparison the baseline realized model performs better than the DCC model. Note that a poor performance of the reduced form realized model reflects the absence of the return distribution in the estimation objective.

Table A.6.8 provides the variances of hedging tracking errors. As we may see, the realized models demonstrate significantly better hedging performance than the DCC model. Particularly, it may indicate the superior information content of the variance and correlation signals obtained from realized measures. It is interesting to note that for several funds the realized model with constant correlations outperforms the baseline model both in- and out-of-sample.

DCC with LM-transform	Reduced Form Realized Model	Realized CCC	Baseline Model
	In-sample Mean Predictive Log-likelihood		
-2.619	-3.339	-2.762	-2.658
	Out-of-sample Mean Predictive Log-likelihood		
-1.850	-2.374	-1.840	-1.824
MCS size	1%	5%	10%

Table A.6.7: Density forecasting for the Fama-French factors.

	IWD	IWB	IWF	IWS	IWR	IWP	IWN	IWM	IWO
				In-sample Fund Tracking Variance					
DCC with LM-transform	0.069	0.047	0.062	0.085	0.080	0.090	0.097	0.093	0.090
Reduced Form Realized Model	0.049	0.024	0.056	0.073	0.067	0.092	0.089	0.074	0.099
Realized CCC	0.045	0.022	0.049	0.067	0.060	0.080	0.082	0.072	0.073
Baseline Model	0.044	0.039	0.083	0.067	0.061	0.081	0.083	0.071	0.092
				Out-of-sample Fund Tracking Variance					
DCC with LM-transform	0.031	0.021	0.022	0.055	0.044	0.053	0.058	0.051	0.053
Reduced Form Realized Model	0.022	0.012	0.023	0.042	0.036	0.050	0.029	0.022	0.045
Realized CCC	0.017	0.010	0.016	0.037	0.029	0.036	0.031	0.024	0.034
Baseline Model	0.017	0.018	0.049	0.042	0.032	0.039	0.028	0.021	0.041
MCS size					1%		5%		10%

	XLB	XLE	XLF	XLI	XLK	XLP	XLU	XLV	XLY
				In-sample Fund Tracking Variance					
DCC with LM-transform	0.193	0.212	0.148	0.148	0.133	0.136	0.174	0.146	0.145
Reduced Form Realized Model	0.143	0.156	0.071	0.088	0.110	0.125	0.175	0.122	0.093
Realized CCC	0.187	0.212	0.152	0.137	0.133	0.131	0.182	0.141	0.137
Baseline Model	0.189	0.208	0.145	0.134	0.128	0.131	0.178	0.143	0.138
				Out-of-sample Fund Tracking Variance					
DCC with LM-transform	0.161	0.164	0.079	0.098	0.117	0.125	0.173	0.132	0.111
Reduced Form Realized Model	0.143	0.156	0.071	0.088	0.110	0.125	0.175	0.122	0.093
Realized CCC	0.135	0.150	0.072	0.085	0.107	0.121	0.153	0.118	0.091
Baseline Model	0.140	0.157	0.074	0.089	0.110	0.127	0.167	0.126	0.093
MCS size					1%		5%		10%

Table A.6.8: Beta hedging analysis.

A.7 Fama-French Risk Premia at High Frequency

We estimate risk premia from the cross-section of individual assets at a daily and a monthly frequencies. Daily risk premia is simply estimated with the series of daily cross-section regressions (1.31) using daily excess returns and daily factor loadings filtered from the realized factor model. The monthly premia is estimated with an assumption that the risk premia $\lambda_{0,t}$ and $\lambda_{1,t}$ are locally constant within a month. Therefore, we use monthly excess returns and monthly averaged model-implied factor loadings[3] to estimate (1.31) on a monthly basis.

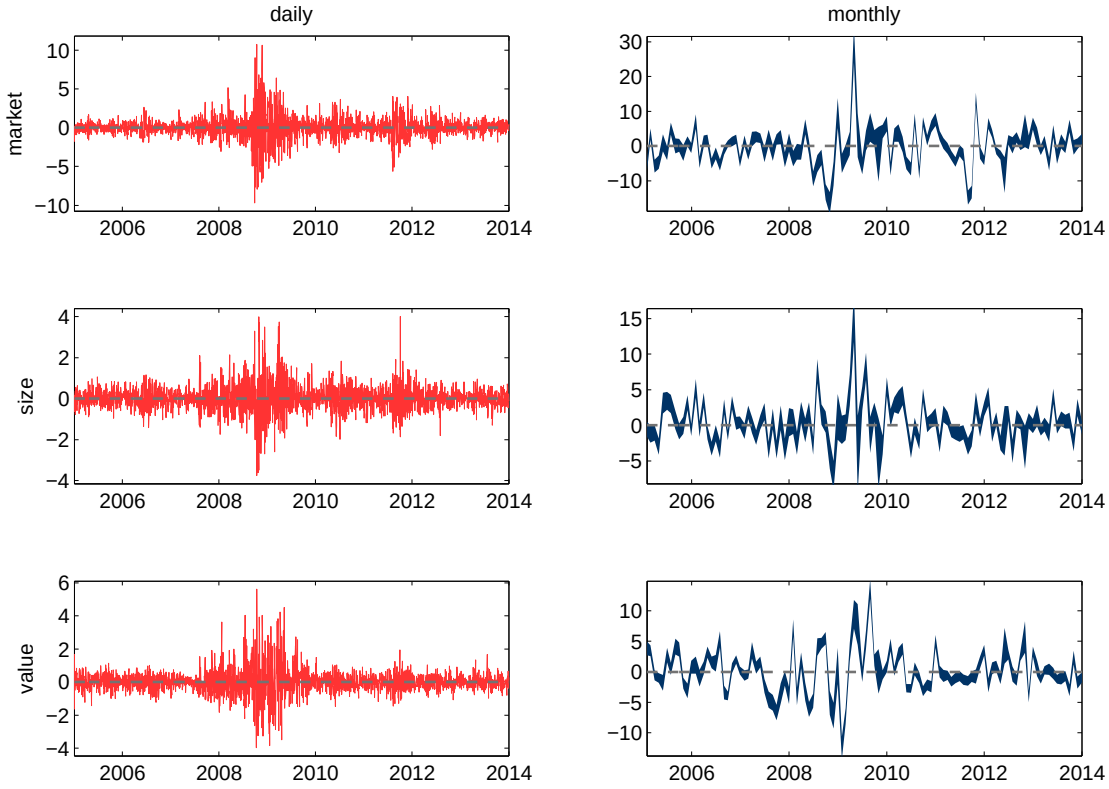

Figure A.7.4: Daily and monthly risk premia estimates (in %) based on the cross-section of 824 stocks and factor loadings filtered from the realized factor model (January, 2005 - December, 2013). For the monthly premia estimates (right column), we plot 95% confidence intervals obtained with robust regression standard errors.

Both the daily and the monthly premia estimates are noisy and not persistent when we obtain them from a sequence of cross-section regressions. This is not surprising because of many adverse factors that are especially pronounced at high frequencies and potentially distort the estimates. These factors include a strong autocorrelaton in the factor loading measurement errors, a large extent of idiosyncratic variation in individual asset returns, the presence of outliers, the error-in-variable attenuation due to the imperfect or not perfectly consistent measurement of factor loadings, etc. Thus, the true unobservable risk premia (as well as the implied pricing kernel) might have more regular and predictable dynamics. This motivates the development of a more elaborate econometric framework for an extraction of the latent dynamic risk premia.

[3]We ignore the compounding effect that causes a slight inconsistency between the averaged daily betas and monthly returns.

Chapter 2

A Realized Dynamic Nelson-Siegel Model with an Application to Crude Oil Futures Prices

Joint with Bo Laursen

2.1 Introduction

Throughout the last decade commodities have become a popular asset class among investors. The commodity markets have experienced a rapid growth in both trading activity and volume due to large inflows of investment capital. Such financialization has directly affected commodity price dynamics and volatility[1] and, as a consequence, has emphasized the role of interactions between commodity prices and macroeconomic variables.[2] Furthermore, the financial crisis spurred investors interest in commodities since they provide a natural alternative for financial securities.[3] Analyses of commodity prices have, therefore, become of increasingly important both from an academic perspective and for practitioners.

The predominant amount of literature related to modeling and forecasting commodity prices focuses on point forecasts, which are informative only about the central tendency of predictions. In contrast to that, density forecasts provide rich information about uncertainty around the central tendency and allow to invoke many additional useful statistics such as interval predictions, upside and downside risk assessments, etc. This appealing feature recently caused a burst of interest in density forecasting of economic and financial variables among researchers and policymakers (see e.g. Hall and Mitchell (2007), Clark (2011)). The density dynamics is of special interest when dealing with commodity markets. The poor point predictability of commodity spot and futures prices enhances the importance of having reliable density forecasts. To accurately evaluate uncertainty associated with point forecasts is of great practical value for risk management purposes and might also be useful for central bankers who are interested in relations between the dynamics of commodity prices and macroeconomic variables.

In this paper, we analyze the role of time-varying volatility in modeling commodity futures prices from the standpoint of forecasting. For this purpose we develop a class of models that allow us to capture the futures

[1] Empirical evidence on the impact of the commodity market financialization on the prices is given in Buyuksahin and Robe (2014), among others. An analysis of such impact from the theoretical perspective is suggested in Basak and Pavlova (2015) (see also references therein).

[2] The links between commodity prices and a real economic activity are discussed in Kilian (2008), Kilian (2009), Collier and Goderis (2012) as well as in many other papers.

[3] An attractive diversification potential of commodities was pointed out in Gorton and Rouwenhorst (2006), for example. However, recent empirical studies provide a cautious assessment of such diversification and hedging abilities identifying an increasing integration between commodity and equity markets around the financial crisis episode (see Ravazzolo and Lombardi (2012), Hansen et al. (2014a), Bhardwaj et al. (2015), etc.).

price dynamics for the full spectrum of maturities. These models perform well both in and out of sample and simultaneously account for the time-varying volatility. The latter is the key element in our framework that allows capturing changes in the second moment of distribution dynamics. We examine to what extent time-varying volatility is able to improve the density predictability and whether we gain from the use of realized measures of volatility.

Commodities are by nature very different from financial securities. The vast majority of trading in commodities happens through exchange traded futures contracts. Analyzing commodity prices therefore often involve modeling the whole term structure of futures prices. A standard approach for modeling the term structure is to specify a dynamic model for the underlying spot price[4] and derive the futures price based on no arbitrage arguments (see Brennan and Schwartz (1985), Schwartz (1997), Geman (2005), etc.). In this paper, we make use of another approach. Instead of building a model based on the spot price dynamics, we explicitly assume a functional form for the term structure curve.

We apply the dynamic Nelson-Siegel (DNS henceforth) factor approach of Diebold and Li (2006). The approach is based on a dynamic version of the classical Nelson-Siegel interest rate model, which is specifically suited for forecasting yield curves. The data on commodity futures share many similarities with interest rate data. Hence, the DNS framework can be directly applied to commodity markets. In this case, the Nelson-Siegel approach is based on the assumption that the futures prices can be described by three unobserved factors interpreted as the level, slope, and curvature of the futures curve. The factors are assumed to follow particular dynamics that allow forecasting the whole term structure of futures prices.

The futures returns of the majority of commodities show a clear sign of time-varying volatility. The turbulence of energy prices around the financial crisis and the recent oil price downfall provide fresh examples of periods with excess volatility. The volatility dynamics of many commodity prices also tend to cluster, that is demonstrating a cyclical behavior where the moderately volatile time intervals alternate with the periods of high price variation.

We introduce time-variation directly into the conditional variance of the three dynamic factors. Therefore, both the conditional mean vector and the conditional covariance matrix of the futures term structure possess factor dynamics determined by the time-varying mean and volatility components of the Nelson-Siegel factors respectively. This factor approach effectively reduces the dimensionality of the problem and allows to parsimoniously model the joint density dynamics of highly multivariate futures prices. Moreover, such parsimonity likely benefits the forecasting performance of the model. Compared with the constant volatility frameworks, time-varying aspects allow to assess more accurately an increased degree of the forecast uncertainty throughout the periods of price turbulence and not to overestimate such uncertainty during the calm periods. We note that an additional challenge comes from the fact that we model the volatility of unobserved variables (Nelson-Siegel factors) opposed to the more standard frameworks (e.g., VAR models with time-varying volatility) where only the volatility is unobserved.

Furthermore, we introduce realized measures of volatility into the model. While doing this we closely follow the Realized GARCH framework of Hansen et al. (2012). Namely, we not only insert the realized measures directly into the conditional dynamics of the volatility components, but also "close" our model by adding measurement equations for the realized measures. The use of precise volatility information extracted from high frequency data often provides a better signal about the current level of volatility than what could otherwise be extracted from the lower frequency data. Hence, models that exploit realized measures should be better at "catching up" with the hidden volatility dynamics and therefore more suited to periods where the volatility changes rapidly. We expect these aspects to improve the density forecasting properties of the model.

We apply our modeling framework to the futures prices of the most traded commodity from the energy sector - WTI light sweet crude oil. We demonstrate that ime-varying volatility significantly improves density forecasting of the oil futures prices. On top of that, the use of realized volatility measures provides moderate in

[4]In addition, modeling commodity prices potentially involves specifying a certain stochastic variable, which reflects benefits/costs to the holder coming from the possession of a commodity. For example, the net convenience yield is a usual interpretation of such variable.

magnitude, but systematic improvements in the density prediction. We also find that including several dynamic components to control the conditional variance of the DNS factors is preferable to using just a single factor.

The rest of the paper is organized as follows. In the remaining part of this section we provide a brief review of the related literature. In Section 2 we introduce our main modeling framework. In Section 3 we describe the futures data used in the empirical analysis. We discuss the choice and construction of realized volatility measures from the high frequency intra-daily price observations in Section 4. We outline the model evaluation strategy in Section 5 and Section 6 provides the corresponding results. Section 7 concludes and all supplementary materials are found in the Appendices.

2.1.1 Related Literature

Our paper refers to several strands of the literature. First of all, it relates to the area of modeling and forecasting term structure curves with dynamic factor models. Pioneered by Diebold and Li (2006) and Diebold et al. (2006) DNS models have become a popular framework for modeling and forecasting term structures of interest rates. The original models primarily focused on point forecasting and assumed a constant conditional variance for the term structure curve.

Recently, several versions of the DNS model with time-varying volatility have been proposed in the literature. Koopman et al. (2010) added a time-varying volatility component to the idiosyncratic conditional variances of the term structure. Alternatively, Bianchi et al. (2009), Hautsch and Ou (2012), Shin and Zhong (2015) introduced stochastic volatility directly into the dynamics of the common Nelson-Siegel factors. Applied to the interest rate yield dynamics, these models confirmed the usefulness of a time-varying volatility for in-sample fit and revealed some gains for point and density predictability. In our paper, we suggest a modification of the DNS modeling framework, which introduces an observation driven[5] conditional volatility into the Nelson-Siegel factors and exploits intra-daily realized volatility measures.

The DNS model has previously been applied to commodity markets. Karstanje et al. (2015) examine comovements in different commodities by assuming that the factors of each commodity can be decomposed into a market, sector, and idiosyncratic components. Hansen and Lunde (2013) analyze futures prices on oil by modeling the factors using a GARCH model with Normal Inverse Gaussian innovations within a copula framework. The method provides a flexible model suitable for matching heavy tails in the conditional density of the commodity prices. Although we rely on a conditionally Gaussian density, we note that by simply allowing the volatility of the factor innovations to be time-varying the model may also replicate some of the heavy-tailed aspects of futures returns. The procedure in Hansen and Lunde (2013) demands extracting the latent factors in a first step and subsequently treating them as if they were observable. This method becomes problematic when the cross-sectional dimension of the futures data is sparse. In our approach the factors are extracted directly using the Kalman filter. Thus, while filtering the latent factors, we effectively exploit both cross-sectional and time-series dimensions of the data.

Secondly, we add to the spacious and still growing area of volatility models that exploit realized measures as an extra source of information about the latent volatility level. The idea originates from Andersen and Bollerslev (1998b) where the realized variance based on intra-daily asset returns was used as an ex-post return variance proxy to assess the performance of the ARCH-type dynamic models. Shortly thereafter, realized variance measures based on high frequency data were directly introduced into the dynamic volatility specifications in (G)ARCH models (Engle (2002)) and stochastic volatility models (Barndorff-Nielsen and Shephard (2002b)). The more accurate estimation and the forecasting gains associated with the inclusion of realized measures into the modeling frameworks inspired a rapidly growing research activity in this area.[6]

[5]According to the classification in Cox (1981) (see also a related discussion in Koopman et al. (2015)), an observation driven specification of a time-varying parameter implies that its current realization is a deterministic function of its lagged realizations and lagged observable variables (the most notable examples include (G)ARCH, ACD and Dynamic Conditional Score (DCS) models). In contrast, a parameter driven specification suggests that a time-varying parameter is driven by its own stochastic innovations (as in unobserved component models, stochastic volatility models, etc.).

[6]Some selected examples related to asset return modeling include observation driven MEM of Engle and Gallo (2006), HEAVY model of Shephard and Sheppard (2010) as well as parameter driven (stochastic volatility) models of Takahashi et al. (2009),

The way of volatility modeling suggested in this paper is closely related to the Realized GARCH model of Hansen et al. (2012). We also incorporate realized measures into the GARCH dynamics of the latent conditional volatility and use measurement equations to directly establish a link between realized measures and the volatility process. Loosely speaking, our framework can be considered as a multivariate version of the Realized GARCH where the state-space DNS factor model plays a role of the return equation linking the dynamic volatility components with the conditional distribution of the low frequency price vector. Naturally, the presented framework can be easily extended for the case of a more general dynamic factor model in a place of the return equation.[7]

A closely related paper from a methodological perspective is Shin and Zhong (2015). They incorporated realized measures into the DNS framework with stochastic volatility and documented the corresponding gains for bond yield density prediction. Similarly to their framework, we introduce a time-varying conditional volatility into the Nelson-Siegel factor dynamics and link it with realized measures of volatility obtained from observations of a higher frequency. Our modeling strategy, however, is different in several aspects. Opposed to the prevailing practice in the literature, we model the conditional volatility dynamics not as a stochastic volatility process, but in an observation driven manner. The stochastic volatility models often lead to non-linear non-Gaussian state space models and thereby complicate the estimation dramatically. In contrast, our observation driven specification can be estimated straightforwardly by maximum likelihood via the Kalman filter. Moreover, we suggest the use of *direct* realized volatility measures for the latent Nelson-Siegel factor innovations to avoid imposing structural non-linear relations between the vector of term structure realized volatility measures and the three volatility components related to the Nelson-Siegel factors, as in Shin and Zhong (2015).

Finally, we contribute to the empirical literature on the forecasting of commodity futures prices. In doing so, we apply our modeling framework to the prices of futures contracts and evaluate how well the model is able to predict the conditional density for the whole range of maturities. So far, the literature related to density forecasting of commodities is scarce and the corresponding findings are not very conclusive. Using the DNS framework Hansen and Lunde (2013) found some improvements from the time-varying volatility for the point forecasts of crude oil futures prices and for the corresponding value-at-risk forecasts. Lunde and Olesen (2014) used the Realized GARCH framework and documented only slight gains in the forecasting of return densities of univariate electricity forwards from the use of realized volatility measures. Despite that our framework comprises certain elements from these two settings, it differs from both. On one hand, we complement to Hansen and Lunde (2013) by focusing on the implications of time-varying volatility on the density prediction and by integrating realized volatility measures into the framework. On the other hand, although we also exploit realized measures and invoke Realized GARCH dynamics in our study, we model the joint multivariate conditional density dynamics of the whole term structure and not only the univariate density of the shortest contract as in Lunde and Olesen (2014).

2.2 Modeling Strategy

2.2.1 Nelson-Siegel Framework

We model the term structure of futures contracts within a dynamic Nelson-Siegel framework. In their original paper Nelson and Siegel (1987) suggested a parsimonious and flexible three-component polynomial approximation for the term structure of Treasury bond yield curves. The choice of polynomial functions was associated with the solution to a second-order differential equation and provided an ability to fit a range of shapes typically revealed by yield curves including monotonic, hump-shaped and inverted hump-shaped patterns.

Subsequently, Diebold and Li (2006) interpreted the Nelson-Siegel framework as a factor model where the factor loadings are defined by the Nelson-Siegel polynomials and are treated as latent dynamic factors. Such extension made it possible to filter the dynamically evolving latent Nelson-Siegel factors and to produce out-of-sample forecasts of the entire yield curve. The class of dynamic Nelson-Siegel factor models has been then

Koopman and Scharth (2012), etc.

[7] From this perspective, our model is a hybrid of the Realized GARCH of Hansen et al. (2012) and the framework of Harvey et al. (1992) (or the Factor GARCH of Engle et al. (1990) where the common factors are latent and have to be extracted).

successfully applied to modeling and forecasting different types of assets and financial instruments with a term structure embedded.

We denote the logarithmic price at time period t of a contract with τ periods until expiration by $y_t(\tau)$. Assume that at time t there are n_t types of contracts in the market that differ in their times until expiration: $\tau_{t,1} < \tau_{t,2} < ... < \tau_{t,n_t}$. Note that n_t likely changes with t since the number of contracts available for trade is varying from day to day. We also introduce the $n_t \times 1$ vector $y_t(\tau_t)$, which collects the futures prices for all n_t contracts traded at time t, $y_t(\tau_t) = (y_t(\tau_{t,1}), y_t(\tau_{t,2}), ..., y_t(\tau_{t,n_t}))'$. In accordance with the dynamic Nelson-Siegel framework, a scalar log-price component $y_t(\tau)$ for some given τ obtains the following factor representation[8]

$$y_t(\tau) = f_{l,t} + \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau}\right) f_{s,t} + \left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau}\right) f_{c,t} \tag{2.1}$$

where $f_{l,t}$, $f_{s,t}$ and $f_{c,t}$ are latent dynamic factors and $\lambda > 0$ is a parameter of the Nelson-Siegel factor loadings.[9]

The model in (2.1) links futures prices with the maturities by means of the three common factors and the associated factor loadings. The loadings to the first factor $f_{l,t}$ are constant across the whole range of maturities and a change in $f_{l,t}$ translates to the same change for all prices in y_t. Factor $f_{l,t}$ is therefore usually interpreted as a level of the term structure curve. The second factor, $f_{s,t}$, has loadings which are monotonically decreasing as time to maturity τ grows, so it is commonly interpreted as a slope of the term structure. Parameter λ regulates the steepness of this slope. The higher λ is, the faster decay is exhibited by the loading function. The loadings to the factor $f_{c,t}$ also depend on τ and λ. They converge to 0 for $\tau \to 0$ and $\tau \to \infty$ and reach a maximum at some intermediate value of τ. Parameter λ governs the location of the extremum. When λ gets larger, the peak of the loading function shifts towards lower values of τ. This hump-shaped loading function associated with $f_{c,t}$ leads to an interpretation of this factor as the curvature of the term structure. Whereas a change in the level factor $f_{l,t}$ affects all the contracts in y_t uniformly, an increase in the slope factor $f_{s,t}$ imposes an effect mainly on short-term contract prices and an increase in the curvature factor $f_{c,t}$ increases prices for mid-term contacts.

The dynamic Nelson-Siegel framework treats factors as time-varying processes, which explain the variation in the term structure. Most of the related literature, however, is focused exclusively on the dynamics of the conditional mean of the factor process. As a consequence, an associated analysis is usually limited to modeling a dynamic location of the term structure variable and its point forecasting, so the uncertainty around the extracted and predicted term structure curves is left uncovered. This paper shifts the focus to the dynamic conditional variance of the factors. Such aspect naturally allows to model and forecast the conditional density dynamics of the term structure curve, which is implied by the conditional density dynamics of Nelson-Siegel factors.

2.2.2 Dynamic Nelson-Siegel Model with Time-varying Volatility and Realized Measures

2.2.2.1 Dynamics of Conditional Mean

Following Diebold et al. (2006) we cast the dynamic Nelson-Siegel model from (2.1) in a state-space representation.[10] The observation equation in a matrix form is

[8] In equation (2.1) we use a slightly reparametrized version of the original Nelson-Siegel curve, as it is suggested in Diebold and Li (2006).

[9] We note that parameter λ can be considered as a time dependent parameter, λ_t. In most of related studies, however, λ is treated as a constant parameter. Although being convenient for estimation, such assumption is quite restrictive, especially if the time span of the analyzed period is sufficiently long. There are several papers where time-varying λ_t was examined in the context of term structure modeling of interest rates (see Creal et al. (2008) (section 4.1.3), Koopman et al. (2010)). Nevertheless these studies do not provide an unambiguous answer on whether the gains from time-varying λ_t are significant for the model fit and forecasting. Due to the different focus of the present paper, we will treat λ as a constant parameter as it is common in the literature on term structure modeling.

[10] State-space representation of a dynamic Nelson-Siegel model is closely related to the traditional factor model representation (see Stock and Watson (2011) for a related survey). According to the factor model interpretation, as it is conventional in the literature, term $\Lambda_t f_t$ from (2.2) is called a common component and ε_t is an idiosyncratic component. A part of variation in y_t explained by the common component is due to the latent factors f_t which are common for all variables in y_t. The idiosyncratic

$$y_t = \Lambda_t f_t + \varepsilon_t \tag{2.2}$$

where $f_t = (f_{l,t}, f_{s,t}, f_{c,t})'$ is a 3×1 vector of dynamic latent factors, $\varepsilon_t = (\varepsilon_{t,1}, \varepsilon_{t,2}, ..., \varepsilon_{t,n_t})'$ is a $n_t \times 1$ vector of idiosyncratic components, and $\Lambda_t = \Lambda_t(\lambda; \tau_{t,1}, \tau_{t,2}, ..., \tau_{t,n_t})$ is $n_t \times 3$ matrix of factor loadings

$$\Lambda_t = \begin{pmatrix} 1 & \frac{1-e^{-\lambda\tau_{t,1}}}{\lambda\tau_{t,1}} & \frac{1-e^{-\lambda\tau_{t,1}}}{\lambda\tau_{t,1}} - e^{-\lambda\tau_{t,1}} \\ 1 & \frac{1-e^{-\lambda\tau_{t,2}}}{\lambda\tau_{t,2}} & \frac{1-e^{-\lambda\tau_{t,2}}}{\lambda\tau_{t,2}} - e^{-\lambda\tau_{t,2}} \\ \vdots & \vdots & \vdots \\ 1 & \frac{1-e^{-\lambda\tau_{t,n_t}}}{\lambda\tau_{t,n_t}} & \frac{1-e^{-\lambda\tau_{t,n_t}}}{\lambda\tau_{t,n_t}} - e^{-\lambda\tau_{t,n_t}} \end{pmatrix} \tag{2.3}$$

The state equation describes the dynamics of latent factors. We suggest a simple random walk specification[11]

$$f_t = f_{t-1} + \eta_t \tag{2.4}$$

where η_t is a 3×1 shock to the state vector f_t orthogonal to ε_t. We select a random walk for several reasons. At first, it is an appealing choice due to its parsimony and is a pretty suitable assumption in our context, since we mainly focus on the time-varying volatility of the futures term structure and not on the dynamics of its conditional mean. Such simple specification not only facilitates the estimation process, but is also less prone to overfitting that may have an adverse effect on the out-of-sample model performance. At second, the Nelson-Siegel factors extracted from commodity futures prices on a daily basis are, in fact, close-to-unit-root processes (see Hansen and Lunde (2013) for an empirical evidence).

We assume that the idiosyncratic component ε_t is an autoregressive process which reads

$$\varepsilon_t = \beta C_t \varepsilon_{t-1} + w_t \tag{2.5}$$

where C_t is a selection matrix of dimensions $n_t \times n_{t-1}$ filled with zeros and ones. An entry (i, j) of matrix C_t is equal to 1 if some contract that is traded at $t-1$ (with maturity $\tau_{t-1,j}$) is also traded at t (with maturity $\tau_{t,i}$ which is equal to $\tau_{t-1,j} - 1$). If otherwise, an entry (i, j) of matrix C_t is equal to 0. Parameter $\beta \in (0, 1)$ is a constant scalar that determines the persistence of the price misspecification generated by the model. We then suppose that $w_t \sim$ i.i.d.$N(0, \Sigma_t)$, where Σ_t is a constant diagonal matrix with time-dependent dimensions $n_t \times n_t$. We use a simple parametrization $\Sigma_t = \sigma_w^2 I_{n_t}$, where I_{n_t} is the identity matrix of size n_t.[12]

The conditional mean dynamics of the futures prices suggested in this subsection have an appealing and intuitive interpretation. The common Nelson-Siegel term $\Lambda_t f_t$ from (2.2), which is governed by the random walk factors, can be treated as a persistent long-run dynamic component of the futures price vector y_t. The autoregressive idiosyncratic component ε_t, in turn, can be related to the transitory short-run fluctuations in y_t. As a result, our model implies that the futures prices are mean-reverting to its long-run (equilibrium) dynamics.[13]

2.2.2.2 Dynamics of Conditional Variance

Time-varying volatility is a central element in our analysis. We introduce it directly into the dynamics of the Nelson-Siegel factors

$$\eta_t \big| \mathcal{F}_{t-1} \sim N(0, \Omega_t) \tag{2.6}$$

component is a zero mean stationary process, which comprises variable specific variation of y_t that is not explained by the common factors and absorbs measurement errors.

[11] Note that a more general VARMA process (with proper identification restrictions) can be used to model transition dynamics within a linear state-space framework. In particular, a popular choice is to assume a first-order autoregressive process as in Diebold et al. (2006).

[12] We note that the time dependence in Λ_t, C_t, and Σ_t is related to the number of contracts n_t traded at a given day t and their maturities $\{\tau_{t,1}, \tau_{t,2}, ..., \tau_{t,n_t}\}$. Since this information is exogenous in our model, matrices Λ_t, C_t and Σ_t change with time in a deterministic way.

[13] From this perspective our specification is conceptually similar to the modeling framework of Schwartz and Smith (2000).

where Ω_t is a 3×3 time-varying conditional covariance matrix of the factor increment $\Delta f_t = \eta_t$ and $\mathcal{F}_t$ is a filtration that contains all the relevant observable information up to day t inclusively. This way of modeling imposes a factor structure on the conditional variance of the multivariate price process, y_t

$$\text{Var}\big(y_t \big| f_{t-1}; \mathcal{F}_{t-1}\big) = \Lambda_t \Omega_t \Lambda_t^{'} + \Sigma_t \tag{2.7}$$

Thus, the large $n_t \times n_t$ conditional covariance matrix of the whole term structure of futures prices $\text{Var}\big(y_t \big| f_{t-1}; \mathcal{F}_{t-1}\big)$ is fully characterized by the deterministic matrices Λ_t and Σ_t as well as the small scale 3×3 time-varying conditional covariance matrix Ω_t.

A similar approach was used in Bianchi et al. (2009) and Hautsch and Ou (2012) among others, where stochastic volatility was introduced directly into the dynamics of the Nelson-Siegel factors. Koopman et al. (2010) pointed out that this approach is somewhat restrictive since in this case the Nelson-Siegel loading matrix Λ_t serves as a weighting matrix for both conditional mean and variance of elements in y_t. As an alternative, they suggested to endow a covariance matrix of idiosyncratic components Σ_t with a time-varying volatility factor, so each element from y_t has a time-varying volatility component.

We, nonetheless, maintain the assumption from (2.6) and incorporate time-varying conditional volatility in the factor dynamics. In doing so, we are motivated by the argument that the total variation of the term structure variable y_t is mainly explained by the common factor dynamics f_t and only to a little extent by the idiosyncratic component ε_t. Therefore, uncertainty associated with the futures prices y_t is largely incorporated in f_t. As a result, we may expect that time-varying volatility of the factors is more relevant for density modeling and forecasting than time-varying volatility in the idiosyncratic term ε_t. Moreover, these two approaches have been compared in Shin and Zhong (2015) in a similar study in the context of bond yield density prediction. Their results confirm that whereas time-varying volatility in factor dynamics helps to sufficiently improve the model performance, dynamic components in the variance of the idiosyncratic term do not seem to play a significant role.

We use a constant conditional correlation (Bollerslev (1990)) structure to parametrize Ω_t and treat the conditional factor volatilities as time-varying processes. Thus, the conditional variance matrix can be decomposed as follows

$$\Omega_t = H_t R H_t = H_t \begin{pmatrix} 1 & \rho_{ls} & \rho_{lc} \\ \rho_{ls} & 1 & \rho_{sc} \\ \rho_{lc} & \rho_{sc} & 1 \end{pmatrix} H_t \tag{2.8}$$

where R is a constant correlation matrix and H_t is a diagonal matrix with conditional volatilities of the factor increments Δf_t on the main diagonal. Therefore, the dynamics of Ω_t is fully driven by the dynamics of the 3×3 diagonal matrix H_t. We will suggest two alternative dynamic specifications for H_t - with one and three time-varying parameters.

One of the key features of our framework is that we use realized measures of variance based on intra-daily high frequency data while formulating the dynamics for the time-varying parameters from Ω_t. The use of realized measures of volatility is commonly motivated by the precise information content of the high frequency data (e.g., Dobrev and Szerszen (2010)). We, thus, expect that the realized signals of the futures price volatility will help to formulate and extract the latent conditional variance dynamics more accurately and to improve the predictive performance of the model.

We incorporate realized measures into the model along the lines of Realized GARCH of Hansen et al. (2012). Thus, we formulate a transition (GARCH) equation that links the latent volatility with own lags and lagged realized measures. The realized measures, in turn, are linked to the latent volatility through the specially formulated measurement equation. The measurement equation provides a contemporaneous relationship between the latent dynamic variable and its realized proxy allowing for a stochastic measurement error. The role of such measurement equation, therefore, is to directly relate the latent volatility process to the corresponding measures and, thus, to incorporate the precise information about the conditional factor variance Ω_t into the model. Apart

from that, this equation may account for a systematic measurement bias and can be accommodated with some stylized features of financial data, such as the leverage effect.

Note that another possibility could be to specify the dynamics for all six distinct elements of the 3×3 matrix Ω_t, thus reflecting a time variation in both conditional variances and correlations.[14] However, an approach with fully specified dynamics of Ω_t is hardly implementable in our application due to the difficulties with the extraction of a full realized matrix of Nelson-Siegel factor covariances at intra-daily frequency. We describe these challenges in details in the section related to the construction of realized volatility measures.

2.2.2.3 One-component Specification for the Conditional Variance Matrix

We start with a simple parametrization of the conditional covariance matrix Ω_t. We assume that the diagonal volatility matrix H_t from (2.8) depends on a single time-varying parameter controlling the volatilities of all 3 factors

$$H_t = \begin{pmatrix} \sqrt{\exp(h_t)} & 0 & 0 \\ 0 & \sqrt{a_s + b_s \exp(h_t)} & 0 \\ 0 & 0 & \sqrt{a_c + b_c \exp(h_t)} \end{pmatrix} \tag{2.9}$$

where a_s, b_s, a_c, b_c are non-negative scaling coefficients and h_t is a scalar time-varying parameter. Therefore, the variances of all three Nelson-Siegel factors are linear functions of the single time-varying variance component $\exp(h_t)$. Note that for identification we directly set $\exp(h_t)$ equal to the variance of the level factor. An exponential transformation simplifies the formulation of the transition and measurement equations for h_t since for any real values of h_t the positivity of the factor variances is preserved. Closely following the specification of the Realized GARCH model we set up the transition (GARCH) equation as an autoregressive process[15]

$$h_t = \gamma_0 + \gamma_1 s_{t-1} + \gamma_2 h_{t-1} \tag{2.10}$$

where s_t is an appropriately chosen realized measure for h_t. The measurement equation is formulated as

$$s_t = \xi + \phi h_t + \delta(z_t) + u_t \tag{2.11}$$

where $u_t \sim N\left(0, \sigma_u^2\right)$ is a stochastic measurement error, ξ and ϕ are constant parameters which determine the systematic bias of the realized measure s_t, $\delta(\cdot)$ is the so-called leverage function that introduces a dependence between variance and returns of futures prices, z_t is a suitable $\mathcal{F}_t$-measurable variable that reflects the price movement occurred at day t. We discuss the choice of z_t below.

An empirical usefulness of the leverage component introduced in such way was confirmed in Hansen et al. (2012) and Hansen and Huang (2016) in the context of modeling the conditional variance dynamics of asset returns. A standardized asset return plays the role of variable z_t there, so the measurement equation explicitly establishes the link between the conditional variance and price shocks. The specification for the leverage component $\delta(\cdot)$ suggested in Hansen et al. (2012) is the following polynomial function

$$\delta(z_t) = \delta_1 z_t + \delta_2 (z_t^2 - 1) \tag{2.12}$$

where δ_1 and δ_2 are constant parameters. Such functional form is convenient for several reasons. As long as z_t is a standardized random shock, the leverage component $\delta(z_t)$ has zero mean. It, therefore, does not induce an additional bias for s_t in the measurement equation (2.11). Also, the quadratic specification in (2.12) allows to capture asymmetry effects in the volatility responses to the price shocks.

[14] For example, a proper dynamics for all elements from Ω_t (which preserves a positive definiteness of Ω_t) can be formulated along the lines of Chiriac and Voev (2011) or Golosnoy et al. (2012). In the former case, the vector of Cholesky elements of Ω_t is modeled and then these elements can be related to the corresponding realized measures. In the latter case, Ω_t is treated as a matrix process and the realized measures of Ω_t are assumed to be distributed as a Wishart matrix random variable.

[15] In this paper, we consider only the parsimonious (1,1) specification for lags. The lag order can be extended if needed.

The one-component specification of the conditional variance matrix Ω_t given by (2.8) and (2.9) implies that

$$\mathrm{Var}\big(y_t|f_{t-1};\mathcal{F}_{t-1}\big)= \Lambda_t H(h_t) R H(h_t) \Lambda_t^{'} + \Sigma_t \tag{2.13}$$

Particularly, it means that conditional variances of all futures contracts are driven by the single dynamic volatility component h_t. As a consequence, we may expect that the realized variances of the futures prices can serve as good signals about the unobservable volatility state and, thus, can be related to h_t. More specifically, for some selected $\tau \in \{\tau_{t,1}, \tau_{t,2}, ..., \tau_{t,n_t}\}$[16]

$$s_t \equiv \log RV\big(\Delta y_t(\tau)\big) \simeq \log \mathrm{Var}\big(y_t(\tau)|f_{t-1};\mathcal{F}_{t-1}\big) \tag{2.14}$$

For such choice of s_t we suggest the following choice of variable z_t to establish a dependence between volatility component h_t and the futures price change

$$z_t = \frac{\Delta y_t(\tau) - \mathrm{E}\big(\Delta y_t(\tau)|f_{t-1};\mathcal{F}_{t-1}\big)}{\sqrt{\mathrm{Var}\big(\Delta y_t(\tau)|f_{t-1};\mathcal{F}_{t-1}\big)}} \tag{2.15}$$

so z_t is a conditionally standardized $\mathcal{F}_t$-measurable Gaussian random variable, which represents a price shock to the contract τ.

Summarizing, equations (2.9), (2.10), and (2.11), where s_t and z_t are defined as in (2.14) and (2.15) respectively, constitute the one-component dynamic specification for the conditional covariance matrix Ω_t.

2.2.2.4 Three-component Specification for the Conditional Variance Matrix

We also suggest a richer dynamic parametrization for Ω_t where three time-varying parameters are introduced. We link these dynamic components with the three conditional variances of the level, slope and curvature. The matrix H_t reads

$$H_t = \begin{pmatrix} \sqrt{\exp(h_{l,t})} & 0 & 0 \\ 0 & \sqrt{\exp(h_{s,t})} & 0 \\ 0 & 0 & \sqrt{\exp(h_{c,t})} \end{pmatrix} \tag{2.16}$$

where $h_{l,t}$, $h_{s,t}$ and $h_{c,t}$ are the three time-varying logarithmic conditional variances of the 3-variate shock η_t to the dynamic factor process f_t. We denote a 3×1 vector of dynamic volatility components as $h_t = \big(h_{l,t}, h_{s,t}, h_{c,t}\big)'$ and note that $\exp(h_t) = diag\big(\Omega_t\big)$. Similar to the previous specification, we write the transition equation as

$$h_t = \Gamma_0 + \Gamma_1 s_{t-1} + \Gamma_2 h_{t-1} \tag{2.17}$$

where $s_t = \big(s_{l,t}, s_{s,t}, s_{c,t}\big)'$ is a 3×1 vector of the proper realized measures for h_t and Γ_0, Γ_1, Γ_2 are constant parameter matrices of suitable dimensions. The measurement equation is defined as

$$s_t = \zeta + \Phi h_t + D(z_t) + u_t \tag{2.18}$$

where $u_t \sim N\big(0, \Sigma_u\big)$ is a 3×1 vector of measurement errors, $D(\cdot)$ is a multivariate leverage function, z_t is a $\mathcal{F}_t$-measurable 3×1 vector and ζ, Φ are constant matrices of proper dimensions. The leverage function in a multivariate version has the following form

$$D(z_t) = D_1 z_t + D_2(z_t \circ z_t - \iota_3) \tag{2.19}$$

where D_1, D_2 are 3×3 parameter matrices and ι_3 is the 3×1 vector of ones.

[16]Note that several realized measures related to distinct contracts $\tau \in \{\tau_{t,1}, \tau_{t,2}, ..., \tau_{t,n_t}\}$ can be used to extract h_t. In this case, in analogy with Hansen and Huang (2016), several measurement equations (2.11) which correspond to different realized measures $s_t(\tau)$ have to be formulated.

We attempt to directly extract realized measures s_t of the latent factor variances from intra-daily data by exploiting the recent advances in high frequency volatility estimation. The description of the estimation procedure is referred to the next section. Thus, signals s_t in this case are the realized variance measures for Nelson-Siegel factor increments Δf_t

$$s_{i,t} \equiv \log RV\big(\Delta f_{i,t}\big) \simeq \log \mathrm{Var}\big(\Delta f_{i,t} | f_{t-1}; \mathcal{F}_{t-1}\big), \qquad i \in \{l, s, c\} \tag{2.20}$$

We define the corresponding 3×1 vector z_t, which is required for the leverage function, as

$$z_t = \mathrm{Var}\big(\mathrm{E}\big(f_t|\mathcal{F}_t\big) - \mathrm{E}\big(f_{t-1}|\mathcal{F}_{t-1}\big)|\mathcal{F}_{t-1}\big)^{-\frac{1}{2}} \big(\mathrm{E}\big(f_t|\mathcal{F}_t\big) - \mathrm{E}\big(f_{t-1}|\mathcal{F}_{t-1}\big)\big) \tag{2.21}$$

so z_t is a conditionally standard Normal random vector.

As a result, the three-component specification for the conditional covariance matrix Ω_t is represented by equations (2.16), (2.17) and (2.18) with s_t and z_t specified in (2.20) and (2.21) respectively. In our empirical analysis we stick to the parsimonious model parametrization. In particular, matrices $\mathit{\Gamma}_1$, $\mathit{\Gamma}_2$, Φ, D_1, D_2 from (2.18)-(2.20) are assumed to be diagonal.

2.2.3 Estimation

Suppose that we observe futures prices y_t and realized measures s_t within a sample period $t = 1, ..., T$. The joint conditional density function of observable variables for some given trading day t can be factorized as follows

$$p(y_t, s_t|\mathcal{F}_{t-1}; \theta) = p(y_t|\mathcal{F}_{t-1}; \theta) p(s_t|y_t, \mathcal{F}_{t-1}; \theta) \tag{2.22}$$

where $\mathcal{F}_t = \sigma(y_t, s_t, y_{t-1}, s_{t-1}, ...)$ is a σ-field generated by all observables up to t and θ is a vector of model parameters.

The first term in the decomposition, $p(y_t|\mathcal{F}_{t-1}; \theta)$, is the conditional density of observed futures prices. Equation (2.2) and the dynamic Nelson-Siegel specification imply that this conditional density is Gaussian and can be evaluated by means of the Kalman filter[17] given the sequence $\{\Omega_r,\ r \leq t\}$. The time-varying covariance matrix $\Omega_t = \Omega(h_t)$ from the state equation (2.6) is a $\mathcal{F}_{t-1}$-measurable variable as it follows from the assumptions of GARCH dynamics for h_t in (2.10) and (2.17). Therefore, for every t we have that $p\big(y_t|\Omega_t, \mathcal{F}_{t-1}; \theta\big) = p\big(y_t|\mathcal{F}_{t-1}; \theta\big)$ and the Kalman filter can be applied to evaluate the density $p(y_t|\mathcal{F}_{t-1}; \theta)$.

The second term in the factorization (2.22), $p(s_t|y_t, \mathcal{F}_{t-1}; \theta)$, relates to the conditional density of the realized measures s_t. From (2.11) and (2.18) it follows that such density is Gaussian. According to (2.15) and (2.21), the leverage variable $z_t = z_t(y_t, \mathcal{F}_{t-1}; \theta)$ for all t is expressed through the variables from the Kalman forward recursions applied to $p(y_t|\mathcal{F}_{t-1}; \theta)$, so the conditional densities for s_t satisfy $p(s_t|y_t, h_t, z_t, \mathcal{F}_{t-1}; \theta) = p(s_t|y_t, \mathcal{F}_{t-1}; \theta)$.

Finally, the log-likelihood function of the model can be written as

$$\log\ p(y, s|\theta) = \sum_{t=1}^{T} \log\ p\big(y_t|\mathcal{F}_{t-1}; \theta\big) + \sum_{t=1}^{T} \log\ p\big(s_t|y_t, \mathcal{F}_{t-1}; \theta\big) \tag{2.23}$$

The function in (2.23) is effectively a sum of conditional Gaussian log-densities. It can be maximized by means of standard numerical optimization routines and the corresponding quasi-maximum likelihood (QML) parameter estimates, $\hat{\theta}$, can be obtained. While estimating the model we rely on the QML argument in sense of White (1982). In particular, in case of certain distribution misspecifications we refer to (2.23) as to the quasi-likelihood function, so the corresponding QML estimator is supposed to retain its consistency.

[17] See comprehensive material on the state-space models and estimation in Harvey (1991), Durbin and Koopman (2012).

2.2.4 Discussion

Conceptually our modeling framework consists of two blocks. The first block relates to the observed term structure of futures prices. It is described by the dynamic Nelson-Siegel model with time-varying volatility imposed on the factor dynamics. The Nelson-Siegel factors govern the dynamics of the futures price vector y_t and the time-varying volatility embedded in these factors allows to explicitly model the uncertainty associated with y_t. The second block is auxiliary to the dynamic Nelson-Siegel model and links the latent time-varying factor volatility to the realized measures extracted from the high frequency data. This part of the model is unusual in the literature on term structure modeling and we expect that it improves the identification and forecasting the conditional volatility.

The time-varying volatility component h_t is modeled as an observation driven process and fully determines dynamics of the factor conditional variance matrix Ω_t. Important to note that the process h_t incorporates information from both blocks of the framework, i.e. from both low and high frequency data. First, the daily increments of the latent Nelson-Siegel factors Δf_t (which are extracted from the daily prices y_t) can provide (noisy) signals about the latent conditional variance of f_t. Second, the realized measures s_t (based on intra-daily price observations) directly contribute information about the latent factor variance and these signals are supposed to be more precise.

2.3 Data Description

2.3.1 Transaction Data

The raw dataset is obtained from Tick Data Inc. It consists of transactions on light crude oil (CL) futures traded on the Chicago Mercantile Exchange (CME). The data is analyzed using a period ranging from August 9, 2004 to January 16, 2015.

All futures trading are managed by the CME group and through their clearing house they act as a counterpart to all transactions. The transactions are made through a fully electronic trading system that allows market participants to trade around the clock Monday-Friday except for a 45 minutes trading break each day from 5:15p.m.-6.00p.m. EST. A special feature of these futures contracts is that the settlement at maturity is not by cash but via physical delivery. However, the vast majority of contracts never go to delivery. Most are canceled beforehand by taking the offsetting position. This allows for a huge amount of transactions to take place since it is possible to be a market participant without taking part in the physical delivery of the commodity. The availability of large amounts of transactions data facilitates computation of intra-daily measures, which form the basis of our analysis on volatility of oil futures.

Table 2.1 presents statistics for the transaction data. Generally the number of transactions has been increasing through most of the sample period. The average number of transactions per contract shows a similar pattern suggesting that the increase in transactions is not solely due to the introduction of additional contracts. The majority of the transactions are in futures with a short time to maturity while contracts with medium and long time to maturity are less traded.

The raw data series are cleaned using the algorithm in Barndorff-Nielsen et al. (2009). Specifically we delete the observation if the transaction price is more than five mean absolute deviations from a rolling centered median of the 25 preceding and the 25 subsequent observations. On rare occasions observations may also be reported with a time stamp that is not consistent with the corresponding trading day. These observations are also deleted. In total we delete 732,881 out of 153,883,976 transactions corresponding to nearly 0.5% of the initial sample.

2.3.2 Daily Data

The cleaned high-frequent transaction data is used to construct a dataset with daily observations. For all available contracts the closing price, assumed to be the last transaction of the day, is used as the daily futures

Table 2.1: Average number of transactions

	Average number of transactions											
	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	Total
Shortest	4,892	5,130	7,865	23,348	33,002	33,341	35,396	85,002	98,906	102,910	88,186	49,907
Medium	33	78	85	415	844	674	718	8,662	13,399	16,847	13,855	5,406
Longest	0	10	15	167	1,152	321	258	2,747	5,803	6,228	4,752	2,084
All	4,926	5,217	7,964	23,931	34,998	34,336	36,373	96,412	118,108	125,984	106,794	57,398

	Average number of transactions per contract											
	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	Total
Shortest	699	735	1,148	3,343	4,700	4,741	5,048	12,046	13,951	14,586	12,494	7,080
Medium	10	15	23	80	153	112	125	884	1,322	1,666	1,373	560
Longest	NaN	9	13	51	269	97	40	211	411	450	364	200
All	480	395	652	1,548	2,042	2,072	1,902	3,348	3,776	4,063	3,494	2,273

The first part of the table presents the average number of transactions per day for each year in the sample period (rounded to nearest integer). *Shortest* refers to all futures contracts with time to maturity less than 149 business days. *Medium* contracts have between 149 and 360 business days to maturity and *Longest* refers to all contracts with time to maturity larger than 360. The second part shows the average number of transactions per contract per day (rounded to nearest integer). The year 2014 includes the period January 1, 2015 to January 16, 2015.

price. This procedure results in a daily panel dataset consisting of 56,412 observations.

Table 2.2 presents some descriptive statistics for the constructed dataset. The futures prices are all reported as per barrel of oil. The number of traded contracts pr. day generally increases throughout the sample period. The dataset is therefore an unbalanced panel where the cross-sectional dimension varies with time. As evident from the statistics on time to maturity, the increase in the cross-sectional dimension generally comes from the introduction of new contracts with a very long time until expiration. The *Vol.* columns show that the volatility of the futures returns is generally decreasing with time to maturity. This effect is known as the Samuelson effect (Samuelson (1965)). The argument is that the arrival of news will have a direct impact on the shortest futures while long-term contracts tend to remain unchanged since adjustment of fundamentals is likely to take place before the contracts come to delivery at maturity.

Figure 2.1 presents the daily dataset visually. There are two types of plots. In both types each dot represents an observation. The first plot shows the evolution of all the futures prices as a function of time. The next plot shows the evolution of the corresponding time to maturity (measured in business days). The diagonal "lines" in the maturity plot represent specific contracts and arise since the time to maturity decreases linearly with time. The holes in the "lines" represent days where the contract is not traded and are more common for long contracts.

Generally it seems there are a lot of observations in the short end, while observations in the very long end are more scarce and generally not present before 2010. The long contracts are typically also issued with a longer time span between them, reflected by the longer distance between the diagonal "lines" in the maturity plots.

We divided the dataset into three periods. The first 100 days represent an initialization period used to initialize the Kalman filter (August 9, 2004 to December 30, 2004). The second period is the estimation period (January 3, 2005 to December 31, 2013) and the third period is used for forecasting (January 2, 2014 to January 16, 2015). The three periods are separated by vertical dashed lines in figure 2.1.

Figure 2.2 gives an example of what typical term structures may look like. Each dot represents a futures price observation and they are linearly interpolated for a better visualization. All the futures curves are normalized by the price of the shortest futures. It is clear that the futures curves are able to take very different forms. They can be both increasing (contango), decreasing (backwardation), hump-shaped, smooth, and irregular.

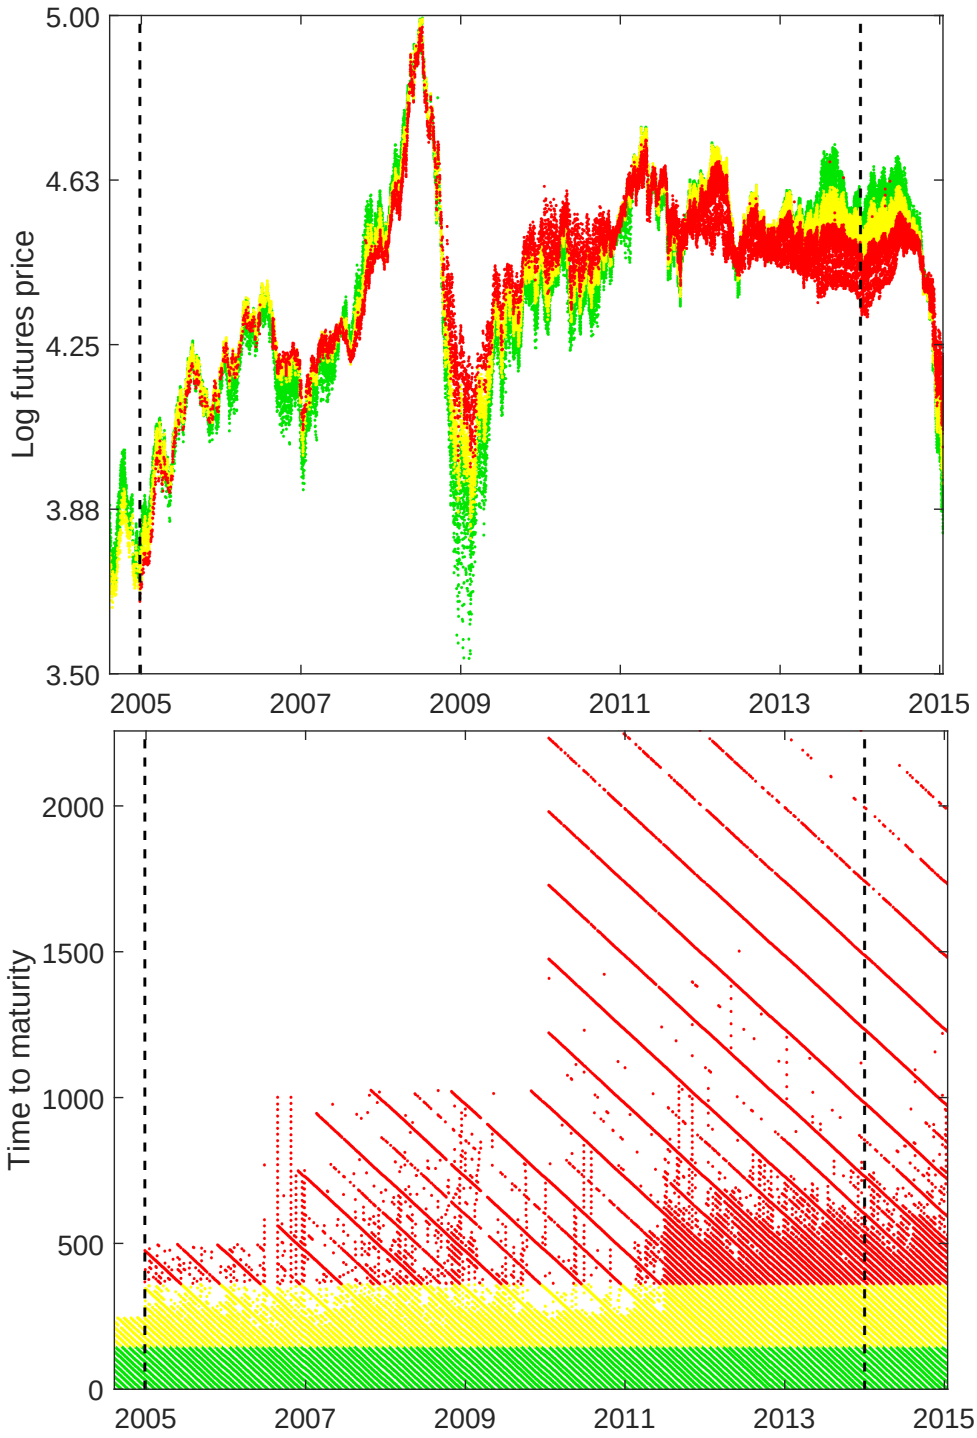

Figure 2.1: The daily dataset. Daily futures prices and time to maturity as functions of time. Each dot represents an observation. For illustration all observations are divided into three groups and colored with respect to time to maturity. Each group consists of approximately one-third of the observations. The first group represents observations with a maturity less than 149 business days and is presented in green while observations with a maturity longer than 360 are colored in red. The second group consists of observations in between and is colored in yellow.

Table 2.2: Descriptive statistics for daily futures data

			Shortest futures				Median futures				Longest futures			
Period	#Cont.	τ	Mean	Min.	Max.	Vol.	Mean	Min.	Max.	Vol.	Mean	Min.	Max.	Vol.
2004	10	108	47.3	40.7	55.2	0.40	43.2	38.6	49.5	0.28	43.2	38.6	49.5	0.27
2005	13	141	56.7	42.2	69.9	0.32	57.5	40.8	69.4	0.24	56.0	39.1	68.2	0.23
2006	12	140	66.3	55.9	77.4	0.27	70.9	64.2	80.8	0.21	70.6	64.5	80.0	0.20
2007	15	168	72.3	50.4	98.3	0.29	72.7	56.1	91.4	0.20	71.4	58.4	88.5	0.18
2008	17	191	99.9	33.2	145.4	0.57	100.8	49.4	146.3	0.40	102.0	64.8	142.2	0.34
2009	16	177	61.9	34.4	81.3	0.51	70.2	48.0	86.0	0.32	77.1	54.3	92.4	0.26
2010	19	205	79.5	69.2	91.4	0.29	84.3	73.9	94.4	0.24	91.5	77.0	100.3	0.18
2011	26	288	95.1	76.4	113.7	0.34	97.7	79.0	114.1	0.28	97.2	87.3	107.1	0.21
2012	31	326	94.2	78.0	109.6	0.25	96.0	81.6	109.2	0.21	87.1	82.0	94.1	0.17
2013	31	325	98.0	86.4	110.2	0.18	93.1	85.5	97.4	0.13	81.5	77.0	86.4	0.11
2014	30	319	91.1	45.7	107.3	0.24	86.4	53.7	98.1	0.18	80.8	66.1	89.8	0.13
Total	19	233	80.4	33.2	145.4	0.35	81.6	38.6	146.3	0.25	80.2	38.6	142.2	0.23

The descriptive statistics are based on the daily data set with sample period August 9, 2004 to January 16, 2015. The three time series *Shortest futures*, *Median futures* and *Longest futures* are constructed by each day taking the shortest contract, the contract with time to maturity closest to 233 (the median number of business days to maturity in the full sample) and the longest contract respectively. *#Cont.* refers to the median number of contracts available per day during the given period. τ shows the median number of business days to maturity for all contracts available. *Vol.* measures the annualized volatility of futures return. The returns are calculated by excluding periods where the shortest contract moves into maturity, a new contract becomes the median contract or a new longest contract is issued.

2.4 Intra-daily Data and Realized Measures

In the preceding section we built a framework for the term structure modeling that introduces time-varying conditional volatility in the dynamics of the Nelson-Siegel factors. The key element in our approach is the use of realized volatility measures constructed with intra-daily futures prices to approximate the latent volatility components in the factor dynamics. In this section, we briefly discuss the features of the high frequency futures data and outline how we use it to retrieve appropriate realized measures of factor volatility.

2.4.1 Intra-daily Data and Practical Challenges

For most business days the futures contracts considered in this paper are traded 24 hours with a short break between 5:15p.m. and 6:00p.m. EST. Thus, we conventionally associate the intra-daily trading period for business day t with the period between 6:00p.m. of $t-1$ and 5:15p.m. of t. Suppose that n_t contracts with times until maturity $\{\tau_{t,1}, \tau_{t,2}, ..., \tau_{t,n_t}\}$ are traded at a given day t. We now assume that there exists n_t-variate process $p_t(v) = \left(p_t^{(\tau_{t,1})}(v), p_t^{(\tau_{t,2})}(v), ..., p_t^{(\tau_{t,n_t})}(v)\right)'$, where $v \in [0,1]$ is an intra-daily time index and $p_t^{(\tau)}(v)$ denotes an intra-daily efficient log-price process of the contract with time to maturity $\tau \in \{\tau_{t,1}, \tau_{t,2}, ..., \tau_{t,n_t}\}$ (contract τ, henceforth). The daily closing price $y_t(\tau)$ which we consider in our main model is, therefore, associated with $p_t^{(\tau)}(1)$.[18]

In practice, an efficient intra-daily price process is never observable. Instead, we observe $p_t(v)$ only at the discrete moments of transactions and these observations are contaminated by microstructural noise. In addition to the market microstructural effects,[19] a multivariate nature of intra-daily futures data suggests several other challenges. In particular, some of these challenges are illustrated in Figure 2.3, which provides examples of typical transaction patterns for the crude oil futures:

- transactions are irregular
- on average, the number of transactions is lower for contracts with higher time to maturity

[18] Practically, however, we use $p_t^{(\tau)}(\nu_{last})$, where ν_{last} is the time of the last observed transaction for the contract τ at day t.

[19] Market microstructural effects are inevitable attributes of financial data observed at high frequency. Among the sources of the microstructural noise researchers emphasize a bid-ask spread, fixed minimal price increment (tick size), etc. Theoretical and empirical implications of market microstructural effects on the observed price process and the analysis of the realized variance are considered in Bandi and Russell (2008) and Hansen and Lunde (2006b).

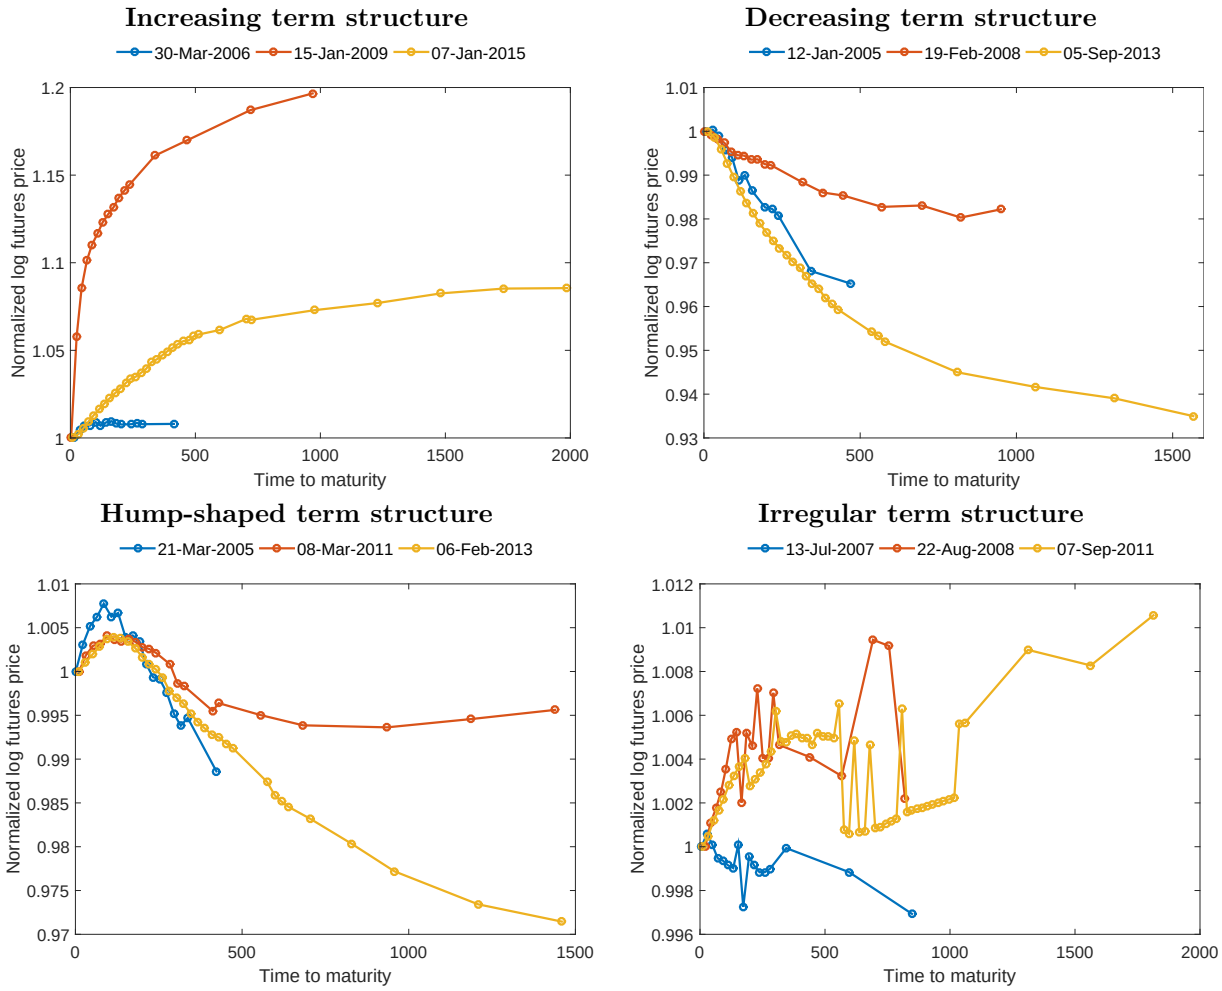

Figure 2.2: Different term structures of futures prices.

- transactions are asynchronous across contracts

Active trading mostly occurs within a relatively short intra-daily sub-period (the distributions of intra-daily trading intensities for the futures contracts are given in Figure 2.4). As a consequence, transactions appear very irregularly over the trading day. In particular, between-trade durations range from milliseconds to several hours depending on the intra-daily period and the time to maturity of the contract.

Another implication of intra-daily futures data relates to the unbalancedness of observed trades across contracts with distinct maturities. Short-maturity contracts are the most traded ones, whereas long-maturity contracts are traded less frequently. This feature is especially pronounced in the first years of our sample. The scarcity of transaction data for long-maturity contracts naturally complicates the extraction of realized volatility signals related to the long end of the term structure.

Transactions on distinct contracts arrive non-synchronously. Partially the asynchronicity can be attributed to the liquidity issues. Namely, we may expect that the fundamental news associated with the underlying commodity are getting incorporated into the futures prices not simultaneously, but with the lags for less liquid long contracts. Another contribution to non-synchronous trading stems from the situations where some information relevant only in the short run does not lead to the trading of long-maturity contracts and conversely. Such situation is typical for multivariate high frequency asset prices and leads to the so-called Epps effect (see Epps (1979)) in the context of covariance estimation.

Irregular, highly unbalanced and asynchronous transaction data substantially complicates the intra-daily analysis of the multivariate factor structure in the futures contracts as well as the construction of the corresponding realized measures of factor volatilities. Therefore, the expected degree of improvements in modeling and forecasting the term structure density from the use of high frequency realized volatility measures should

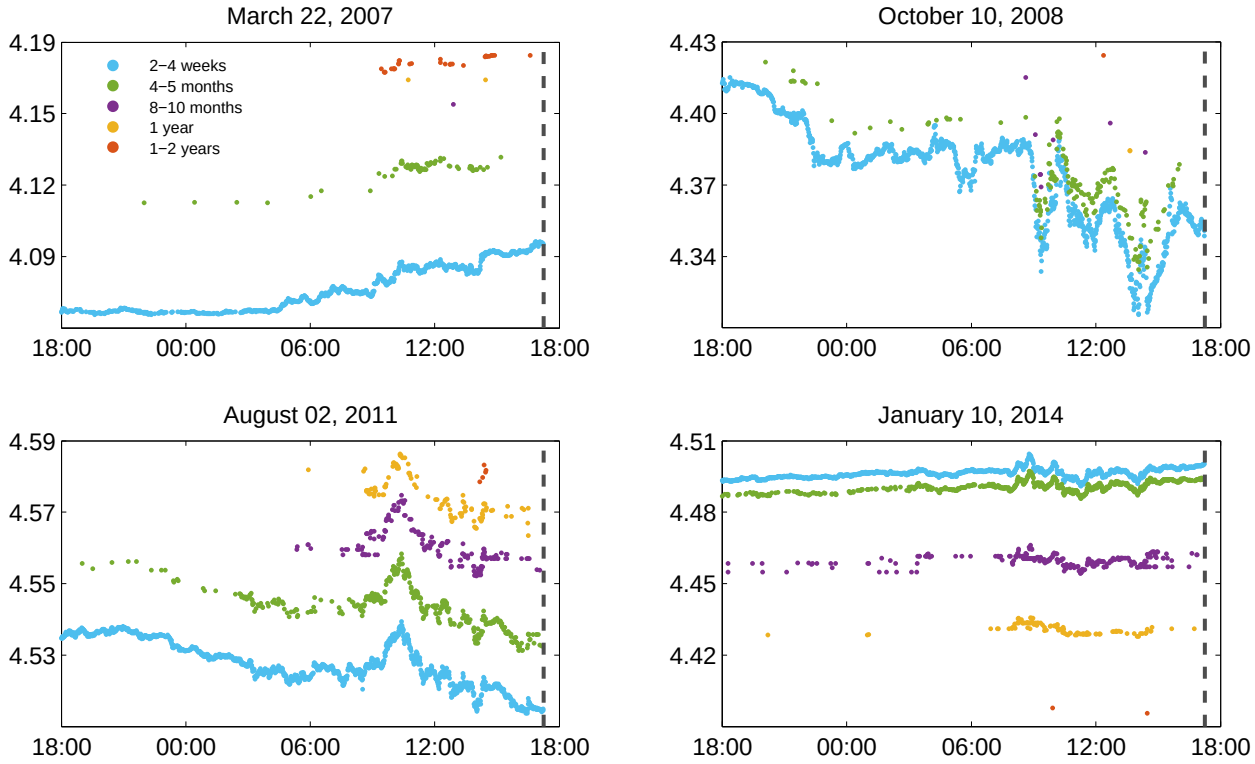

Figure 2.3: Examples of intra-daily transaction data for a 24-hour period which starts at 6p.m. the day before the date specified in the header and lasts until 6p.m. of the specified date. Plots contain log futures prices observed at the moments of transactions for 4 selected trading days from the analyzed sample. For expositional convenience we plot the data for only 5 selected contracts with different times until expiration. Dashed vertical lines correspond to the start of the trading break at 5:15p.m.

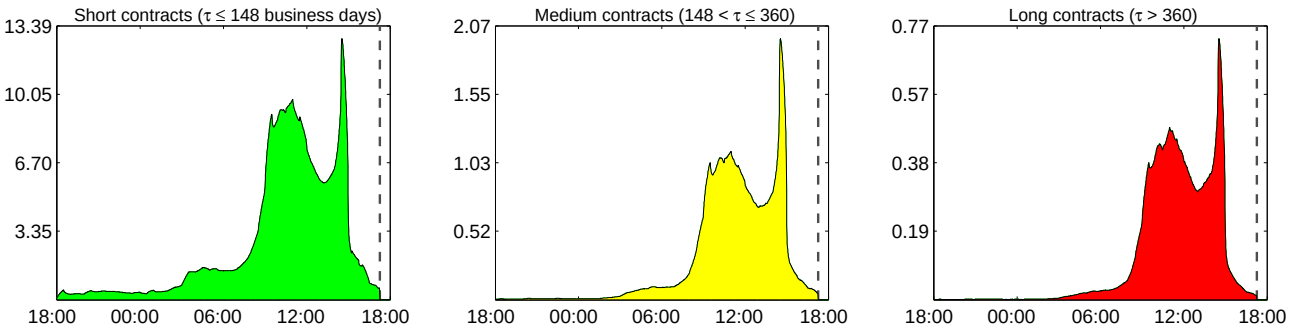

Figure 2.4: Intra-daily distributions of transactions observed over the period 01/2005-12/2014. The left plot corresponds to the futures contracts which expire no later than in 148 business days. The middle plot displays contracts with maturities between 148 and 360 days. The right plot refers to the futures contracts with more than 360 trading days until maturity. On the vertical axis there is a total number of observed transactions (in millions) for the specified contracts from the sample period 01/2005-12/2014 computed over 30-minute rolling intra-daily intervals. Dashed vertical lines correspond to the start of the trading break at 5:15p.m.

be assessed cautiously. In the following subsections we discuss our approach to the construction of realized measures related to the conditional variance of the Nelson-Siegel factors.

2.4.2 Volatility Measures from Intra-daily Data

The aim of our intra-daily developments is to obtain realized measures s_t for the time-varying component h_t from the conditional factor variance matrix $\mathrm{Var}(\Delta f_t | f_{t-1}; \mathcal{F}_{t-1}) = \Omega_t$. As it is outlined in the model description, we consider two specifications of Ω_t based on the one- and three-component dynamic parametrizations. We now suggest and discuss suitable realized measures for each specification. An empirical justification of using few dynamic volatility components for modeling Ω_t is given in Appendix B.1.

2.4.2.1 Realized Measures for the One-component Specification

According to (2.14) the realized variance of a close-to-close futures return on some contract $\tau \in \{\tau_{t,1}, \tau_{t,2}, ..., \tau_{t,n_t}\}$ can serve as a signal of h_t, since all contracts $y_t(\tau_t)$ contain information about the common latent volatility component h_t. The unbalancedness of trading intensity across contracts with different times until expiration suggests that the front contract with the shortest time to maturity is the best candidate for measuring h_t. Such contract is systematically traded more frequently than others and, what is more important, the corresponding transactions appear in the course of the whole trading day. Therefore, the realized variance computed on the basis of intra-daily transactions on the shortest contract $y_t(\tau_{t,1})$ is expected to be a more efficient measure of h_t rather than a realized variance computed using some other contract with longer time to maturity. The realized variance of the shortest contract, thus, can be treated as a state variable which controls all three Nelson-Siegel factor volatilities.

To minimize the influence of microstructural noise effects we suggest to use a sparse sub-sampled 5-minute realized variance estimator.[20] For this purpose we partition a trading period $v \in [0, 1]$ with a 1-minute grid. We denote the set of equidistant gridpoint times as $\{v_i\}_{i=0}^m$, so $0 = v_0 < v_1 < ... < v_m = 1$. Although futures contracts with the shortest maturity $\tau_{t,1}$ are liquid enough and traded frequently, they do not necessarily occur exactly at the moments v_i for all i. Thus, we refer to the previous tick interpolating scheme. In particular, let $\tilde{p}(v_i)$ denote the transaction price in case we observe it at v_i. If we do not observe a transaction at that moment, then $\tilde{p}(v_i)$ stands for the price observed for the nearest transaction prior to v_i.[21] Then, the sub-sampled 5-minute realized variance estimator based on intra-daily prices for contract τ reads

$$RV_t^{ss}(\tau) = \frac{1}{5}\sum_{j=1}^{5} \frac{m}{m-4} \sum_{i=j}^{m-5+j} \left(\tilde{p}_t^{(\tau)}(v_{5i}) - \tilde{p}_t^{(\tau)}(v_{5(i-1)})\right)^2 \tag{2.24}$$

and we define a realized measure for the one-component dynamic variance specification as $s_t \equiv \log RV_t^{ss}(\tau_{t,1})$.

The sub-sampled estimator in (2.24) is essentially an average of realized variances computed for 5 sparse 5-minute grids shifted by a 1-minute interval with respect to each other. Such sub-sampling of the realized variance was proposed in Zhang et al. (2005) and later considered in Andersen et al. (2011b), Liu et al. (2015), among others[22]. The sub-sampling is aimed to improve efficiency of sparsely sampled estimators since more intra-daily information is used in this case.

The estimator in (2.24) is consistent for the quadratic variation of an asset return for the open-to-close trading period. Note, that despite s_t is linked in (2.11) with the close-to-close return variance, we do not account explicitly for the "overnight" (and "overweekend") variation while specifying s_t. Firstly, the length of the overnight period for the considered commodities is very short (45 minutes). Secondly, the specification of

[20] We prefer model-free realized variance estimator since the almost 24-hours trading span provides sufficiently many sparsely sampled intra-daily returns. Also, Liu et al. (2015) demonstrated in an extensive empirical study that a sparsely sampled realized variance estimator exhibits an accuracy which is not significantly outperformed by any of the competing realized estimators. As an alternative, noise robust estimators of the quadratic variation based on more frequently sampled intra-daily returns can be exploited, such as in Barndorff-Nielsen et al. (2008), Jacod et al. (2009).

[21] Previous tick interpolation was introduced in Wasserfallen and Zimmermann (1985). See also a related discussion in Hansen and Lunde (2006b).

[22] Note that estimator (2.24) can also be considered as a case of the realized kernel estimator where Bartlett kernel is used.

the measurement equation (2.11) is able to capture a systematic bias in s_t stemming from the lack of overnight variation embedded in the realized measures.[23]

2.4.2.2 Realized Measures for the Three-component Specification

When we formulate the realized measure for the one-component specification of Ω_t, we discard much of the available intra-daily data. In particular, we use the high frequency price observations of the contract with the shortest maturity, whereas the intra-daily data from all of the traded contracts is available.

According to the three-component dynamic specification of Ω_t, the time-varying elements are the conditional volatilities of the Nelson-Siegel factors, i.e. the components of h_t of the diagonal matrix H_t from (2.16). We now suggest a way to retrieve a proxy for h_t from the intra-daily data on multiple traded contracts with different expiration times τ. Namely, we directly assess factor volatilities by evaluating parametrically the dynamic Nelson-Siegel models intra-daily.

Estimation of the multivariate state-space models with high frequency data closely relates to the recent developments of Shephard and Xiu (2012) and Corsi et al. (2015) in the parametric QML estimation of realized covariances in equity prices. In particular, these methods consider an intra-daily multivariate price process as a local level model and provide consistent QML estimates of the covariance matrix with irregular asynchronous price observations under weak assumptions on the underlying price process. In our case we treat an intra-daily multivariate price process $p_t(v)$ as a factor (dynamic Nelson-Siegel) model. As a consequence, our intra-daily state-space model has an additional challenge of latent factor identification. A typical pattern of the observed high frequency futures prices provides a very harsh field for an intra-daily implementation of such state-space methods. Thus, we reorganize slightly the structure of high frequency observations and impose several simplified assumptions on the intra-daily Nelson-Siegel framework.

First of all, similar to case with the one-component measure, we partition a trading day with a 5-minute grid and define the vector of observed prices as $\tilde{p}_t(v_i) = \big(\tilde{p}_t^{(\tau_{t,1})}(v_i), \tilde{p}_t^{(\tau_{t,2})}(v_i), ..., \tilde{p}_t^{(\tau_{t,n_t})}(v_i)\big)'$ for $i = 1, ..., m$ using the previous tick interpolation scheme.[24] As a result, we obtain a panel of observed prices $\{\tilde{p}_t(v_i)\}_{i=0}^m$. The fact that we use a sparse 5-minute grid is helpful for mitigating microstructural noise effects and the improved balancedness of an intra-daily observations panel enhances a stability of the model identification and estimation.

We then assume that the intra-daily price dynamics of $\tilde{p}_t(v_i)$ is compatible with the constant volatility dynamic Nelson-Siegel model. Thus, the intra-daily dynamic Nelson-Siegel model for a trading day t is given by

$$\tilde{p}_t(v_i) = \Lambda_t g_t(v_i) + \epsilon_t(v_i) \qquad \epsilon_t(v_i) \sim iidN\big(0, \sigma_\epsilon^2 I_{n_t}\big) \tag{2.25}$$

$$g_t(v_i) = g_t(v_{i-1}) + \omega_t(v_i) \qquad \omega_t(v_i) \sim iidN\left(0, \frac{1}{m} H_t R H_t\right) \tag{2.26}$$

where $g_t(v_i)$ is a 3×1 vector of the latent intra-daily Nelson-Siegel factors. For any given trading day t the parameters are fixed and do not depend on the intra-daily index i. Note, that in the intra-daily model specification we do not endow the idiosyncratic term $\epsilon_t(v_i)$ with the autocorrelation effect as we do it in the daily model. The model in (2.25)-(2.26) is a discrete time linear Gaussian state-space model that can be estimated by means of the standard Kalman filter. We also note, that the discrete state-space specification (2.25)-(2.26) complies with the Brownian semimartingale assumption for the latent efficient futures prices $p_t(v)$ which is a standard assumption in the high frequency econometrics literature.

We facilitate the estimation of (2.25)-(2.26) by presetting parameter values for λ, σ_ϵ^2 and R with the corresponding estimates from the constant volatility dynamic Nelson-Siegel model evaluated at a daily frequency. Therefore, only the intra-daily factor covariance matrix $H_t = H(h_t)$ is left for estimation in (2.25)-(2.26). We

[23] Alternatively, it is possible to combine the open-to-close realized variance and the overnight squared return to obtain a consistent measure for a close-to-close quadratic variation, as in Hansen and Lunde (2005). This approach was used in Christoffersen et al. (2014) in a similar context to measure the variation in commodity futures prices.

[24] In addition, if some contract τ that is traded at t was traded at $t-1$ with time to maturity $\tau+1$, we use the close price of the contract $\tau+1$ at $t-1$ in the interpolation scheme when we obtain $\tilde{p}_t^{(\tau)}(v_i)$ for day t.

then define the required realized measure as $s_t \equiv \hat{h}_t$. Vector $\hat{h}_t$ here is the corresponding factor variance estimate obtained from the intra-daily dynamic Nelson-Siegel models (2.25)-(2.26), which are evaluated using the high frequency price observations $\{\tilde{p}_t(v_i)\}_{i=0}^m$ of contracts traded at day t.

Of course, the realized measure s_t is hardly a consistent measure of the conditional variances h_t of the latent Nelson-Siegel factors from the daily model at least due to the difference in specifications between the daily and the simplified intra-daily models (2.25)-(2.26). What is important, we suppose that these measures correlate well with the unobservable factor volatility processes. Thus, they still can be considered as valid and the possible systematic bias can be captured by the parameters of the measurement equations (2.18) in the main model. Informally, we will refer to these measures of factor variance as realized measures.

2.4.3 Remarks

Table (2.3) contains descriptive statistics of the estimated realized measures for the one- and three-component model specifications. Figure (2.5) provides the corresponding time series plots.

								ACF		
Realized measure	Mean	Std.	Skew.	Kurt.	Min	Med	Max	$\rho(1)$	$\rho(5)$	$\rho(22)$
Shortest contract	0.306	0.163	2.823	16.006	0.048	0.275	2.046	0.857	0.803	0.713
Level factor	0.113	0.078	1.590	6.609	0.007	0.095	0.553	0.556	0.437	0.379
Slope factor	0.184	0.102	1.848	7.858	0.021	0.163	0.784	0.646	0.529	0.444
Curvature factor	0.371	0.272	1.890	8.199	0.031	0.308	2.004	0.631	0.553	0.494

Table 2.3: Descriptive statistics of the realized measures of volatility (annualized) for the sample period that starts in January, 2005 and lasts until December, 2014.

The realized measures for the one-component specification based on the non-parametric realized variance estimator demonstrate relatively smooth and highly persistent dynamics. On contrary, parametric volatility measures for the three-component specification reveal noisy and weakly persistent dynamics, which suggests that these measures are likely less accurate.

There are several potential reasons of why the three-component measures are relatively noisy. First of all, since high frequency price observations are unbalanced and we usually do not observe a sufficient number of transactions on long-term contracts, the factor identification in the intra-daily Nelson-Siegel models is weak and unstable.[25] We also have to recognize that the data interpolation may lead to the distortions in the intra-daily estimation of $\hat{h}_t$. We argue, nonetheless, that an adverse impact from these causes is expected to be limited since we preset factor correlation matrix R and only estimate volatilities h_t of the latent factors.

This issue points out an interesting trade-off between the use of the one- and three-component dynamic specifications of the time-varying covariance matrix Ω_t. On one hand, the three-component dynamics provides more flexibility to the term structure volatility, on the other hand the one-component specification suggests a better precision of the realized measures. Consequently, this trade-off makes it ex-ante unclear which type of approach is able to better model and predict the density of the futures prices and we relegate this question to our empirical evaluation.

[25] In practical implementation we use the following method to improve the factor identification. When estimating an intra-daily Nelson-Siegel model for a trading day t we supply it with the interpolated 5-minute prices from 2 consecutive trading days $t-1$ and t, $\{\tilde{p}_{t-1}(v_i); \tilde{p}(v_i)\}_{i=0}^m$. Although only the data related to day t is used for likelihood maximization, additional data from the previous trading day significantly improves the ability of the model to identify the latent Nelson-Siegel factors. This happens due to the critical deficit of price observations for the contracts with long maturities if the data from just one trading day is used. In case we obtain a "suspicious" outlying volatility estimate, we use the data from the two consecutive trading days for likelihood maximization. It helps to smooth out very noisy estimates coming from the days when the intra-daily observation panel is especially scarse.

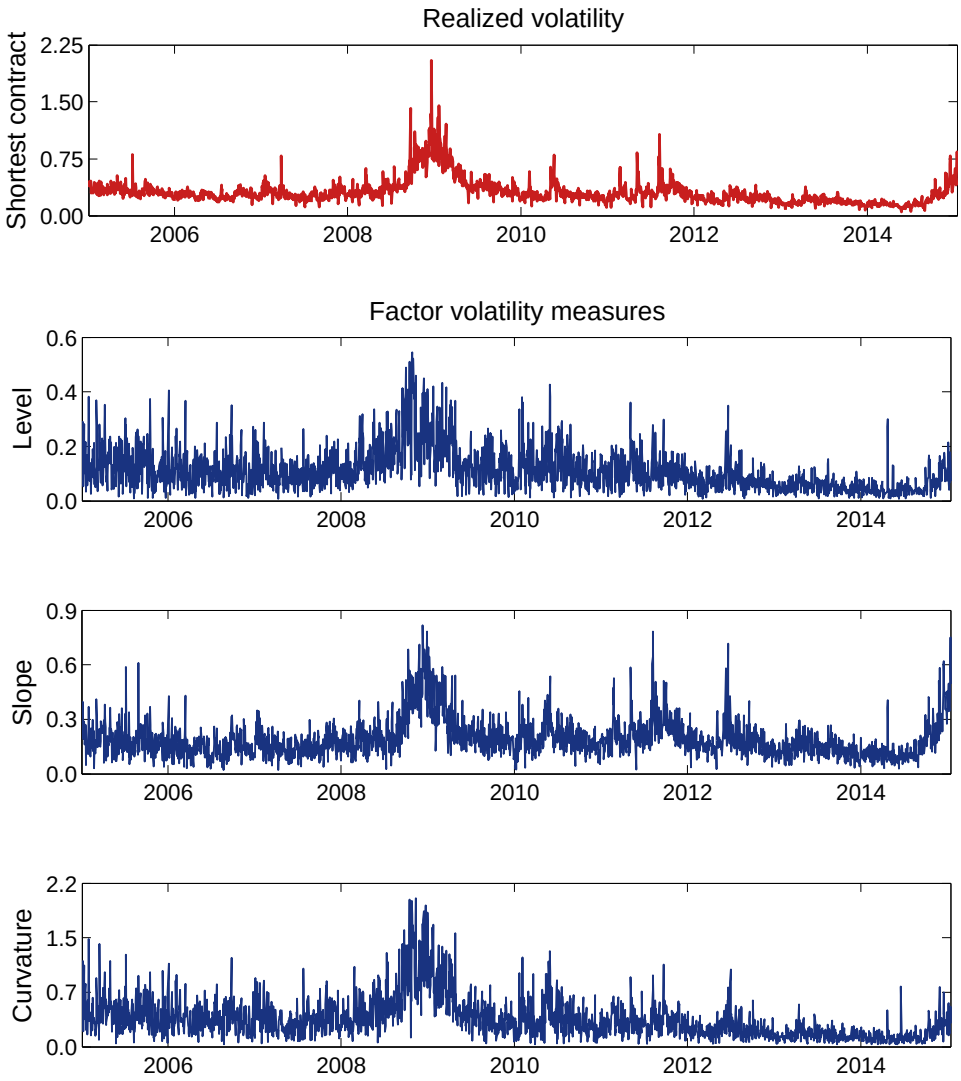

Figure 2.5: In the top panel: annualized subsampled 5-min realized volatility for the front futures contracts. In the 3 bottom panels: annualized volatility measures of the Nelson-Siegel factors estimated parametrically using intra-daily state-space models for each trading day of the sample. The sample period starts in January, 2005 and lasts until December, 2014.

2.5 Evaluation Strategy

In this section, we outline the out-of-sample evaluation of our modeling framework. At first, we briefly describe the set of alternative (benchmark) model specifications. The corresponding state-space model formulations are given in Appendix B.2. Then we explain our strategy for the out-of-sample density forecast evaluation.

2.5.1 Pool of Competing Models

In our empirical evaluation we consider several alternative models along with our baseline framework. The goal is to investigate whether time-varying volatility and realized measures suggested in our main framework provide improvements in the out-of-sample performance compared to the alternative models with no such features. For better comparison, all the alternative models have the DNS factor dynamics in the conditional mean, have autocorrelated idiosyncratic components, and the random walk assumption is maintained for the Nelson-Siegel factor dynamics. In other words, the part described by equations (2.2)-(2.6) is common for all considered models and the differences come in the specification of the factor covariance matrix Ω_t. Particularly, we examine alternative specifications *a)* with constant volatility and *b)* with time-varying volatility, but with

Model	Dynamic covariance components	Realized measures	Leverage effect
DNS-CV	0	No	No
DNS-GARCH-1	1	No	No
DNS-GARCH-3	3	No	No
DNS-RGARCH-1	1	Yes	No
DNS-RGARCH-1-lev	1	Yes	Yes
DNS-RGARCH-3	3	Yes	No
DNS-RGARCH-3-lev	3	Yes	Yes

Table 2.4: Set of competing models used in the paper. All models are the DNS factor models with the random walk dynamics of the Nelson-Siegel factors and the autocorrelated idiosyncratic components. In all models the correlations between the Nelson-Siegel factor innovations are constant.

no use of realized measures.

DNS with constant variance (DNS-CV). This model is the standard state-space DNS framework from Diebold et al. (2006) where all variance parameters are assumed to be constant. The model is characterized by (2.2)-(2.6) with the constant covariance matrix Ω in (2.6). The DNS-CV represents a natural constant volatility benchmark in our analysis.

DNS with GARCH effects (DNS-GARCH). In this class of models we add time-varying volatility components to the factor covariance matrix Ω_t. We suggest two specifications with one and three dynamic volatility components, so they possess the same parametrization of Ω_t as proposed in our baseline framework. Similarly, these time-varying components are supposed to have GARCH-type dynamics, but they are not related to any kind of realized volatility measures. We follow the methodology from Harvey et al. (1992) by incorporating GARCH effects into a linear Gaussian state-space model. Therefore, instead of realized measures we use the conditional expectations of the factor innovation shocks in order to specify the updating terms in the GARCH equations for dynamic volatilities.[26] Since these conditional expectations appear explicitly in the forward Kalman recursions, the models can be estimated directly by means of the Kalman filter. We consider them in order to assess the gains from introducing time-varying volatility in the Nelson-Siegel factor dynamics. At the same time, we rule out the effect from the realized measures which is presented in our baseline models.

DNS with GARCH effects and realized measures (DNS-RGARCH). This class of models is our baseline framework and is described in details in Section 2. We consider specifications of Ω_t with the one and three dynamic volatility components exploiting the information from realized measures of volatility. Therefore, in addition to the time-varying volatility aspect, these models also incorporate the effect from the use of volatility measures based on high frequency data. We examine the baseline models with and without leverage effects specified in (2.12) and (2.19).

Table 2.4 provides the complete list of the models used in our evaluation study.

2.5.2 Density Forecasting

In contrast to point and interval forecasts, which only reflect a central tendency and a probability of falling into a certain interval respectively, density forecasts provide a full probability distribution over the possible future realizations of the forecasted variable. Thus, density forecasting allows to evaluate the model's ability to describe a complete pattern of uncertainty associated with the prediction. Essentially, a specification of the conditional second moment in a model is supposed to affect substantially the quality of the corresponding density forecasts since it is directly related to the degree of uncertainty in the dynamics of the predicted variable. Thus, density forecasting is especially suitable for model evaluation in our context since we aim to investigate and compare the gains from the several dynamic volatility specifications.

[26]The use of conditional expectations of factor volatility shocks implies that the model should be seen as an approximation to a standard GARCH model.

We use the log predictive likelihood as a criterion for the forecast evaluation. The comparison of scoring rules across the pool of competing models to assess their relative ability of producing density forecasts has a long tradition in empirical macro and finance (see e.g. Diebold and Lopez (1996), Lopez (2001), Geweke and Amisano (2010)). The log predictive likelihood of a log futures price $y_{t+h}(\tau)$ reads

$$LPL^{m}_{t,h,\tau}\big(y_{t+h}(\tau)\big) = \log \hat{p}\big(y_{t+h}(\tau)|\mathcal{F}_t, \mathcal{M}_m\big) \tag{2.27}$$

where h denotes a forecast horizon and $\hat{p}\big(\cdot|\mathcal{F}_t, \mathcal{M}_m\big)$ is the forecast of the probability density function of $y_{t+h}(\tau)$ that is made with the model $\mathcal{M}_m$ conditional on the information $\mathcal{F}_t$ available up to day t. As it is evident from (2.27), the log predictive likelihood gets a higher value if the corresponding forecasted density assigns a higher probability to the actual future realization of the predicted variable ($y_{t+h}(\tau)$, in our case). Intuitively, the model that systematically exhibits higher predictive likelihoods has a better density forecasting performance.

Given the fact that our models are set at a daily frequency, we focus on short-term forecasting. We consider 8 forecast horizons: 1, 2, 3, 4, 5, 10, 15, and 22 business days ahead. Note that for commodity futures the maturities of traded contracts τ change from day to day, so we can not observe the same set of contracts in the course of the forecasting period. Thus, we partition the traded contracts into several maturity ranges (e.g. 1-3 months, 1-2 years, etc.) and consider average log predictive likelihoods across the traded contracts from the given maturity range.

The forecasting period lasts from January 2, 2014 until January 16, 2015 and contains 269 trading days in total. The oil prices demonstrate very diverse dynamic regimes throughout the forecasting period (see Figure 2.6). In particular, it spans both the relatively calm period in the first half of 2014 and the sharp downfall of oil prices that occurred during the second half of 2014. In addition to the full out-of-sample period we also evaluate our models based on the 6 month sub-period that lasts from March 3, 2014 until August 29, 2014 (129 trading days). During this sub-period the futures prices behave in a relatively regular way. We consider this sub-period separately for two reasons. First, since the full out-of-sample period contains the episode with the oil price collapse in autumn 2014, the sharp negative trend in prices might cause severe density mispredictions for a number of days. As a consequence, we get too large negative values of log predictive likelihoods for these days and this complicates the comparative analysis.[27] Thus, we may expect that the model comparison based on the calm forecasting period will be less noisy. Second, it is interesting by itself to inspect the relative forecasting performance of our models based on the out-of-sample sub-period when the futures prices have more stable and regular dynamics.

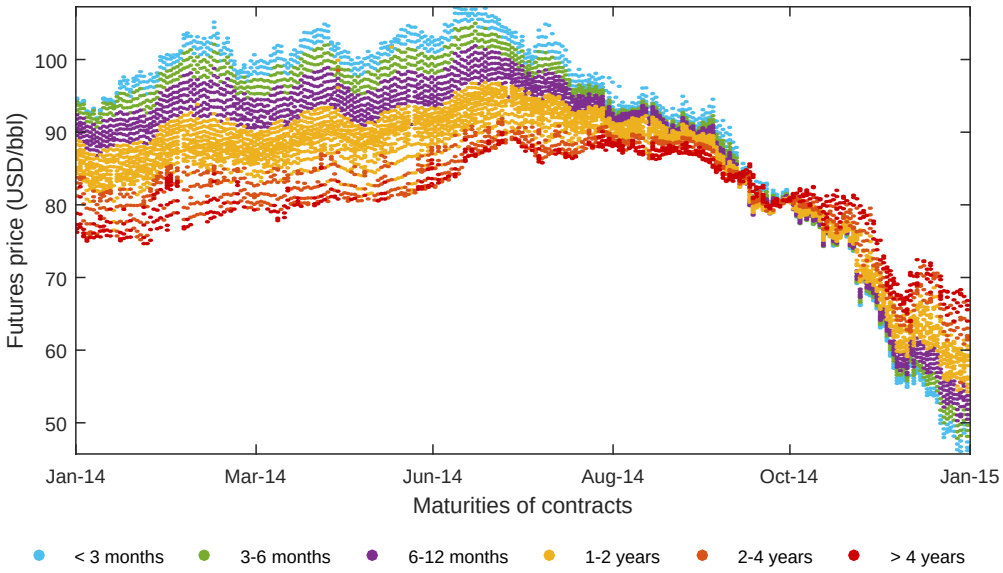

Figure 2.6: Futures prices during the forecasting period (January 2, 2014 - January 16, 2015)

[27] In the forecast evaluation based on the full out-of-sample period we delete such days where the log predictive likelihoods get extremely large negative values for all of considered models (14 trading days in total).

In order to forecast the futures price density from a given model we refer to the simulation techniques (apart from the case of the constant volatility DNS model where the predicted density can be derived analytically). More precisely, in order to produce a density forecast $\hat{p}_{t+h,\tau}\big(\cdot|\mathcal{F}_t, \mathcal{M}_m\big)$ for the price $y_{t+h}(\tau)$, we simulate S price sequences of length h, which are initialized at day t given information $\mathcal{F}_t$. While simulating, we use the stochastic law of motion defined by the model $\mathcal{M}_m$ with the pre-estimated constant parameters. As we obtain S forecasts $\{\hat{y}^s_{t+h}(\tau)\}_{s=1}^S$, we apply a kernel density estimator to get $\hat{p}_{t+h,\tau}\big(\cdot|\mathcal{F}_t, \mathcal{M}_m\big)$ and evaluate the log predictive likelihood from (2.27) at the actual futures price realization $y_{t+h}(\tau)$. In our analysis we take $S = 10,000$ simulations and use the MATLAB built-in function *ksdensity* with Normal kernel smoother and the corresponding optimal bandwidth.

In order to statistically assess the differences across the log predictive likelihoods from the pool of competing models, we construct the Model Confidence Set (MCS) of Hansen et al. (2011). Compared to the Weighted Likelihood Ratio Test in Amisano and Giacomini (2007), which is designed for a pairwise model comparison and is a popular tool for the statistical evaluation of density forecasts in the literature, MCS provides the framework for *multiple* comparison of the competing models. Taking the log-predictive density as a "gain" criterion, we define the MCS as follows

$$MCS_{h,\tau} = \big\{m^* \in \mathcal{M} :\ E(LPL^{m^*}_{t,h,\tau}) \geq E(LPL^{m}_{t,h,\tau}),\ \forall m \in \mathcal{M}\big\} \tag{2.28}$$

so MCS consists of the models that exhibit significantly superior log-predictive likelihoods. We use Kevin Sheppard's MFE Toolbox[28] for MATLAB in order to obtain the corresponding p-values for MCS of 99, 95, and 90% coverage, while doing so we set the number of stationary bootstrap replications at 10,000 and the average window length at 12.

2.6 Empirical Results

2.6.1 Parameter Estimates and Factors

Table 2.5 presents the parameter estimates for the seven different models. Significance at a 5% level is indicated in bold face.

The parameter estimates are obtained by numerically maximizing the log-likelihood (2.23) using data from the estimation period (January 3, 2005 to December 31, 2013). To start the estimation with appropriate values of the three Nelson-Siegel factors and the volatility components we make use of the 100 days initialization period (August 9, 2004 to December 30, 2004).[29]

The parameters $\lambda, \sigma^2_w, \beta, \rho_{ls}, \rho_{lc}, \rho_{sc}$ are common for all models and their estimates tend to be relatively close across the models. Hansen and Lunde (2013) fix $\lambda = 0.005$, which is slightly smaller than the estimates found in Table 2.5. The parameter β is significant and relatively high for all models, which reflects the importance of incorporating autocorrelation in the idiosyncratic component especially when applying the model to daily data. The most pronounced correlation parameter is the large negative correlation between the level and the curvature. Hence, an increase in the overall level of futures prices tends to be offset by a decrease in the prices of the medium maturity contracts.

The parameter estimates controlling the volatility dynamics in the standard GARCH models are not directly comparable to their realized counterparts. The volatility dynamics in the GARCH-1 and GARCH-3 models are linear while the models based on realized measures have a log-linear volatility specification.

If we consider the realized models, we immediately see that the realized measure of the RGARCH-1(-lev) model is less noisy than the realized measures of the RGARCH-3(-lev) model by comparing the σ^2_u parameter estimates. The parameter γ_{1l} governing the impact on the volatility from the realized measure is also substan-

[28] https://www.kevinsheppard.com/MFE_Toolbox

[29] In the initialization period the filter is started by assuming the factors are multivariate normal with a mean equal to their empirical counterparts (see Diebold and Li (2006)) and a high variance equal to the identity matrix. The volatility components are initialized at their unconditional means.

tially larger in the RGARCH-1(-lev) model compared to the corresponding parameters in the RGARCH-3(-lev) model.

In the RGARCH-1-lev model we find a negative δ_1, which is the common finding in finance literature. Here, a negative return in the shortest futures often leads to an increase in the volatility of all three factors. In the RGARCH-3-lev model the sign of the δ_1 parameters is not unambiguous and for the level factor we observe an inverse leverage effect $\delta_{1l} > 0$. All δ_1 parameter estimates are, however, insignificant in this model. The δ_2 parameter is positive in the RGARCH-1-lev, which indicates that large standardized returns of either sign in the shortest futures often lead to an increase in the general volatility level. For RGARCH-3-lev we find an even more pronounced effect for all factors.

Figure 2.7 shows the filtered factors and the annualized volatility series of the different models. Here we only present the filtered factors for the RGARCH-3 model including its 99% confidence interval. The filtered factors from the other models are visually identical. In the volatility plots we leave out the volatility series of two models including leverage since they are visually identical to their non-leverage counterparts.

By comparing the level factor with the data (Figure 2.1) we may see that the factor tracks the overall tendency of the futures prices. The slope factor is slightly positive (corresponding to a decreasing futures curve) in the beginning of the estimation period and rapidly turns negative to reach its lowest values during the financial crisis. After the financial crisis the slope factor gradually returns to positive values. The curvature is positive in nearly 75% of the days in the estimation period. A hump-shaped or inverted hump-shaped futures curve may appear if the curvature is relatively large compared to the slope. Most of the hump-shaped futures curves are observed in late 2005 to 2006 and again in 2012. A few inverted hump-shaped futures curves, with a negative curvature and a relatively small slope, are observed in late 2007.

The time-varying volatility series of the three factors all share the same pattern. All factors are most volatile during the financial crisis. The curvature is by far the most volatile factor followed by the slope. The level is generally the factor with the smallest volatility. It is evident that the volatility series of the RGARCH-3 models are less "rough" than the series of the competing models. This is consistent with the relatively low estimates of the γ_1 parameters.

2.6.2 Density Forecasting

Tables 2.6-2.9 contain the results of our density forecasting analysis. Results based on the whole forecasting period are reported in Tables 2.6-2.7. In Tables 2.8-2.9 we provide the results based on the 6 month sub-period (March, 2014 - August, 2014) where the oil prices exhibit relatively stable dynamics with no long negative or positive trends (see Figure 2.6). We report the mean predictive log-likelihoods in the first column that are related to the constant volatility DNS model. For the remaining models we provide the changes in mean predictive log-likelihoods relative to the constant volatility benchmark from the first column. Since we compare log-likelihoods, a difference in 0.01 approximately stands for a 1% relative change in the density forecasting performance. The background colors indicate that the corresponding models belong to MCS. It means that for a given forecast horizon and a maturity range such models produce higher predictive log-likelihoods and the gains are statistically significant.

The analysis of the density forecasting results reveals several important implications. First of all, it is evident that models with time-varying volatility strongly outperform the constant volatility benchmark.[30] This finding is quite expected in the density predictability context and agrees with the corresponding results on the term structure density forecasting of interest rates (see Carriero et al. (2014), Shin and Zhong (2015)). As far as the forecasting horizon increases, the difference between constant and time-varying volatility models disappears. This fact is also not surprising since in the absence of long-memory effects the GARCH volatility components converge to their unconditional means relatively fast. We also note that such difference is more pronounced for

[30]We note that the constant volatility DNS model produces pretty conservative density forecasts and, therefore, is not an easy competitor. This is because the estimation sample includes the exceedingly volatile Great Recession episode, so it increases the estimated values of constant variance parameters.

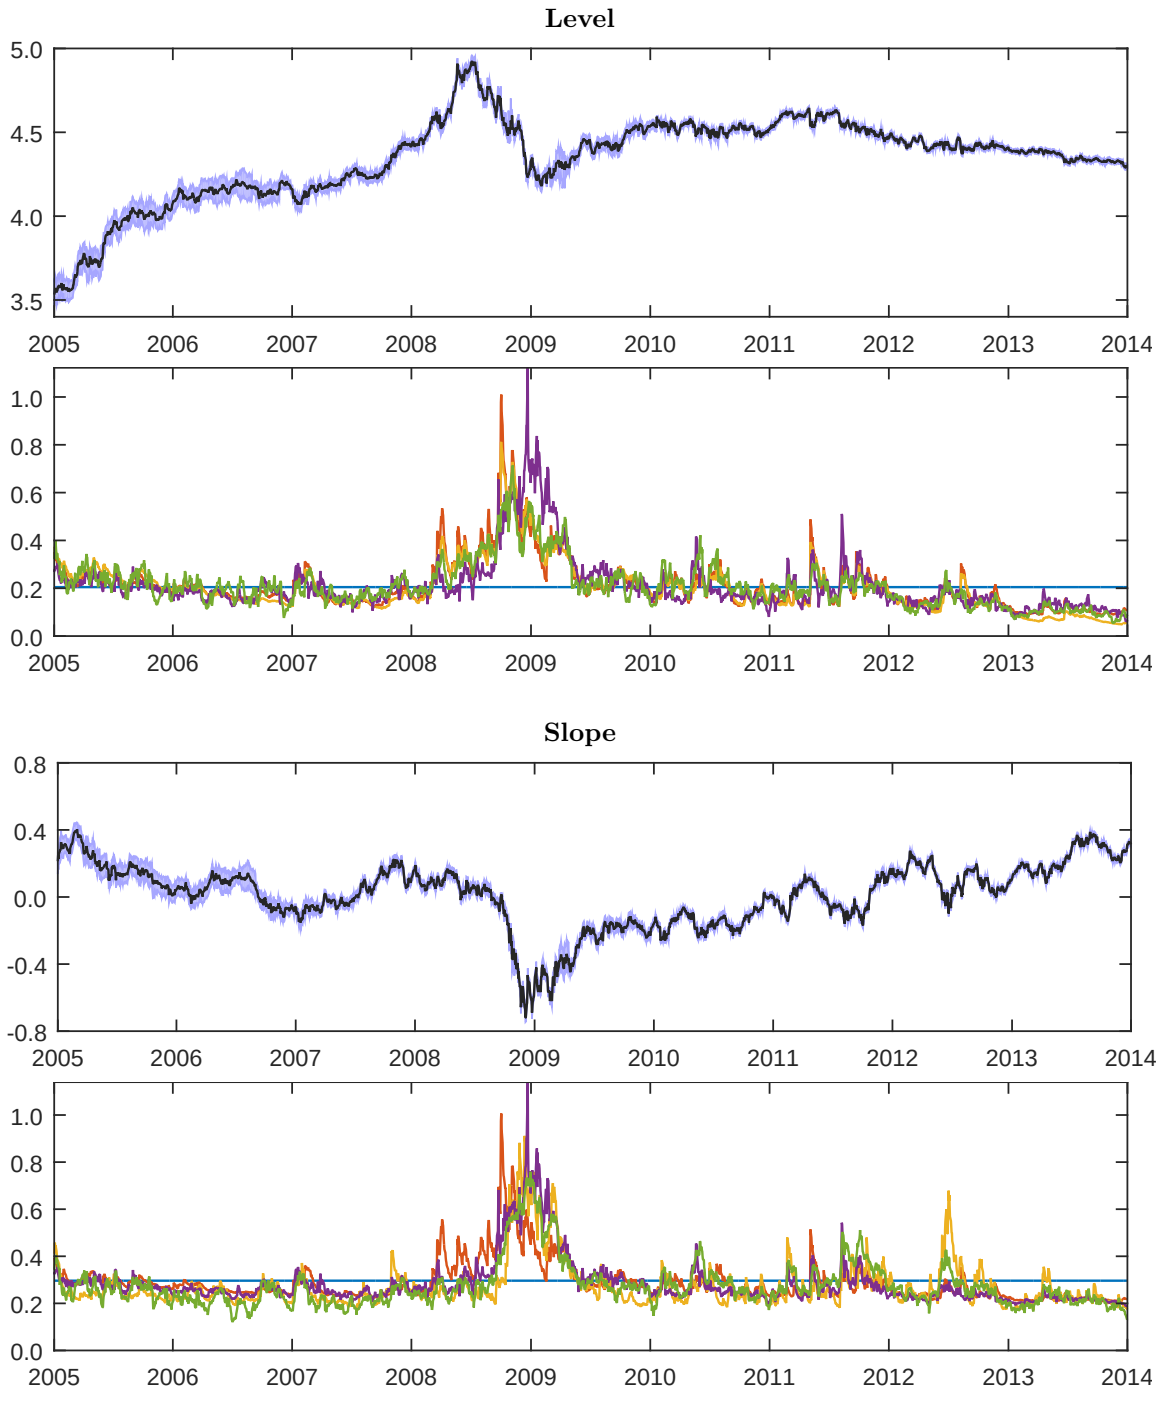
Level
5.0
4.5
4.0
3.5
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
1.0
0.8
0.6
0.4
0.2
0.0
Slope
0.8
0.4
0.0
-0.4
-0.8

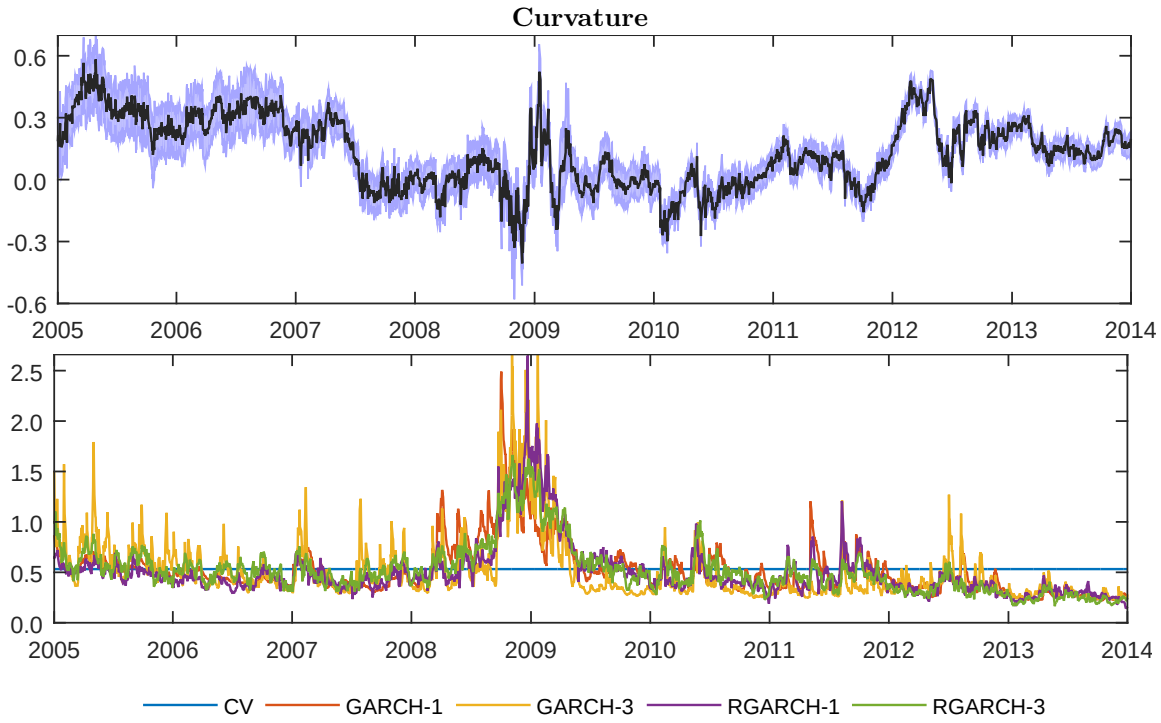

Figure 2.7: Filtered factors with 99% confidence intervals and the annualized filtered conditional factor volatilities.

the futures with short and medium time to maturity, whereas for the long contracts (more than 4 years until expiration) an advantage of time-varying volatility is significant only up to the 2-week forecasting horizon.

Secondly, we find that the use of realized volatility measures provides not large, but systematic gains in density forecasting. Since the volatility is persistent, realized measures make an appreciable difference only when the latent volatility state is rapidly changing. Whereas the standard models can take a while to adjust the volatility process, the use of realized measures can speed the adjustment up. By comparing DNS-GARCH-1 and DNS-GARCH-3 models to DNS-RGARCH-1(-lev) and DNS-RGARCH-3(-lev) respectively, we may see that the realized models demonstrate higher predictive likelihoods at all the considered horizons for the short- and mid-term contracts. Qualitatively, this finding is similar to the results of Shin and Zhong (2015) who documented significant gains from the use of realized measures in the bond yield density forecasting within a DNS framework. However, in our case these gains are quantitatively much smaller.[31] The "realized" models always appear in the MCS (for both full and calm out-of-sample periods), which suggest that the gains from realized measures are systematically significant, though they are not so large. The relative performance of the models with and without realized measures is getting less discernible as the forecasting horizon increases. This is natural since the realized information about the current volatility state is more relevant for the short-term volatility (and density) predictions rather than for the far ahead forecasts.

Third, the comparison of models with the one and three dynamic volatility components (DNS-(R)GARCH-1(-lev) to DNS-(R)GARCH-3(-lev)) reveals a significant advantage of the three component specifications in the ability to produce density forecasts for the short and mid term contracts. This is especially evident from the results based on the calm forecasting sub-period where for the contracts below 2 years until maturity only 3-component models appear in MCS at all considered horizons. It may imply that the use of three dynamic

[31] For instance, realized models with the one dynamic volatility component, DNS-RGARCH-1(lev), provide about 8% gain in predictive likelihood compared to its non-realized counterpart DNS-GARCH-1 for the 1-day ahead horizon for the short and mid term futures. The corresponding gain for the 3-component DNS-RGARCH-3(lev) is much smaller and is about 2.5%. We note, however, that the latter result can be explained by relatively noisy realized measures used in the 3-component model specifications. The relatively small benefits from using realized measures in forecasting the oil futures densities resemble the findings of Lunde and Olesen (2014) from the electricity markets. In particular, they demonstrated that the use of realized volatility measures does not give any significant improvement in the density predictions of the nearby NOMXC forward returns in the context of the univariate Realized GARCH framework.

volatility components allows to better control the distribution of the futures term structure than the use of just one dynamic volatility component. This result conforms with the study of Creal and Wu (2014) who found that one volatility component is not sufficient to properly capture the distribution dynamics of the interest rate term structure. At the same time, we may see that for the long-term contracts ($>$ 4 years until maturity) the one component models provide superior predictive likelihoods at the horizons above 1 week ahead. The possible explanation of this effect is the following. The volatility of the level factor implied by the one-component models is systematically more conservative (higher) than the level volatility in the three factor models. Since the conditional variance of long term futures is mostly explained by the volatility of the level factor, one-component models provide more conservative density forecasts for the long-end contracts. Thus, the large mispredictions of the long futures prices at turbulent periods lead to less severe losses in the corresponding predictive likelihoods for the models with the one volatility component.[32]

Finally, we note that the leverage effect does not seem to improve the forecasting performance. We observe only slight gains from the introduction of the leverage effect for the short-term contracts up to 3 days-ahead forecasting horizons and only for the calm out-of-sample period.

2.7 Conclusion

In this paper, we analyze the role of time-varying volatility in modeling commodity futures prices. Our approach is based on a dynamic version of the classical Nelson-Siegel model. The model effectively reduces the dimensionality by modeling the whole term structure as a function of three factors: level, slope, and curvature. The dynamic factor structure allows for direct forecasting of the whole term structure of futures prices. We improve the density forecasting properties of the model by introducing time-varying volatility directly into the conditional variance of the three unobserved factors. When modeling the volatility we follow the Realized GARCH framework by utilizing precise volatility information extracted from high-frequency data. The realized measures provide a better signal about the latent volatility than the daily returns used in conventional GARCH models. The volatility specification is therefore better suited to periods where the volatility changes rapidly. This makes the model very attractive for forecasting future volatility and in turn providing superior density forecasts.

Empirically we apply the model to data on light crude oil futures. By an extensive out-of-sample forecasting exercise we examine to what extent time-varying volatility improves the density predictability and whether we gain from the use of realized measures. We evaluate the forecast performance based on two out-of-sample periods: one sub-period with regular price dynamics and the full out-of-sample period containing both calm periods and periods with sharp price changes. We examine the forecasting performance for different forecast horizons and for different maturity categories of the futures prices. We use the log predictive likelihood as a criterion for forecast evaluation and compare the performance of the competing models using the model confidence set.

We find that the models with time-varying volatility strongly outperform the constant volatility benchmark. The benefit from time-varying volatility is generally larger for short horizon forecasts and for futures contracts with short/medium time to maturity. The use of realized measures provides small but systematic gains in density forecasting relative to standard GARCH volatility. When comparing the one- and three factor specifications of the volatility we find that the three component specification generally produces better density forecasts except for long maturity contracts at long forecast horizons. Finally, we found that leverage effects do not seem to

[32] As we can see from the results, for the long-term contracts the gains from the time-varying volatility (and the realized measures) are small and insignificant for the long enough forecasting horizons (which is especially evident from the results based on the full out-of-sample period). One possible reason is an inability of the DNS models to appropriately capture the variance dynamics of the long-term futures prices. When the time-varying volatility is introduced into the three Nelson-Siegel factors, the variance of the term structure long-end is regulated by the variance of the level factor solely, whereas the volatilities of the short and mid futures are additionally explained by the variances of the slope and curvature factors respectively. Since the fraction of the futures with long maturities is relatively small in our panel, the volatility of the level factor may be not very well adjusted to the volatility dynamics of the long term futures prices. Therefore, an introduction of an additional volatility factor associated with the long-term futures prices could lead to a potential improvement in the term structure modeling.

improve density forecasting significantly.

Table 2.5: Maximum likelihood parameter estimates.

		CV	GARCH-1	GARCH-3	RGARCH-1	RGARCH-1-lev	RGARCH-3	RGARCH-3-lev
λ	10^3	**5.800**	**5.734**	**5.829**	**5.774**	**5.772**	**5.770**	**5.803**
σ_w^2	10^5	**4.493**	**4.466**	**4.471**	**4.472**	**4.472**	**4.467**	**4.455**
β		**0.633**	**0.634**	**0.632**	**0.632**	**0.633**	**0.633**	**0.631**
ρ_{ls}		**0.030**	**−0.092**	−0.047	**−0.113**	**−0.111**	**−0.081**	**−0.097**
ρ_{lc}		**−0.504**	**−0.535**	**−0.505**	**−0.531**	**−0.533**	**−0.563**	**−0.537**
ρ_{sc}		**−0.139**	−0.030	−0.047	−0.005	−0.005	−0.011	−0.003
σ_l^2	10^3	**0.168**						
σ_s^2	10^3	**0.351**						
σ_c^2	10^3	1.132						
a_s	10^4		**1.458**		1.274	1.140		
b_s			**0.957**		**1.006**	**1.094**		
a_c	10^7		88.809		0.044	0.000		
b_c			**6.078**		**5.578**	**5.571**		
γ_{0l}			**0.000**	**0.000**	−0.023	−0.062	**0.147**	0.127
γ_{0s}				**0.000**			−0.068	0.009
γ_{0c}				**0.000**			−0.033	−0.042
γ_{1l}			**0.146**	**0.093**	**0.431**	**0.436**	**0.135**	**0.143**
γ_{1s}				**0.119**			**0.138**	**0.156**
γ_{1c}				**0.347**			**0.155**	**0.120**
γ_{2l}			**0.845**	**0.907**	**0.607**	**0.597**	**0.860**	**0.849**
γ_{2s}				**0.850**			**0.838**	**0.828**
γ_{2c}				**0.629**			**0.822**	**0.859**
ξ_l					−0.382	−0.414	**−2.166**	**−2.117**
ξ_s							−0.537	**−1.107**
ξ_c							−0.631	−0.478
ϕ_l					**0.864**	**0.861**	**0.913**	**0.916**
ϕ_s							**1.047**	**0.981**
ϕ_c							**1.026**	**1.055**
$\sigma_{u,l}^2$					**0.154**	**0.138**	**1.278**	**1.085**
$\sigma_{u,s}^2$							**0.632**	**0.538**
$\sigma_{u,c}^2$							**1.082**	**0.921**
δ_{1l}						**−0.049**		0.040
δ_{1s}								−0.041
δ_{1c}								−0.030
δ_{2l}						**0.079**		**0.232**
δ_{2s}								**0.190**
δ_{2c}								**0.191**
$\log \mathcal{L}$		160,551.5	161,130.2	161,255.1	160,101.6	160,230.1	151,578.7	152,128.9
$\sqrt{250E[V_l]}$		0.205	0.235	0.195	0.223	0.221	0.224	0.220
$\sqrt{250E[V_s]}$		0.296	0.299	0.273	0.286	0.286	0.279	0.286
$\sqrt{250E[V_c]}$		0.532	0.582	0.559	0.526	0.522	0.549	0.487

Maximum likelihood parameter estimates for the seven different models. Significance at a 5% level is indicated with bold face. For visualization some rows have been scaled by the number given in the second column.

Full forecasting period (January 2, 2014 - January 16, 2015)

		Specification of factor variance dynamics within DNS model					
Maturity Category	CV	GARCH-1	GARCH-3	RGARCH-1	RGARCH-1-lev	RGARCH-3	RGARCH-3-lev
				1 day-ahead			
$<$ 3 months	*2.656*	0.174	0.248	0.242	0.249	0.270	0.271
3-6 months	*2.784*	0.146	0.243	0.237	0.244	0.268	0.272
6-12 months	*2.921*	0.168	0.257	0.253	0.259	0.285	0.289
1-2 years	*3.094*	0.148	0.286	0.274	0.277	0.308	0.306
2-4 years	*3.126*	0.246	0.326	0.322	0.317	0.344	0.326
$>$ 4 years	*3.210*	0.239	0.247	0.241	0.240	0.272	0.274
all maturities	*3.001*	0.178	0.275	0.268	0.271	0.299	0.296
				2 days-ahead			
$<$ 3 months	*2.368*	0.200	0.249	0.228	0.237	0.265	0.263
3-6 months	*2.500*	0.191	0.249	0.230	0.237	0.271	0.270
6-12 months	*2.633*	0.203	0.271	0.254	0.260	0.300	0.301
1-2 years	*2.787*	0.238	0.308	0.281	0.285	0.327	0.327
2-4 years	*2.861*	0.261	0.333	0.301	0.300	0.324	0.337
$>$ 4 years	*2.884*	0.246	0.223	0.245	0.246	0.273	0.275
all maturities	*2.707*	0.227	0.286	0.265	0.270	0.305	0.307
				3 days-ahead			
$<$ 3 months	*2.176*	0.207	0.255	0.233	0.240	0.272	0.269
3-6 months	*2.305*	0.187	0.250	0.233	0.239	0.275	0.273
6-12 months	*2.435*	0.196	0.279	0.261	0.266	0.310	0.309
1-2 years	*2.585*	0.203	0.302	0.291	0.293	0.333	0.339
2-4 years	*2.679*	0.245	0.307	0.296	0.295	0.331	0.332
$>$ 4 years	*2.696*	0.248	0.190	0.249	0.247	0.270	0.273
all maturities	*2.512*	0.211	0.279	0.271	0.274	0.311	0.313
				4 days-ahead			
$<$ 3 months	*2.033*	0.200	0.246	0.223	0.226	0.259	0.256
3-6 months	*2.157*	0.187	0.243	0.224	0.229	0.269	0.266
6-12 months	*2.286*	0.188	0.267	0.251	0.254	0.299	0.299
1-2 years	*2.443*	0.211	0.301	0.284	0.286	0.326	0.328
2-4 years	*2.528*	0.249	0.304	0.299	0.298	0.322	0.322
$>$ 4 years	*2.565*	0.243	0.137	0.239	0.239	0.247	0.250
all maturities	*2.368*	0.212	0.270	0.264	0.266	0.301	0.301
MCS coverage			99%		95%		90%

Table 2.6: Comparison of predictive log-likelihoods across the pool of competing models. In the first column, the mean predictive log-likelihoods for the constant volatility DNS model are reported. In the remaining columns, we provide changes in mean predictive log-likelihoods relative to the values reported in the first column. The background color indicates that the model is in MCS with a coverage 99%, 95%, and 90% depending on the color density.

Full forecasting period (January 2, 2014 - January 16, 2015)

	Specification of factor variance dynamics within DNS model						
Maturity Category	CV	GARCH-1	GARCH-3	RGARCH-1	RGARCH-1-lev	RGARCH-3	RGARCH-3-lev
				1 week-ahead			
$<$ 3 months	*1.916*	0.192	0.240	0.219	0.221	0.254	0.253
3-6 months	*2.039*	0.175	0.234	0.215	0.220	0.254	0.256
6-12 months	*2.167*	0.179	0.254	0.239	0.243	0.282	0.284
1-2 years	*2.321*	0.201	0.286	0.275	0.280	0.311	0.311
2-4 years	*2.404*	0.212	0.272	0.294	0.294	0.294	0.295
$>$ 4 years	*2.456*	0.228	0.092	0.235	0.235	0.230	0.231
all maturities	*2.248*	0.197	0.251	0.256	0.259	0.284	0.285
				2 weeks-ahead			
$<$ 3 months	*1.501*	0.166	0.222	0.217	0.219	0.264	0.261
3-6 months	*1.627*	0.147	0.209	0.212	0.212	0.249	0.250
6-12 months	*1.763*	0.138	0.210	0.222	0.224	0.244	0.249
1-2 years	*1.939*	0.158	0.218	0.249	0.249	0.240	0.243
2-4 years	*2.052*	0.167	0.147	0.250	0.249	0.191	0.202
$>$ 4 years	*2.154*	0.197	-0.104	0.184	0.187	0.146	0.158
all maturities	*1.867*	0.159	0.177	0.231	0.232	0.228	0.233
				3 weeks-ahead			
$<$ 3 months	*1.256*	0.123	0.235	0.240	0.238	0.294	0.295
3-6 months	*1.382*	0.094	0.217	0.227	0.229	0.270	0.270
6-12 months	*1.524*	0.075	0.198	0.226	0.231	0.244	0.243
1-2 years	*1.710*	0.088	0.181	0.239	0.239	0.202	0.203
2-4 years	*1.835*	0.037	0.063	0.209	0.208	0.111	0.115
$>$ 4 years	*1.945*	0.138	-0.382	0.138	0.142	0.091	0.102
all maturities	*1.637*	0.085	0.125	0.222	0.223	0.201	0.204
				1 month-ahead			
$<$ 3 months	*0.925*	-0.216	0.232	0.199	0.214	0.272	0.282
3-6 months	*1.027*	-0.399	0.209	0.183	0.208	0.234	0.242
6-12 months	*1.172*	-0.434	0.180	0.195	0.211	0.199	0.203
1-2 years	*1.371*	-0.276	0.107	0.201	0.206	0.105	0.129
2-4 years	*1.490*	-0.415	-0.238	0.159	0.170	-0.073	-0.000
$>$ 4 years	*1.650*	-0.003	-1.543	0.090	0.106	-0.072	-0.038
all maturities	*1.297*	-0.312	-0.055	0.182	0.194	0.109	0.135
MCS coverage			99%		95%		90%

Table 2.7: Comparison of predictive log-likelihoods across the pool of competing models. In the first column, the mean predictive log-likelihoods for the constant volatility DNS model are reported. In the remaining columns, we provide changes in mean predictive log-likelihoods relative to the values reported in the first column. The background color indicates that the model is in MCS with a coverage 99%, 95%, and 90% depending on the color density.

Forecasting sub-period with relatively mild price dynamics (March 3, 2014 - August 29, 2014)

		Specification of factor variance dynamics within DNS model					
Maturity Category	CV	GARCH-1	GARCH-3	RGARCH-1	RGARCH-1-lev	RGARCH-3	RGARCH-3-lev
				1 day-ahead			
$<$ 3 months	*2.805*	0.290	0.368	0.333	0.345	0.386	0.391
3-6 months	*2.947*	0.286	0.354	0.327	0.338	0.379	0.382
6-12 months	*3.077*	0.309	0.373	0.352	0.360	0.406	0.408
1-2 years	*3.212*	0.349	0.413	0.389	0.394	0.443	0.442
2-4 years	*3.305*	0.360	0.445	0.395	0.396	0.435	0.433
$>$ 4 years	*3.287*	0.300	0.325	0.321	0.320	0.331	0.334
all maturities	*3.136*	0.326	0.391	0.364	0.370	0.412	0.413
				2 days-ahead			
$<$ 3 months	*2.469*	0.305	0.398	0.352	0.368	0.417	0.422
3-6 months	*2.622*	0.305	0.386	0.351	0.363	0.416	0.418
6-12 months	*2.756*	0.339	0.415	0.388	0.398	0.456	0.459
1-2 years	*2.890*	0.382	0.457	0.426	0.431	0.494	0.493
2-4 years	*2.996*	0.412	0.513	0.448	0.448	0.507	0.504
$>$ 4 years	*2.963*	0.307	0.302	0.319	0.321	0.330	0.336
all maturities	*2.814*	0.355	0.429	0.397	0.404	0.458	0.459
				3 days-ahead			
$<$ 3 months	*2.282*	0.313	0.408	0.362	0.373	0.434	0.439
3-6 months	*2.431*	0.316	0.393	0.360	0.372	0.425	0.429
6-12 months	*2.567*	0.352	0.427	0.401	0.411	0.475	0.476
1-2 years	*2.707*	0.401	0.479	0.445	0.452	0.522	0.521
2-4 years	*2.815*	0.435	0.543	0.470	0.470	0.539	0.536
$>$ 4 years	*2.781*	0.313	0.282	0.317	0.315	0.332	0.334
all maturities	*2.629*	0.370	0.444	0.411	0.418	0.479	0.479
				4 days-ahead			
$<$ 3 months	*2.140*	0.309	0.399	0.357	0.369	0.422	0.423
3-6 months	*2.288*	0.316	0.391	0.359	0.371	0.423	0.426
6-12 months	*2.427*	0.351	0.427	0.402	0.410	0.475	0.475
1-2 years	*2.571*	0.410	0.492	0.457	0.460	0.533	0.532
2-4 years	*2.684*	0.448	0.561	0.482	0.482	0.554	0.553
$>$ 4 years	*2.655*	0.326	0.255	0.316	0.318	0.324	0.331
all maturities	*2.493*	0.376	0.448	0.416	0.422	0.483	0.483
MCS coverage			99%		95%		90%

Table 2.8: Comparison of predictive log-likelihoods across the pool of competing models. In the first column, the mean predictive log-likelihoods for the constant volatility DNS model are reported. In the remaining columns, we provide changes in mean predictive log-likelihoods relative to the values reported in the first column. The background color indicates that the model is in MCS with a coverage 99%, 95%, and 90% depending on the color density.

Forecasting sub-period with relatively mild price dynamics (March 3, 2014 - August 29, 2014)

		Specification of factor variance dynamics within DNS model					
Maturity Category	CV	GARCH-1	GARCH-3	RGARCH-1	RGARCH-1-lev	RGARCH-3	RGARCH-3-lev
			1 week-ahead				
$<$ 3 months	*2.022*	0.307	0.391	0.356	0.362	0.413	0.415
3-6 months	*2.172*	0.313	0.383	0.356	0.363	0.411	0.413
6-12 months	*2.317*	0.351	0.422	0.400	0.407	0.467	0.466
1-2 years	*2.464*	0.413	0.494	0.459	0.463	0.535	0.533
2-4 years	*2.580*	0.458	0.574	0.492	0.492	0.566	0.563
$>$ 4 years	*2.557*	0.332	0.240	0.312	0.316	0.332	0.335
all maturities	*2.384*	0.379	0.446	0.418	0.422	0.482	0.482
			2 weeks-ahead				
$<$ 3 months	*1.680*	0.295	0.364	0.344	0.351	0.393	0.391
3-6 months	*1.829*	0.301	0.357	0.348	0.349	0.388	0.384
6-12 months	*1.976*	0.343	0.398	0.395	0.399	0.444	0.442
1-2 years	*2.133*	0.414	0.480	0.468	0.469	0.521	0.518
2-4 years	*2.254*	0.453	0.546	0.486	0.483	0.531	0.529
$>$ 4 years	*2.230*	0.312	0.068	0.232	0.242	0.239	0.257
all maturities	*2.050*	0.372	0.412	0.409	0.412	0.455	0.454
			3 weeks-ahead				
$<$ 3 months	*1.490*	0.299	0.371	0.346	0.347	0.405	0.405
3-6 months	*1.639*	0.302	0.362	0.347	0.347	0.400	0.395
6-12 months	*1.786*	0.338	0.391	0.388	0.389	0.444	0.437
1-2 years	*1.941*	0.402	0.464	0.457	0.451	0.507	0.499
2-4 years	*2.052*	0.422	0.497	0.432	0.425	0.471	0.465
$>$ 4 years	*2.017*	0.276	-0.197	0.157	0.171	0.160	0.178
all maturities	*1.856*	0.360	0.375	0.390	0.389	0.436	0.432
			1 month-ahead				
$<$ 3 months	*1.302*	0.294	0.380	0.342	0.343	0.405	0.401
3-6 months	*1.448*	0.290	0.358	0.335	0.335	0.391	0.385
6-12 months	*1.602*	0.330	0.387	0.381	0.378	0.443	0.436
1-2 years	*1.753*	0.385	0.443	0.433	0.425	0.487	0.479
2-4 years	*1.857*	0.380	0.438	0.380	0.376	0.420	0.415
$>$ 4 years	*1.814*	0.221	-0.477	0.083	0.112	0.088	0.116
all maturities	*1.666*	0.339	0.334	0.364	0.363	0.414	0.411
MCS coverage			99%		95%		90%

Table 2.9: Comparison of predictive log-likelihoods across the pool of competing models. In the first column, the mean predictive log-likelihoods for the constant volatility DNS model are reported. In the remaining columns, we provide changes in mean predictive log-likelihoods relative to the values reported in the first column. The background color indicates that the model is in MCS with a coverage 99%, 95%, and 90% depending on the color density.

Appendix B

Supplement to Chapter II

B.1 Realized Variances for a Panel of Contracts

In this appendix, we investigate the co-dynamics of realized variances computed for contracts with different maturities. The purpose of this inspection is to validate the assumption that the variance of the oil futures term structure can be adequately captured by means of just few dynamic volatility components.

We consider the period between January 3, 2012 and January 16, 2015, which amounts to 785 trading days. This period is chosen due to availability of the large cross-section of traded contracts, so we are able to consider a broad spectrum of maturities. We partition this spectrum into 10 maturity categories: up to 2 month, 2-4, 4-6, 6-9, 9-12, 12-18, 18-24, 24-36, 36-48 and more than 48 months until expiration. For each trading day from the considered sample period and for each maturity category we compute a subsampled realized variance as in (2.24) using high frequency transaction data. If at a given day t we observe several traded contracts that belong to the same maturity category, we take an average of the corresponding realized variances. As a result, we have a balanced panel of realized variances that spans 10 maturity ranges and 785 trading days (see Figure B.1.1).

The matrix of correlations between the log realized variance series of different maturity ranges is provided in Table B.1.1. We can see that the realized variances are highly correlated across all maturity categories.

< 2	1.000	0.990	0.961	0.885	0.788	0.766	0.697	0.740	0.703	0.541
$2-4$		1.000	0.970	0.892	0.799	0.775	0.702	0.749	0.709	0.551
$4-6$			1.000	0.900	0.796	0.758	0.688	0.725	0.692	0.515
$6-9$				1.000	0.832	0.726	0.634	0.686	0.657	0.453
$9-12$					1.000	0.789	0.647	0.657	0.671	0.492
$12-18$						1.000	0.753	0.696	0.678	0.553
$18-24$							1.000	0.643	0.588	0.463
$24-36$								1.000	0.669	0.508
$36-48$									1.000	0.653
> 48										1.000

Table B.1.1: Correlation matrix of log realized variances computed for distinct maturity categories (January 3, 2012 - January 16, 2015). The maturity range is specified in the left column.

These sample correlations, however, are silent about the co-dynamic properties of the realized variance series as well as about the presence of some common factors in such dynamics. To gauge the possible co-variation we implement a principal component analysis. The results are provided in Table B.1.2.

PC	Variation	Total
1	0.690	0.690
2	0.137	0.826
3	0.048	0.874
4	0.038	0.913
5	0.033	0.945
6	0.027	0.972
7	0.019	0.991

Table B.1.2: Principal component analysis of the log realized variance panel. Seven main principal components appear in the first column. In the second column there is a percentage of the variation explained by the corresponding principal components. In the third column we report the cumulative fraction of the explained variation.

As it can be seen, the first principal component explains roughly two third of the total variation in the joint dynamics of the log realized variance series, whereas the first 4 explain more than 90% of the total variation.

On one hand, the principal component analysis indicates that the volatility panel reveals a significant degree of comovement. To a certain extent this finding supports the idea that even only one time-varying component can describe an appreciable amount of the common realized variance dynamics. Note that the realized variance of the shortest contract that is used as a volatility state variable in our one-component dynamic volatility specification has a correlation about 0.9 with the first principal component.

On the other hand, if just a one component is used, a sufficient amount of realized variance dynamics (about 30%) is left unexplained. This motivates to make use of several dynamic components in order to better capture the conditional second moment of the whole term structure distribution.

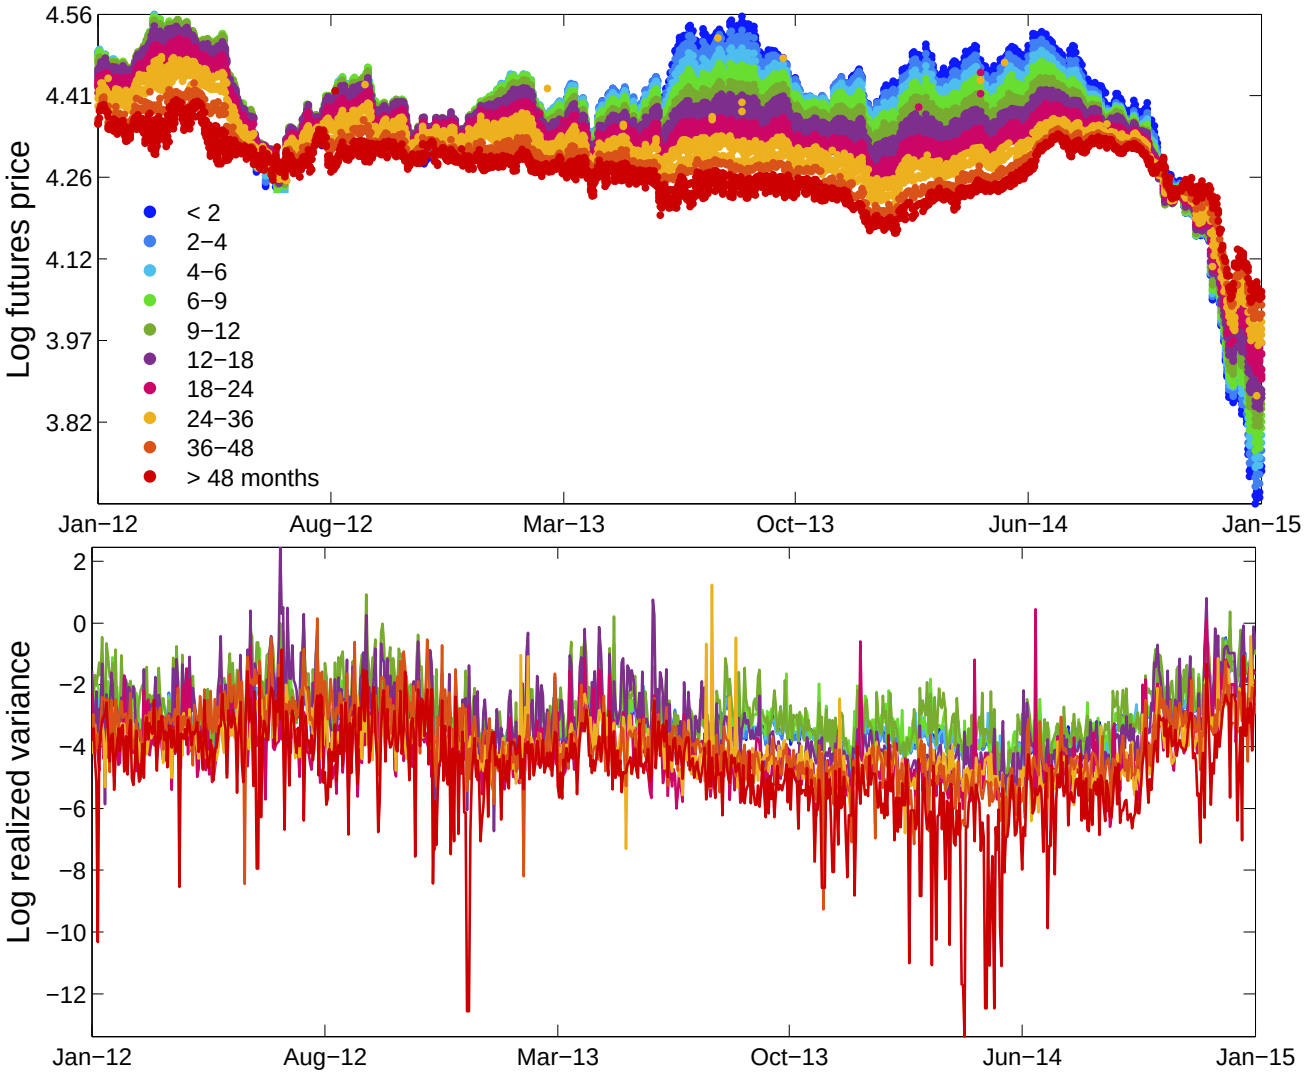

Figure B.1.1: Dynamics of the log prices (top panel) and the corresponding log annualized realized variances (bottom panel) for oil futures (January 3, 2012 - January 16, 2015).

B.2 State-Space Formulations

In the technical appendix we use the following notation inspired by Durbin and Koopman (2012).

The standard linear Gaussian state space system is given by

Observation equation

$$x_t = d_t + Z_t\alpha_t + \nu_t \qquad \nu_t \sim \mathcal{N}(0, V_t)$$

State equation

$$\alpha_{t+1} = T_t\alpha_t + c_t + W_{t+1}q_{t+1} \quad q_{t+1} \sim N(0, Q_{t+1})$$

Where the Kalman filter is defined by the following recursions:

Kalman filter

$$e_t = x_t - Z_t a_t - d_t \qquad F_t = Z_t P_t Z_t^{'} + V_t$$

Updating step

$$a_{t|t} = a_t + P_t Z_t^{'} F_t^{-1} e_t \quad P_{t|t} = P_t - P_t Z_t^{'} F_t^{-1} Z_t P_t$$

Prediction step

$$a_{t+1} = T_t a_{t|t} + c_t \qquad P_{t+1} = T_t P_{t|t} T_t^{'} + W_{t+1} Q_{t+1} W_{t+1}^{'}$$

where e_t denotes the one-step ahead forecast error. $F_t = \mathrm{Var}(e_t \,|\mathcal{F}_{t-1})$ is the conditional variance of the forecast errors. $a_{t|t} = \mathrm{E}(\alpha_t \,|\mathcal{F}_t)$ is the filtered states and $P_{t|t} = \mathrm{Var}(\alpha_t \,|\mathcal{F}_t)$ denotes the time t conditional variance of the states. $a_{t+1} = \mathrm{E}(\alpha_{t+1} \,|\mathcal{F}_t)$ is the one-step ahead predicted states and $P_{t+1} = \mathrm{Var}(\alpha_{t+1} \,|\mathcal{F}_t)$.

The log-likelihood function of a standard linear Gaussian state space system is given by

$$\log p(x_1, ..., x_T) = \sum_{t=1}^{T} \log p(x_t | \mathcal{F}_{t-1})$$

where

$$\log p(x_t | \mathcal{F}_{t-1}) = -\frac{n_t}{2}\log(2\pi) - \frac{1}{2}\left(\log|F_t| + e_t^{'} F_t^{-1} e_t\right) \tag{B.2.1}$$

The next sections contain the specific state space formulations of each model.

B.2.1 Constant Volatility

The constant volatility dynamic Nelson-Siegel model can be formulated directly in the standard linear Gaussian state space formulation. Matching the notation of the standard linear Gaussian state space system with the notation of the constant volatility dynamic Nelson-Siegel model we have that:

State space notation		Dynamic NS notation
x_t	$=$	y_t
Z_t	$=$	$(\Lambda_t f_t, -\beta C_t \Lambda_{t-1})$
α_t	$=$	$\begin{pmatrix} f_t \\ f_{t-1} \end{pmatrix}$
d_t	$=$	$\beta C_t y_{t-1}$
ν_t	$=$	w_t
T_t	$=$	$\begin{pmatrix} I_{3\times3} & 0_{3\times3} \\ I_{3\times3} & 0_{3\times3} \end{pmatrix}$
W_t	$=$	$\begin{pmatrix} I_{3\times3} \\ 0_{3\times3} \end{pmatrix}$
q_{t+1}	$=$	η_{t+1}
Q_{t+1}	$=$	$\begin{pmatrix} \sigma_l^2 & \rho_{ls}\sigma_l\sigma_s & \rho_{lc}\sigma_l\sigma_c \\ \rho_{ls}\sigma_l\sigma_s & \sigma_s^2 & \rho_{sc}\sigma_s\sigma_c \\ \rho_{lc}\sigma_l\sigma_c & \rho_{sc}\sigma_s\sigma_c & \sigma_c^2 \end{pmatrix}$
c_t	$=$	$0_{3\times1}$

B.2.2 GARCH-1

In the GARCH-1 model we follow Harvey et al. (1992).

State space notation		Dynamic NS notation
x_t	$=$	y_t
Z_t	$=$	$(\Lambda_t f_t, -\beta C_t \Lambda_{t-1}, 0_{n_t \times 1})$
α_t	$=$	$\begin{pmatrix} f_t \\ f_{t-1} \\ \eta_{t,l} \end{pmatrix}$
d_t	$=$	$\beta C_t y_{t-1}$
ν_t	$=$	w_t
T_t	$=$	$\begin{pmatrix} I_{3\times 3} & 0_{3\times 3} & 0_{3\times 1} \\ I_{3\times 3} & 0_{3\times 3} & 0_{3\times 1} \\ 0_{1\times 3} & 0_{1\times 3} & 0_{1\times 1} \end{pmatrix}$
W_t	$=$	$\begin{pmatrix} I_{3\times 3} \\ 0_{3\times 3} \\ (1,0,0) \end{pmatrix}$
q_{t+1}	$=$	η_{t+1}
Q_{t+1}	$=$	$H_{t+1} R H_{t+1}$
c_t	$=$	$0_{3\times 1}$

where

$$
\begin{aligned}
H_{t+1} &= \begin{pmatrix} \sqrt{h_{t+1,l}} & 0 & 0 \\ 0 & \sqrt{h_{t+1,s}} & 0 \\ 0 & 0 & \sqrt{h_{t+1,c}} \end{pmatrix} \\
R &= \begin{pmatrix} 1 & \rho_{ls} & \rho_{lc} \\ \rho_{ls} & 1 & \rho_{sc} \\ \rho_{lc} & \rho_{sc} & 1 \end{pmatrix} \\
h_{t+1} &= \gamma_{0l} + \gamma_{1l}\left(a^2_{t|t,7} + P_{t|t\,7,7}\right) + \gamma_{2l} h_t
\end{aligned}
$$

Here $a_{t|t,7}$ denotes the 7th element in $a^2_{t|t}$ and $P_{t|t\,7,7}$ is the $(7,7)$ element in $P_{t|t}$. Since $\eta^2_{t,l}$ is not directly observable we use $\mathrm{E}\left(\eta^2_{t,l}\,|\mathcal{F}_t\right) = a^2_{t|t,7} + P_{t|t\,7,7}$ as the "signal" in the volatility process. Hence, the model should be seen as an approximation to a "standard GARCH" model.

B.2.3 GARCH-3

State space notation		Dynamic NS notation
x_t	=	y_t
Z_t	=	$(\Lambda_t f_t, -\beta C_t \Lambda_{t-1}, 0_{n_t \times 3})$
α_t	=	$\begin{pmatrix} f_t \\ f_{t-1} \\ \eta_t \end{pmatrix}$
d_t	=	$\beta C_t y_{t-1}$
ν_t	=	w_t
T_t	=	$\begin{pmatrix} I_{3\times3} & 0_{3\times3} & 0_{3\times3} \\ I_{3\times3} & 0_{3\times3} & 0_{3\times3} \\ 0_{3\times3} & 0_{3\times3} & 0_{3\times3} \end{pmatrix}$
W_t	=	$\begin{pmatrix} I_{3\times3} \\ 0_{3\times3} \\ I_{3\times3} \end{pmatrix}$
q_{t+1}	=	η_{t+1}
Q_{t+1}	=	$H_{t+1} R H_{t+1}$
c_t	=	$0_{3\times1}$

where

$$
\begin{aligned}
H_{t+1} &= \begin{pmatrix} \sqrt{h_{t+1,l}} & 0 & 0 \\ 0 & \sqrt{h_{t+1,s}} & 0 \\ 0 & 0 & \sqrt{h_{t+1,c}} \end{pmatrix} \\
R &= \begin{pmatrix} 1 & \rho_{ls} & \rho_{lc} \\ \rho_{ls} & 1 & \rho_{sc} \\ \rho_{lc} & \rho_{sc} & 1 \end{pmatrix} \\
h_{t+1,l} &= \gamma_{0l} + \gamma_{1l}\left(a^2_{t|t,7} + P_{t|t\,7,7}\right) + \gamma_{2l} h_{t,l} \\
h_{t+1,s} &= \gamma_{0s} + \gamma_{1s}\left(a^2_{t|t,8} + P_{t|t\,8,8}\right) + \gamma_{2s} h_{t,s} \\
h_{t+1,c} &= \gamma_{0c} + \gamma_{1c}\left(a^2_{t|t,9} + P_{t|t\,9,9}\right) + \gamma_{2c} h_{t,c}
\end{aligned}
$$

Here $a_{t|t,i}$ denotes the i'th element in $a_{t|t}$ and $P_{t|t\,i,j}$ is the (i,j) element in $P_{t|t}$. Since $\eta^2_{t,k}$ is not directly observable we use

$$
\mathrm{E}\left(\eta^2_{t,k} \,|\mathcal{F}_t\right) = \begin{cases} a^2_{t|t,7} + P_{t|t\,7,7} & \text{for } k = l \\ a^2_{t|t,8} + P_{t|t\,9,9} & \text{for } k = s \\ a^2_{t|t,8} + P_{t|t\,9,9} & \text{for } k = c \end{cases}
$$

as the "signal" in the volatility process. Hence, the model should be seen as an approximation to a "standard multivariate GARCH" model.

B.2.4 RGARCH-1-lev

State space notation		Dynamic NS notation
x_t	$=$	y_t
Z_t	$=$	$(\Lambda_t f_t, -\beta C_t \Lambda_{t-1})$
α_t	$=$	$\begin{pmatrix} f_t \\ f_{t-1} \end{pmatrix}$
d_t	$=$	$\beta C_t y_{t-1}$
ν_t	$=$	w_t
T_t	$=$	$\begin{pmatrix} I_{3\times3} & 0_{3\times3} \\ I_{3\times3} & 0_{3\times3} \end{pmatrix}$
W_t	$=$	$\begin{pmatrix} I_{3\times3} \\ 0_{3\times3} \end{pmatrix}$
q_{t+1}	$=$	η_{t+1}
Q_{t+1}	$=$	$H_{t+1} R H_{t+1}$
c_t	$=$	$0_{3\times1}$

Hence, the standard space formulation is equivalent to the constant volatility model except for the fact that Q_{t+1} is time-varying and driven by realized measures. Here we have that:

$$H_{t+1} = \begin{pmatrix} \sqrt{\exp(h_t)} & 0 & 0 \\ 0 & \sqrt{a_s + b_s \exp(h_t)} & 0 \\ 0 & 0 & \sqrt{a_c + b_c \exp(h_t)} \end{pmatrix}$$

with

$$\begin{aligned} h_t &= \gamma_0 + \gamma_1 s_{t-1} + \gamma_2 h_{t-1} \\ s_t &= \begin{cases} \xi + \phi h_t + \delta(z_t) + u_t & \text{with leverage} \\ \xi + \phi h_t + u_t & \text{without leverage} \end{cases} \end{aligned}$$

In the Realized models we need to account for s_t being an observable. Hence, we need to add

$$\log p(s_t | y_t, \mathcal{F}_{t-1}; \theta) = -\frac{1}{2}\log(2\pi) - \frac{1}{2}\left(\log(\sigma_u^2) + \frac{u_t^2}{\sigma_u^2}\right)$$

to the NS log-likelihood contribution (B.2.1).

B.2.5 RGARCH-3-lev

State space notation		Dynamic NS notation
x_t	=	y_t
Z_t	=	$(\Lambda_t f_t, -\beta C_t \Lambda_{t-1})$
α_t	=	$\begin{pmatrix} f_t \\ f_{t-1} \end{pmatrix}$
d_t	=	$\beta C_t y_{t-1}$
ν_t	=	w_t
T_t	=	$\begin{pmatrix} I_{3\times 3} & 0_{3\times 3} \\ I_{3\times 3} & 0_{3\times 3} \end{pmatrix}$
W_t	=	$\begin{pmatrix} I_{3\times 3} \\ 0_{3\times 3} \end{pmatrix}$
q_{t+1}	=	η_{t+1}
Q_{t+1}	=	$H_{t+1} R H_{t+1}$
c_t	=	$0_{3\times 1}$

Hence, the standard space formulation is equivalent to the constant volatility model except for the fact that Q_{t+1} is time-varying and driven by realized measures. Here we have that:

$$H_{t+1} = \begin{pmatrix} \sqrt{\exp(h_{t,l})} & 0 & 0 \\ 0 & \sqrt{\exp(h_{t,s})} & 0 \\ 0 & 0 & \sqrt{\exp(h_{t,c})} \end{pmatrix}$$

with

$$\begin{aligned} h_{t,j} &= \gamma_{0,j} + \gamma_{1,j} s_{t-1,j} + \gamma_{2,j} h_{t-1,j} \\ s_{t,j} &= \begin{cases} \xi_j + \phi_j h_{t,j} + \delta_j (z_{t,j}) + u_{t,j} & \text{with leverage} \\ \xi_j + \phi_j h_{t,j} + u_{t,j} & \text{without leverage} \end{cases} \\ & \text{for } j = l, s, c \end{aligned}$$

In the Realized models we need to account for s_t being an observable. Hence, we need to add

$$\begin{aligned} \log p(s_t | y_t, \mathcal{F}_{t-1}; \theta) &= -\frac{1}{2}\log(2\pi) - \frac{1}{2}\left(\log(\sigma^2_{u,l}) + \frac{u^2_{t,l}}{\sigma^2_{u,l}}\right) \\ &\quad -\frac{1}{2}\log(2\pi) - \frac{1}{2}\left(\log(\sigma^2_{u,s}) + \frac{u^2_{t,s}}{\sigma^2_{u,s}}\right) \\ &\quad -\frac{1}{2}\log(2\pi) - \frac{1}{2}\left(\log(\sigma^2_{u,c}) + \frac{u^2_{t,c}}{\sigma^2_{u,c}}\right) \end{aligned}$$

to the NS log-likelihood contribution (B.2.1).

Chapter 3

Intra-daily Volatility Flow: How Fast Does the Information Arrive?

3.1 Introduction

The persistence of financial volatility is a firmly established empirical fact that is exploited intensively for modeling and forecasting purposes. Related studies primarily focus on forecasting uncertainty in financial and macro variables for the horizons from one day ahead and more. In this paper, we shift the focus to the opposite side and investigate how well financial volatility and uncertainty can be predicted within the daily span based on the information from high frequency transaction data observed over the short intra-daily intervals.

A natural approach to this question is to consider how much we can learn about the ex-post daily variance from the sub-periods between the market opening and a certain intra-daily moment of time. Indeed, we may expect that even the first few minutes after the start of trading may deliver a lot of valuable predictive information about the volatility that will be generated in course of the following hours. Clustering of intra-daily asset volatility around the market open and close is a theoretically supported and well documented stylized evidence (Admati and Pfleiderer (1988), Andersen and Bollerslev (1997)). The fact that in the beginning of a trading day a relatively large fraction of daily volatility is generated mechanically induces high correlation between the morning intra-daily and the ex-post daily variance measures. Additionally, intense morning transactions reflect a reaction of market participants to overnight (or overweekend) news as well as to the signals that come from other markets. Such information may have a long-lasting effect on the variance of intra-daily returns and might be a useful signal for daily variance prediction.

We study how fast a reasonably precise projection of a daily aggregate volatility level can be constructed by means of intra-daily information that arrives after stock market opening. Particularly, we compare the forecasting projections (both direct and indirect) that are made upon the realized volatility measures computed over gradually expanding intra-daily intervals that are supposed to incorporate increasingly more and more intra-daily information about the aggregated daily variance. We find that the first 15 minutes of trading are able to explain up to 70% of the variation in the ex-post daily variances in terms of the Mincer-Zarnowitz R^2. Moreover, the predictive performance of intra-daily measures based on high frequency data from the first 15-30 minutes of trading are on average equivalent to the standard model-based daily variance forecasts that employ information from the preceding days.

The study appeals to several empirical and methodological concerns. Firstly, we found that the forecasting performance differs significantly across different realized volatility estimators. It turns out that the intra-daily variance forecasting design might serve as a field for the comparison of small sample performance of different realized estimators. This is of significant methodological interest since it might help to better explore the advantages and limitations of distinct classes of estimators when those are used in constrained conditions. At second, the paper makes a step towards the framework that combines volatility forecasts that are made in

advance (by means of any appropriate models, such as GARCH or Stochastic Volatility) with the current intra-daily information in order to improve the projection of the latent variance level in real time. An ability to promptly update the current state of uncertainty in the very beginning of the day is definitely important for the risk management purposes. In addition, we discuss a dynamic approach to the intra-daily volatility updating and suggest a simple dynamic model to illustrate the idea.

This paper contributes to the vast literature on measuring, modeling and forecasting financial volatility. We now discuss in more details the connection with several closely related papers and briefly outline our main findings.

First, the analysis implemented in the current paper extends the work of Frijns and Margaritis (2008) who use the realized variance computed over a starting intra-daily interval for the ex-post daily realized variance forecasting in a regression framework. We expand their analysis by using more recent data sample, more fine intra-daily time intervals, extra information and multiple classes of realized measures. In addition, we investigate the predictive performance not only by forecasting ex-post daily realized measures, but also by forecasting the return density. We strengthen the results of Frijns and Margaritis (2008) by demonstrating a remarkably high predictive signal contained within the short (15-30 min) opening intervals for forecasting aggregated daily variance. Being large in the beginning of a day, the marginal improvement of the forecast quality from adding an extra intra-daily bin of data fades very quickly.

Conceptually, our study is closely interrelated with the study in Hansen and Lunde (2005) who exploit close-to-open (squared overnight return) and open-to-close (realized measure) information in order to construct an optimal (mean-squared-error minimizing) estimator of a close-to-close unobservable asset return variance. The clear analogy with our framework can be established if we note that in both studies the daily volatility measures based on "incomplete" information (that is, measured over just a part of the period of interest) are used as signals for an unobservable (at least by the moment of the signal extraction) volatility over the whole period of interest (the day).

Our paper suggests a path to complement an extensive comparison of realized volatility estimators performed in Liu et al. (2015). In our analysis, however, we focus on the small sample performance of estimators and assess it indirectly through the intra-daily forecasting framework. Liu et al. (2015) demonstrated that the classical 5-min realized variance estimator is hardly significantly outperformed by more recently developed estimators in measuring daily return variance. We provide an additional evidence of a systematic advantage of the estimators that exploit data on transaction frequency over the estimators that rely on sparsely sampled data when the information is limited and the role of estimation efficiency is better pronounced.

Finally, we contribute to the literature on adaptive time series models that use score driven dynamic specifications to control a drifting of parameter values over time (Blasques et al. (2014), Petrella and Delle Monache (2016)). We propose an extension of EWMA model that exploits realized measures as signals of the latent variance and is able to adjust the distributed lag weights in response to the changing volatility regimes. The model serves as an illustration of a dynamic optimal weighting scheme and closely interacts with our intra-daily forecasting framework.

We proceed as follows. In Section 3.2 we discuss the realized estimators that will be used throughout the analysis and describe the data. Section 3.3 contains a simple forecasting framework that is based on a linear projection of the daily realized measure on the intra-daily measure of variance with and with no presence of additional signals. In Section 3.4 we forecast the return density using a predictive likelihood approach. We outline a dynamic approach to volatility measurement and forecasting in Section 3.5. Section 3.6 concludes and provides an outlook of some further research directions. Several robustness checks are referred to Appendices.

3.2 Volatility Measures and Data

3.2.1 Econometric Preliminaries and Notation

As it is standard in the literature, we suppose that the efficient (no-arbitrage) logarithmic price $\{p_s^*\}_{s\geq 0}$ follows a semimartingale process (see, for example, Protter (2004)). The notable example of a semimartingale commonly used in economics and finance is a Brownian semimartingale process that satisfies

$$p_s^* = \int_0^s \mu_z dz + \int_0^s \sigma_z dW_z + N_s \tag{3.1}$$

where μ_s is an instantaneous drift that is predictable and locally bounded, σ_s is a strictly positive stochastic volatility process with a proper regularity assumptions, W_s is a standard Brownian motion independent from μ_s and σ_s, N_s is a pure jump process.

The equilibrium price process p_s^* is inherently latent. The price that is observed at transaction times $s(k)$ (where $k = 0, 1, 2, ..., K$ is the transaction index) during the open market period represents a contaminated signal of the latent efficient price

$$p_{s(k)} = p^*_{s(k)} + u_{s(k)} \tag{3.2}$$

The additive component $u_{s(k)}$ reflects the market microstructural noise arising from the rounding and data errors, price staleness, bid-ask bounces etc.[1] In the present paper we stick to the mild assumptions on the noise properties. These assumptions correspond to those made in Barndorff-Nielsen et al. (2008) and Hansen and Horel (2009), since we will use the related estimators in our analysis.[2]

The central object of interest from the perspective of the financial risk assessment is the quadratic variation over a certain period of time. In particular, for the period $s \in [s_a, s_b]$ the quadratic variation (QV) is formally defined as

$$QV_{[s_a,s_b]} = \int_{s_a}^{s_b} \sigma_s^2 ds + \sum_{s\in I_J} \left(p_s^* - p_{s_-}^*\right)^2 \tag{3.3}$$

where I_J denotes the set of time moments within $[s_a, s_b]$ when jumps occur and $p_{s_-}^*$ stands for the left limit of p_s^*.

Another important quantity is the integrated variance (IV) that comprises only continuous ("small" price moves) variation of the price process and ignores the variation stemming from jumps ("large" price moves)

$$IV_{[s_a,s_b]} = \int_{s_a}^{s_b} \sigma_s^2 ds \tag{3.4}$$

Thus, in the absence of the jump component in (3.1), QV and IV represent the same quantity. The motivation to consider the integrated variance in isolation partly comes from the evidence that the jump volatility component is far less persistent than the continuous volatility component. This is reflected in better forecasting properties of the latter (see Andersen et al. (2007)).

For notational convenience, we will denote trading days from the sample by indexes $t = 1, ..., T$ and normalize the open-to-close trading interval to $[0, 1]$. A typical trading period at NYSE is between 9:30a.m. and 4:00p.m. EST. We then partition it into N equidistant sub-intervals, $0 = s_0 < s_1 < s_2 < ... < s_N = 1$. In our empirical analysis we will use $N = 26$ which corresponds to 15 minutes between two adjacent knots. Though, it is possible to partition the trading period with an arbitrarily fine grid, e.g. with 1-min or even 1-sec intervals between the adjacent knots.

For each day t we obtain the sequence $X_t = \{x_{t,i}\}_{i=1}^N$, where $x_{t,i}$ is a realized measure that approximates

[1] The effects of microstructural noise on the realized volatility estimation were analytically and empirically investigated in Hansen and Lunde (2006b) and Bandi and Russell (2008).

[2] In particular, the nose is assumed to be stationary, but allowed to be serially dependent and correlate with the efficient price.

QV, IV or a similar variance quantity over an intra-daily interval $[s_0, s_i]$. Thus, a measure $x_{t,i}$ is constructed using the information that is available at day t for the period from the trade opening and until the intra-daily knot s_i. Below we discuss the choice of estimators that we use to obtain realized measures in our empirical analysis.

3.2.2 Estimators

The minutes after the stock market opening are considered as notoriously hard time for investors and traders. Activity of heterogeneously informed participants and high general uncertainty about the true price level increase the costs of trade and the risk of taking a wrong position. These and other issues lead to severe market microstructural effects, such as enormously wide bid-ask spreads, that complicate retrieving accurate signals of the latent volatility from the observed price changes.

The quality of our intra-daily forecasting projections is supposed to heavily depend on the precision of the extracted intra-daily volatility. Particularly, as long as we assume that an intra-daily variance measure contains a strong predictive signal about the ex-post variance level aggregated over the whole day, a more precisely extracted measure is supposed to provide an improved projection. To obtain realized measures we use distinct estimators. Thus, the intra-daily variance forecasting design might serve as a field for an empirical comparison of the small sample performance across different estimators as well as their abilities to mitigate microstructural noise effects that are especially severe in the beginning of a trading day. This is of significant methodological interest and might help to better explore the advantages and limitations of distinct classes of estimators.

We consider 4 types of realized estimators that differ both in frequencies of exploited data and in quantities they attempt to measure. Although in our empirical analysis we inspect just a single estimator from each of the considered classes, most of these estimators are widely representative and actively used in the empirical research.

Realized Variance. Our first estimator is the classical realized variance (RV) estimator (Andersen et al. (2001b), Andersen et al. (2001a), Barndorff-Nielsen and Shephard (2002b), etc.) computed at 5-minute frequency. This estimator measures QV of the efficient price process and minimizes the effects of microstructural noise by utilizing sparsely sampled observations. Since we consider only relatively liquid assets, we expect that 5-minute sampling is a sufficiently low frequency to mitigate the bias caused by microstructural noise effects. Although the design of RV estimator is probably the simplest if compared to the most of the existed and constantly introduced realized measures in the literature, the performance of RV in measuring daily return volatility is still remains to be a strong benchmark.[3] Particularly, in the extensive comparative study where more than 600 distinct volatility estimators were investigated, Liu et al. (2015) have found just a few realized measures that significantly outperform standard 5-min RV.

Threshold Variance. Secondly, we consider a truncated version of the realized variance estimator (TV) suggested in Mancini (2001), Mancini (2009). Similar to the realized variance, TV estimator represents a sum of squared intraday returns, but only those that do not exceed a certain threshold.[4] The truncation aims to "filter out" the portion of variance due to the jump component in (3.1), so TV can be treated as an estimator of the integrated variance (3.4). As in the case of RV, TV is computed with sparsely sampled returns reducing, thus, the influence of microstructural noise.

Realized Kernel. The third type of volatility estimators that we include in our analysis operates with frequently sampled price observations.[5] The main challenge of using very high frequency data in estimating asset return volatility relates to the disentangling the latent return variation from the microstructural noise effects. The realized kernel estimator (RK) suggested in Barndorff-Nielsen et al. (2008) is designed to capture

[3] We note, that in our analysis we consider a classical RV, but not a popular sub-sampled analogue where the average of realized variances computed at many distinct "grids" is used (see Zhang et al. (2005)). The reason is that such estimator can be treated as a case of the realized kernel estimator with the Bartlett kernel. Instead, we separately consider the realized kernel estimator with the Parzen kernel that has been shown to be more efficient (Barndorff-Nielsen et al. (2011b)).

[4] A widely accepted practice is to relate the threshold to a local IV estimate (see Corsi et al. (2010)). Since a local IV is difficult to obtain conditional on a small number of observations over short intra-daily periods, we use an approach outlined in Bollerslev and Todorov (2011).

[5] In particular, we use transaction data that is converted to a grid with a minimal interval of 1 second.

the long-run variance of the intra-daily return process in the presence of a stationary noise (that can be both autocorrelated and "endogenous" to the efficient price process, see Hansen and Lunde (2006b)) and, from this perspective, it is close in spirit to the kernel type HAC estimator (Newey and West (1987)). RK has been demonstrated to be a consistent estimator of QV with a decent small sample properties. Importantly, the fact that RK utilizes more observations than sparsely sampled estimators (such as RV or TV) promises better performance in measuring return volatility over the short intra-daily intervals. In our analysis, we use the "non-flat-top" Parzen kernel and implement "jittering" to dampen the end-effects (see Barndorff-Nielsen et al. (2009) for details).

Markov Chain Estimator. The last measure is the Markov chain estimator (MC) of Hansen and Horel (2009) which also employs data at a transaction frequency, but is constructed upon the specific assumptions about the observed price process $p_{s(k)}$. In particular, it is assumed that observed price increments, $\Delta p_{s(k)}$, follow a finite ergodic Markov chain process. Thus, MC naturally exploits the discreteness of the market price changes caused by the presence of the so-called "tick size" - a minimal possible price increment.[6] Since the discreteness of the observed price changes is usually treated as a factor that generates microstructural noise, it is even more interesting that MC estimator directly exploits this feature to disentangle the efficient price variation. The central idea rests on the fact that a finite ergodic Markov chain can be uniquely decomposed into the sum of a martingale difference sequence and a stationary process (see Hansen (2015)). The former component is associated with increments of an efficient martingale price process.[7] The latter can be referred to the microstructural noise $u_{s(k)}$ and, as in the case of RK, is allowed to be both autocorrelated and endogenous with respect to the efficient price. The Markov chain assumption allows to derive the long-run variance of the martingale component in a closed form, so it can be easily computed using an empirical estimate of the transition probability matrix of the observed price increments.[8]

The feasible implementation of MC estimator, however, requires to reasonably restrict the set of affordable price increments in order to keep the number of possible transitions moderate. More specifically, there should be specified some number z_{max} and all price increments that are larger than z_{max} tick size values (in absolute terms) are discarded.[9] Therefore, especially large price increments are getting censored and, hence, MC can't be considered as a consistent QV estimator.[10] As long as z_{max} is specified ad hoc, MC is obviously not a precise IV estimator as well. The conceptual difficulty of specifying such a threshold number z_{max} that would guarantee a consistency of IV estimation relates to the absence of a clear empirical definition of jumps when the price process is considered at higher frequencies.[11] Furthermore, the presence of microstructural effects would require to introduce additional assumptions on the noise dynamics for z_{max} to be rigorously specified. As a bottom line, MC estimator that we use here provides a measure of the quantity that is somewhere "in between" of IV and QV. On one hand, MC truncates large outlying price changes that may be caused by jumps or data errors. On the other hand, some episodes like short-lived local trends or volatility bursts (when the price changes considerably within a short time period), that would be potentially classified as jumps at 5-minute sampling, are not censored by MC.

[6] Starting from 2001 the NYSE tick size is equal to 1 cent.

[7] We note that due to the Markov chain assumption on the observed process, the latent martingale price process is also a Markov chain. Hence, the efficient price under assumptions of MC estimator does not follow a jump-diffusion process (3.1), but has similar long-run properties.

[8] In spite of the quite distinct assumptions, both RK and MC are based on the conceptually similar grounds. Namely, both estimators filter the variation associated with the latent martingale dynamics exploiting the long-run properties of the contaminated (observed) process.

[9] In our empirical application we set up a number z_{max} in such way that no more than 3% of transactions are censored.

[10] Alternatively, it would be possible to add the sum of squared truncated increments to the quantity obtained with MC estimator and get (possibly noisy) measure of QV. Nonetheless, we will not do it this way in order to avoid the contribution from possible noise effects such as bounce backs and data errors.

[11] In particular, Christensen et al. (2014) (and also Christensen et al. (2016)) demonstrated that most of episodes that are detected as "jumps" when the sparsely sampled price data is used can be attributed to the local bursts in volatility process σ_s (or in drift process μ_s), but not to price discontinuities.

Acronym	Estimator	Sampling	Jump censoring
RV	Realized Variance	Sparse	No
TV	Threshold Estimator	Sparse	Yes
RK	Realized Kernel	Tick	No
MC	Markov Chain Estimator	Tick	Yes

Table 3.1: List of volatility estimators used in the empirical analysis.

We note that when applying the realized estimators to the short intra-daily periods we expect the presence of an appreciable small sample bias. We, however, use the estimators in their standard forms since our main focus is not in the unbiased estimation of QV/IV over a given intra-daily interval, but rather in extracting the signal that correlates well with the aggregated ex-post daily variance.

3.2.3 Data

In our empirical analysis, we consider data on 12 blue chips. The selected stocks represent different sectors of the U.S. economy and are traded at NYSE and NASDAQ exchanges (see Table 3.2). The sample period starts on January 2, 2004, lasts until December 31, 2013 and comprises 2517 trading days. Due to the specific focus of our analysis we remove the days when the market was not open for the full trading day of 6.5 hours (from 9:30a.m. until 4:00p.m. EST). Such days are usually correspond to the national holidays (Independence Day, Thanksgiving Day, Christmas, etc.). In addition, we remove the days with the exceptionally high market volatility (5 trading days in total, including the Flash Crash day).

Ticker	Company	Sector	Exchange
AXP	American Express Company	Consumer finance	NYSE
BA	The Boeing Company	Aerospace and defense	NYSE
BAC	Bank of America	Banking	NYSE
CSCO	Cisco Systems, Inc.	Computer networking	NASDAQ
CVX	Chevron Corporation	Oil & gas	NYSE
GE	General Electric	Conglomerate	NYSE
INTC	Intel Corporation	Semiconductors	NASDAQ
JPM	JPMorgan Chase & Co.	Banking	NYSE
PFE	Pfizer Inc.	Pharmaceuticals	NYSE
UTX	United Technologies	Conglomerate	NYSE
WMT	Walmart	Retail	NYSE
XOM	Exxon Mobil Corp.	Oil & gas	NYSE

Table 3.2: List of stocks used in the empirical analysis.

Daily returns adjusted for dividends are obtained from the CRSP US Stock Database. Intra-daily returns and realized measures of variance are constructed using Trade and Quote (TAQ) data. The transaction data is cleaned accordingly to the guidelines suggested in Barndorff-Nielsen et al. (2009).

We intentionally focus our attention on liquid stocks. Most of the selected stocks are constituents of the Dow Jones Industrial Average (DJIA) index. Therefore, for each asset we have a plenty number of intra-daily transactions that allows to efficiently use RK and MC realized estimators.

3.3 Intra-daily (Realized) Variance Forecasting

3.3.1 Intra-daily Variance Distribution

We begin our analysis with the investigation of a periodic intra-daily "allocation" of an asset return variance. This exercise is somewhat similar to the estimation of intra-daily periodic volatility patterns (see Andersen and Bollerslev (1997), Andersen and Bollerslev (1998a)), though we address the issue from a different angle.

Particularly, we estimate average variance fractions that are generated within $[0, s_i]$, $i = 1, ..., N$, and are measured by distinct estimators (see details in Appendix C.1).

The upper numbers in the cells in Tables 3.4-3.5 report average estimated fractions of variance realized before the specified intra-daily time points (with 15-min step). The variance fractions are largely agree for different estimators and indicate that on average 15-17% of variance is generated within the first 15 minutes after market opening (9:45a.m. EST) that is less than 4% of the typical daily trading period. After one hour of trading (10:30a.m. EST) the average variance fraction constitutes 33-35%. It is interesting to note that stocks from the energy sector (CVX, XOM) are relatively less volatile in the morning. Thus, their average variance fraction for the first 30 minutes is about 16-17%, whereas for other stocks it is about 25%.

One interesting observation regarding the comparative performance of our realized estimators can also be made. An average bid-ask spread right after the stock market opening exceeds by several orders of magnitude an average bid-ask spread in the mid/end of the day.[12] Since a bid-ask spread is considered among the main sources of market microstructural noise (Roll (1984)), we may suspect that the first 15-min interval is about the most challenging intra-daily period for the volatility estimation. The average volatility fractions for the considered estimators are 16.6% for RV, 17.3% for TV, 15.9% for RK and 14.8% for MC. Thus, the realized measures based on transaction data (especially MC estimator) document systematically lower fraction of variance generated during the first 15 minutes of trading than measures that use sparsely sampled data. This may imply that the "high frequency" estimators filter a microstructural noise contribution better than the "low frequency" ones. It may follow, in particular, that the 5-min sampling frequency, that is used to construct RV and TV measures, is too high to alleviate the bid-ask spread effects in the opening trading interval.

3.3.2 Forecasting the End-of-Day Realized Variance

A high concentration of volatility near the market opening indicates a potential for making relatively precise projections of the ex-post daily variance using intra-daily measures obtained quickly after 9:30a.m. In this subsection we empirically assess the time that is needed to construct such projections.

We will start with predicting end-of-day realized volatility measures conditionally on realized measures obtained in the beginning of the day. We use a simple predictive regression of the following form

$$x_t = \gamma_{1,i} x_{t,i} + e_{t,s} \tag{3.5}$$

for $i = 1, ..., N$. The model (3.5) represents the linear projection of the end-of-day measures x_t onto the subspace spanned by the beginning-of-day measures $x_{t,i}$. Obviously, such forecasting design is far from being theoretically elaborated and rigour[13], but it may serve as a reasonable starting point.

We estimate (3.5) based on a 3-year rolling window by least squares. We re-estimate it every month in order to account for possible changes in the average distribution of the intra-daily variance. We denote by $\hat{x}_t$ the out-of-sample projections obtained with the coefficient $\gamma_{1,i}$ estimated in-sample. Projections $\hat{x}_t$ are treated then as forecasts of the end-of-day realized volatility measures conditionally on the measures computed within the intra-daily interval $[0, s_i]$. In order to assess the quality of such predictions we use R^2 of the Mincer-Zarnowitz regression (Mincer and Zarnowitz (1969), Meddahi (2002)).

The out-of-sample period lasts since January, 2007 until December, 2013 and consists of 1741 business days. The approach outlined above is similar to the one implemented in Frijns and Margaritis (2008). However, we use more recent data, exploit a range of principally different realized measures and consider a finer step for intra-daily intervals. In Tables 3.4-3.5 we provide the Mincer-Zarnowitz R^2 (middle numbers in cells). Results aggregated across all considered stocks are also illustrated in Figure 3.1. The main implications can be summarized as follows.

[12] Such phenomenon reflects the relatively large uncertainty about the latent efficient stock price during the first trading minutes. Then, as traders learn more about it during the trading session, the spread considerably shrinks. See, for example, this article in the Wall Street Journal (http://www.wsj.com/articles/early-birds-suffer-in-market-1442273794).

[13] At least due to an implicit assumption $E[x_t - \gamma_{t,i} x_{t,i} | x_{t,i}] = 0$ that likely does not hold in reality.

First, the opening intra-daily period is indeed informative about the aggregate daily volatility level. Although after the first 30 minutes of trading less than 25% of daily variance is on average generated, the measures constructed over the same interval are able to explain up to 60-80% of variation in the aggregate daily realized measures. These numbers are significantly higher than those reported in Frijns and Margaritis (2008).[14]

Second, the marginal gain from the use of an additional intra-daily information is sharply decreasing in the morning and then settles at a roughly constant rate. Although the gains are large during a relatively short opening interval, adding a 15-min intra-daily bin to the information set after 11:00 leads to about 1-2% of an average increase in R^2.

Third, the predictive power of the intra-daily volatility measures is very heterogeneous across considered estimators. If we look at the starting 15-min interval, RV attains on average 49.7%, TV - 34.5%, RK - 65.3%, and MC - 65.9%. The reason why the "sparse" measures (RV and TV) obtained after the first 15 minutes perform significantly worse in predicting themselves at the end of the day can be naturally attributed to their higher small sample error. In contrast, the "dense" measures (RK and MC) exploit much more intra-daily information about price dynamics and, as a consequence, demonstrate about 1.5-2 times better predictive accuracy.

A comment should be made about the jump-robust estimator TV. The predictive fit of TV is systematically lower than that of RV, even though both measures use sparsely sampled data. This evidence may be explained by a lower efficiency of the truncated measure and by the presence of important predictive information that is possibly contained in jumps. Thus, the jump truncation might reduce a short-run predictive ability. This finding can be considered as an additional argument in favour of the predictive potential of the jump variation.[15] Intuitively, jumps during the opening period might indicate an arrival of important news (that come in the overnight break, for instance). In a short-run it may cause an additional uncertainty about the new level of the equilibrium asset price and, as a consequence, a short-lived (intra-daily) burst in volatility and an increase in the overall daily variance. Therefore, if jumps are ignored, such volatility bursts are left partially unexplained by the forecast projection and it leads to a lower predictive R^2. Note, that the similar reasoning might be applied to justify the use of overnight returns as additional regressors for the short-run volatility forecasting. However, this hypothesis should be investigated more carefully in the future.

For robustness purposes, in Appendix C.2 we provide a subsample analysis of the forecasting exercise. In particular, we compare the predictive RMSE based on different subsets of days from the out-of-sample period. In Appendix C.3 we report the Mincer-Zarnowitz R^2 for an alternative log-specification of the predictive regression model (upper numbers in cells). More precisely, we run

$$\log x_t = \gamma_{0,i} + \gamma_{1,i} \log x_{t,i} + e_{t,s}$$

The predictive fit is on average slightly better in this case, possibly due to a regularization of the realized measures by the logarithmic transformation that dampens effects from outliers and near "non-stationary" episodes. Important to note that such specification implies the non-linear relationship between intra-daily and daily variances, $E[x_t - \gamma_{0,i} x_{t,i}^{\gamma_{t,i}} | x_{t,i}] = 0$.

3.3.3 Intra-daily Realized Variance Nowcasting

Next, we consider how an extra information about the daily variance level can affect the intra-daily forecasting from the previous section. In fact, even before the market opening we may possess many potentially valuable signals about the volatility that will realize over the trading day. For instance, Bollerslev et al. (2000), Andersen et al. (2003b), Andersen et al. (2007) among many others investigated the reaction of different securities traded

[14] They found that the corresponding average R^2 for NYSE and NASDAQ stocks is 40-45% for the RV estimator over the period 1996-2005. A part of the discrepancy can be attributed to the presence of the financial crisis 2007-09 in our out-of-sample period. A large increase of variation in daily realized volatility measures can mechanically rise R^2.

[15] Whereas Andersen et al. (2007) didn't find a strong predictive content in jumps, Corsi et al. (2010) challenged this finding and reported a significant impact of jumps on the future volatility by using newly introduced threshold multipower estimator to separate continuous and jump variations. Patton and Sheppard (2015) discovered that a jump direction also affects an impact on the future volatility.

on distinct markets on macroeconomic and financial announcements and found that volatility responds significantly. Since many relevant announcements are scheduled, the effect can be anticipated and this information can be used for the volatility prediction at a given day. In addition, the presence of an asymmetry of volatility responses to positive and negative news may be exploited as an additional information about the volatility level in the aftermath of announcements (Pagan and Schwert (1990), Engle and Ng (1993)).

Another valuable signal is contained in overnight and overweekend (close-to-open) returns. The large close-to-open absolute return may indicate a high volatility generated at night time (due to an arrival of important news, for example). Then, we may expect that the volatility still be high during the open-to-close period as well, since investors and traders will probably keep to react to the overnight news in the course of the trading session. Hansen and Lunde (2005) demonstrated the usefulness of overnight information in the ex-post measurement of the variance for the whole day. Andersen et al. (2011a) incorporated overnight returns in the joint dynamic model together with continuous and jump components of an asset return volatility.

In this paper we exploit the most straightforward predictive signals - volatility forecasts based on the observations from the previous days, or the model-based forecasts. We expect that this source suggests a considerable amount of the predictive content due to the strongly persistent dynamics of the asset return volatility. Conceptually, our exercise is similar to the nowcasting procedure where the forecast of some low frequently observed variable is updated as long as a more frequently observed information is arrived. Also, the predictive performance of the model-based forecasts can serve as a good benchmark for assessing the predictive performance of the intra-daily forecast projections.

To obtain the volatility forecasts based on the past information we exploit heterogeneous autoregressive framework (HAR) suggested by Corsi (2009). It represents a reduced-form model for the realized volatility measures and for the log-transformed measures the HAR model is formulated through the following linear regression

$$\log x_t = \beta_0 + \beta_1 \log x_{t-1} + \beta_2 \log x_{t-5:t-1} + \beta_3 \log x_{t-22:t-1} + \varepsilon_t \tag{3.6}$$

where $\log x_{t-k:t-1} = \frac{1}{k}\sum_{i=1}^{k} \log x_{t-i}$. Despite its simplicity, the HAR performs remarkably well in forecasting realized volatility and has recently become a widely applied tool in the related empirical research. Three autoregressive components in (3.6) stand for the preceding long, mid, and short-run volatility levels and reflect the cascade nature of volatility.[16] Moreover, the combination of just three autoregressive factors is able to successfully capture long memory effects in the dynamics. A log-transformation of the realized measures in (3.6) not only ensures positivity of the dependent variable, but also is in line with the empirical findings about the distribution of realized variance as long as the error term is assumed to be Gaussian.[17]

Denote by $h_t \equiv \exp(E_{t-1}[\log x_t])$ an ex-ante forecast of the realized measure x_t that is made using the HAR model (3.6).[18] We incorporate h_t into the analysis by including it in the regression (3.5) as an additional explanatory variable

$$x_t = \gamma_{1,i} x_{t,i} + \gamma_{2,i} h_t + e_{t,i} \tag{3.7}$$

for $i = 1, ..., N$. Equation (3.7) can also be considered as a slightly restricted version of the HAR model

[16] Volatility cascade is associated with an hierarchical dependence between the long-run and the short-run volatility levels. Particularly, long-run components are supposed to be useful in forecasting short-run components.

[17] More particularly, Andersen et al. (2001b) and Barndorff-Nielsen and Shephard (2002c) demonstrate that the realized variance estimator has an approximately log-normal asymptotic distribution. Although it is not necessarily the case for other realized measures, the log-transformed version of the HAR model demonstrates more stable performance (e.g., see Andersen et al. (2007)).

[18] Note that in case of the log-specification (3.6), we have that $\exp(E_{t-1}[\log x_t]) \neq E_{t-1}[x_t]$. Therefore, the Jensen effect will contribute a bias to the exponential transforms of the log-measure forecasts. Suppose that the error terms ε_t are iid Normal and have a constant variance σ_ε^2. Then, using that $\log x_t$ is conditionally log-normally distributed according to (3.6), the conditional expectation (day-ahead forecast) of the variance level is

$$E_{t-1}[x_t] = \exp(E_{t-1}[\log x_t] + \frac{1}{2}\sigma_\varepsilon^2) = A \exp(E_{t-1}[\log x_t])$$

where $A > 1$ is a constant term. Therefore, if we take $h_t \equiv \exp(E_{t-1}[\log x_t])$ as an ex-ante forecast, it will have a linear multiplicative bias that does not deteriorate the forecasting content of h_t in the predictive regression.

augmented by an additional intra-daily component $x_{t,i}$.[19] We naturally expect that as long as s_i approaches the end of a trading day, the importance of the ex-ante forecast h_t relative to $x_{t,i}$ in prediction of x_t is falling due to $x_{t,i} \to x_t$ as $s_i \to 1$ (therefore, $\gamma_{1,i} \to 1$ as $s_i \to 1$). Though, in the beginning of a trading day (when s_i is close to 0) the situation is opposite. We follow the same rolling estimation and forecasting scheme as before.

Tables 3.4-3.5 contain the Mincer-Zarnowitz R^2 from regressing the out-of-sample daily realized measures x_t on the forecasts $\hat{x}_t$ (bottom numbers in the table cells). In Figure 3.1 we also plot results averaged across all considered stocks. First of all, we note that the Mincer-Zarnowitz R^2 of the HAR volatility forecasts h_t alone is on average between 0.6 and 0.7 and the predictive performance of RV and TV forecasts is systematically lower than that of RK and MC.

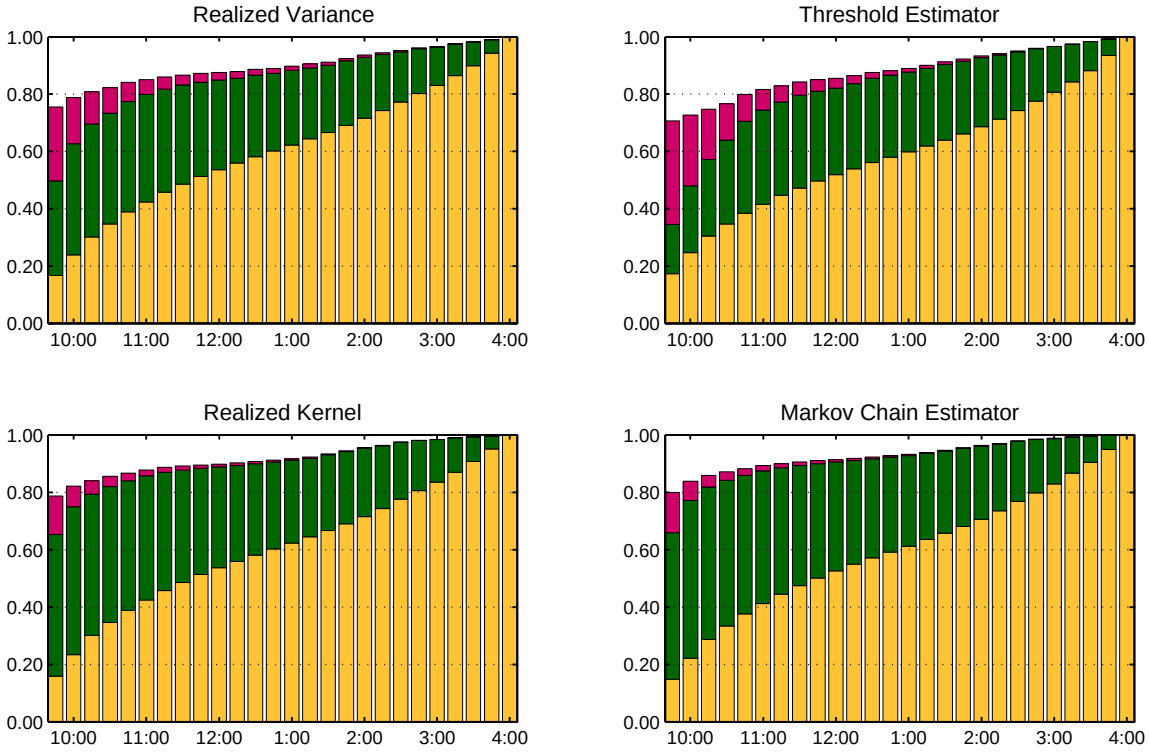

Figure 3.1: Subplots contain averaged (across all considered stocks) estimates of the intra-daily variance fractions and averaged results of the open-to-close realized measures forecasting for the specified estimators. The bottom (yellow) bars correspond to intra-daily variance fractions. The middle (green) bars correspond to the Mincer-Zarnowitz R^2 for the regressions that forecast the end-of-day realized measures from (3.5). The top (red) bars correspond to the Mincer-Zarnowitz R^2 for the regressions that forecast the end-of-day realized measures from (3.7). The out-of-sample period is January, 2007 - December, 2013.

Secondly, after the first 15 minutes of trading the effect from including h_t is considerable. The gains in R^2 constitute around 35% for the "sparse" estimators and 15% for the "dense" estimators. In result, the difference in predictive accuracy between the projections based on distinct estimators is blurred, but still remains appreciable. The average R^2 for RV and TV are 75% and 70%, whereas for RK and MC are 79% and 80% respectively.

Thirdly, although we can observe that the inclusion of the ex-ante volatility forecast h_t provides a significant benefit for the predictive R^2 in the beginning of the trading day, the marginal gains fade very quickly for the estimators that use transaction data. By 10:45a.m. the difference in R^2 between the forecasts based on (3.5) and (3.7) is on average less than 1%. As for the "sparse" estimators, the effect is more long lasting and at 11:30a.m. the average difference is still about 5%. This indicates that the intra-daily "dense" realized measures possibly possess more precise predictive content that crowds out the contribution of the ex-ante forecasts h_t very quickly.

[19] Alternatively, it would be possible to use an augmented HAR model of the following form

$$\log x_t = \beta_0 + \gamma_i \log x_{t,i} + \beta_1 \log x_{t-1} + \beta_2 \log x_{t-5:t-1} + \beta_3 \log x_{t-22:t-1} + \varepsilon_t$$

where all 4 coefficients are estimated jointly.

Also, in Appendix C.3 we provide the Mincer-Zarnowitz R^2 for the specifications with log-transformed variables (bottom numbers in cells).

3.4 Intra-daily Forecasting of Conditional Return Density

In the previous section we investigated how well distinct realized measures obtained over the intra-daily periods $[0, s_i]$, $i = 1, ..., N$, forecast themselves measured over the whole day $[0, 1]$. Although the results provide a good qualitative illustration of an intra-daily arrival rate of the predictive volatility content, there are several issues that are difficult to directly examine using such setting. First, so far we have used realized volatility measures as target variables, but the variable that is usually considered as an object of practical interest is the latent conditional return variance. Second, since for distinct realized estimators we used the corresponding daily realized measures as forecast targets, a direct comparison of the predictive performance across estimators is complicated by the lack of a common objective. In this section we extend our analysis in order to respond to these issues.

An unobservable nature of a conditional return variance significantly complicates an analysis of the forecasting performance. A popular approach in the literature is to evaluate a predictive performance with respect to an observed and possibly noisy proxy of the latent predicted variable (see Hansen and Lunde (2006a), Patton (2011)). Thus, it would be possible, for example, to approximate the latent return variance by a strongly robust estimator of the quadratic variation[20] and specify a proper loss function with respect to this proxy. We will follow a slightly different approach that would allow us to construct a forecast that is consistent with the tractable objective and to evaluate it with respect to this objective. Namely, we investigate an ability of intra-daily volatility measures to predict the conditional density of daily returns.

More precisely, we suppose that a daily open-to-close log-return, r_t, is distributed according to a particular conditional density, $r_t|v_t \sim p(r_t|v_t; \theta)$, where v_t is a (state) variable that corresponds (but not identical) to the conditional variance of r_t and θ stands for the parameter vector. In this case, the chosen forecast objective is the mean log-predictive likelihood, or LPL (see, e.g., Amisano and Giacomini (2007), Geweke and Amisano (2011)). Similar to the model from the previous section, we suppose that the variance parameter v_t is a scaled intra-daily realized measure $x_{t,i}$

$$v_t = \beta^x_{1,i} x_{t,i} \tag{3.8}$$

Coefficient $\beta^x_{1,i} \geq 0$ renders a proper scaling to account for a possible bias. Note that for $s_i = 1$ we use $x_{t,1} = x_t$ (daily realized measure) as a factor and v_t in that case is no more a forecast formally, but rather a measurement.

To estimate $\beta^x_{1,i}$ we maximize the log-likelihood function $\sum_t \log p(r_t|v_t; \theta)$ in-sample. Then we use estimated $\hat{\theta}$ to obtain out-of-sample forecasts of $\hat{v}_t$ conditional on intra-daily realized volatility measures $x_{t,i}$. We then compute the out-of-sample LPL with $\hat{v}_t$. Note that since both the estimation and the evaluation objectives are identical criterion functions, $\hat{\theta}$ is a consistent maximum likelihood estimate with respect to the out-of-sample objective (Dumitrescu and Hansen (2013)). Note that in case of the Gaussian conditional density the maximization of a predictive log-likelihood is identical to the minimization of the "quasi-likelihood" criterion, or QLIKE (see Patton (2011)), where the squared return is used as a proxy for the latent conditional return variance.

As before, we use a 3-year rolling estimation sample and re-estimate parameters every month. In Tables 3.6-3.9 (upper numbers in cells) we report the results of density forecasting for the whole out-of-sample period (1741 business days). The conditional return density is chosen to be Gaussian, so $r_t|v_t \sim \mathcal{N}(0, v_t)$. The values in the first column contain LPLs evaluated using variance forecasts based on the realized volatility measures from the opening 15-min intervals. The values in other columns correspond to the changes (relative to the first column) in LPLs that are evaluated using variance forecasts based on realized measures from the intra-daily

[20] Such as the 30-min RV or the realized range estimator (Martens and van Dijk (2007), Christensen and Podolskij (2007)), for instance.

intervals $[0, s_i]$, where s_i are specified in headers. Since we consider log-likelihoods, such changes approximately equal to the percentage differences in the average predictive likelihoods.

To be consistent with the analysis from the previous section, we also try an alternative 2-factor specification for the conditional variance that is an analogue of model (3.7)

$$v_t = \beta^x_{1,i} x_{t,i} + \beta^x_{2,i} h_t \tag{3.9}$$

where $\beta^x_{1,i}, \beta^x_{2,i} \geq 0$ and h_t is the HAR forecast of the daily realized measure of variance. The results corresponding to this specification are also provided in Tables 3.6-3.9 (bottom numbers in cells). The LPLs averaged across 12 considered stocks for both specifications of the variance parameter are depicted in Figure 3.2.

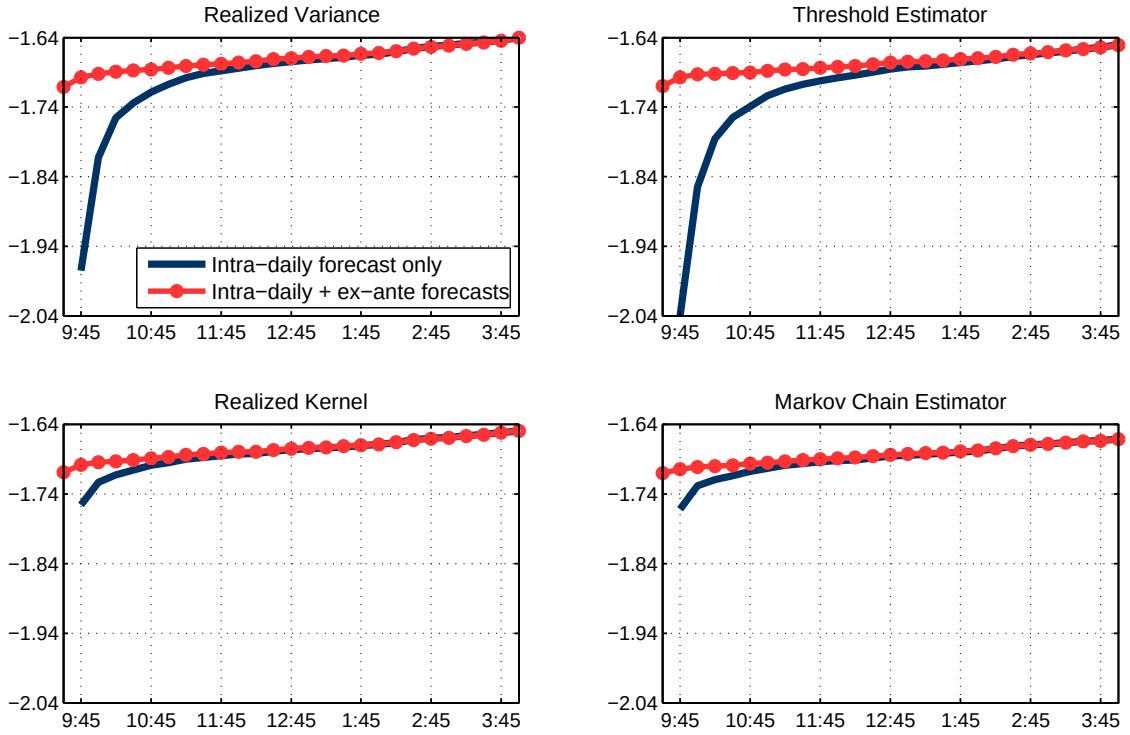

Figure 3.2: Subplots contain averaged (across all considered stocks) forecasting results of a daily return density for the specified estimators. The blue lines correspond to the averaged mean log predictive likelihood (LPL) computed using the variance specification given by (3.8). The red lines correspond to the averaged LPL computed using the variance specification given by (3.9). The out-of-sample period is January, 2007 - December, 2013.

The main implications of the density forecasting results can be summarized as follows. First of all, we corroborate the previous findings and document the fastly arriving predictive content. Particularly, for most of the realized measures and stocks it takes no more than 1-1.5 hours for the return density prediction with the 1-factor specification (3.8) to be on average as precise as the density forecast based by the HAR model (the case where only h_t is used).

Secondly, we confirm a large difference in the forecasting performance between the intra-daily measures obtained with distinct realized estimators. Figure 3.2 serves as a good illustration of large heterogeneity in signal-to-noise ratios implied in different intra-daily realized measures. In particular, lets consider LPLs evaluated after the first 15-min sub-interval with the 1-factor specification (3.8). We can see that the "dense" estimators, RK and MC, demonstrate on average 20% and 30% higher predictive log-likelihoods rather than RV and TV measures respectively. Note, however, that LPLs evaluated with the realized measures computed over the whole trading day ($s_i = 1$) are almost identical across all 4 estimators.

Finally, when we add the extra information h_t and use the 2-factor specification (3.9) to forecast return density, we observe a dramatic improvement in the LPL relative to the predictions made with the 1-factor specification around the market opening period. The difference in performance between realized estimators

almost completely disappears. However, the forecasting gains from using current intra-daily information are not too small and, by the end of the day, contribute on average about 10% to the predictive likelihood.

To provide a robustness check for the obtained results we repeat the density forecasting analysis assuming the conditional return density is Student's t with $\nu = 5$ degrees of freedom. Such distribution implies heavier tails as compared to the Gaussian density. Therefore, large density mispredictions are less severely reflected in LPL. The results are reported in Appendix C.4 and support all our qualitative implications.

In addition, we apply the model confidence set (MCS) of Hansen et al. (2011) to investigate whether we can statistically distinguish between the forecasting performance of different estimators evaluated at distinct intra-daily intervals. We use the sequences of out-of-sample predictive log-likelihoods as "gains" and look at the MCS at 95% significance level. For a given estimator we found that the density predictions made at the end of the trading day are significantly better than predictions made any time before (both for 1 and 2-factor variance specifications). As for the between-estimator comparison, the Markov chain measures are almost always alone in the MCS. This indicates a systematic advantage of MC measures in the forecasting (projection) accuracy.

3.5 Adaptive Approach to Volatility Modeling

The analysis in the previous sections reveals that after a certain intraday point of time the forecasting performance of intra-daily realized measures is tiny improved by adding an extra information from h_t. This result, however, is not surprising provided the way we use to combine the ex-ante (h_t) and the ex-post ($x_{t,i}$) signals. If h_t does not on average contribute extra relevant information to $x_{t,i}$ that can be useful to predict the target variable, then the coefficient β_2 will tend to zero.

The projection approach, though, is inherently static and very imperfect as far as we consider this problem from a dynamic perspective. Particularly, if we suppose that the conditional covariance between h_t and $x_{t,i}$ is non-constant, then both signals may become less or more informative depending on the sample period. In this section we will discuss this issue in the context of volatility measurement and modeling and illustrate the main ideas by a simple model.

3.5.1 Optimal Combination of Measures

Suppose there is a target variable y_t. Assume that variables x_t and h_t represent some noisy measures of y_t. For simplicity, we let these measures to be conditionally unbiased proxies of y_t for all t, that is $\mathrm{E}[x_t - y_t|\mathcal{F}_t] = \mathrm{E}[h_t - y_t|\mathcal{F}_t] = 0$. Further, we assume that $\mathrm{Var}[x_t - y_t|\mathcal{F}_t] = \sigma^2_{x,t}$, $\mathrm{Var}[h_t - y_t|\mathcal{F}_t] = \sigma^2_{h,t}$ and $\mathrm{E}[(x_t - y_t)(h_t - y_t)|\mathcal{F}_t] = \sigma_{xh,t}$, where $\mathcal{F}_t$ is some information set (sigma-algebra) such that both x_t and h_t are adapted to it. An optimal (in mean squared error sense) "estimator" of y_t from the affine class is expressed as (see Bates and Granger (1969), Timmermann (2006))

$$\hat{y}_t = \omega_t x_t + (1 - \omega_t) h_t \tag{3.10}$$

where

$$\omega_t = \frac{\sigma^2_{h,t} - \sigma_{xh,t}}{\sigma^2_{x,t} + \sigma^2_{h,t} - 2\sigma_{xh,t}} \tag{3.11}$$

Therefore, optimal weights for the two proxy variables in the estimator $\hat{y}_t$ depend on the whole matrix of conditional covariance between x_t and h_t. Note that even if just one component of such matrix is time-varying, the optimal weight ω_t is time-varying as well.[21]

In our context, y_t stands for the latent conditional return variance, x_t stands for an intra-daily measure ($x_{t,i}$, for distinct $i = 1, ..., N$), and h_t is an ex-ante model-based variance forecast. Both measures are exploited in the construction of the optimal predictor (3.10) for y_t. Figure 3.4 provides the optimal weights (3.11) averaged

[21] Except for some unrealistic cases, where the covariance components change in time together in a way that ω_t is left unchanged.

across 12 considered stocks for the static case where the covariance structure is assumed constant. Naturally, when the proxy of the latent target variable is noisy and differs from x_t, even at the end of the day both measures, h_t and $x_{t,i}$, have appreciable weights in the optimal predictor.

Lets now discuss an intuition for the case when covariance structure is time-varying. In periods when the ex-post realized measures are expected to be noisy, we put more weight on the ex-ante obtained forecasts. Conversely, when the ex-ante forecast h_t is too uncertain, we reallocate more weight on x_t.

Although the variance $\sigma^2_{x,t}$ of a realized measure can be approximated by means of an asymptotic distribution of x_t and estimated using high frequency data, the remaining elements from (3.11) are usually difficult to obtain. Returning to the static variance specification (3.9) from the previous section, we may suspect that the time-varying coefficients in (3.9) could possibly make a better job in capturing the latent volatility state.

3.5.2 Adaptive Volatility Modeling

In order to illustrate the idea with a particular example we consider a primitive model of asset return volatility and incorporate the intuition described above into the model.

Suppose that the latent conditional variance of an asset return y_t follows the random walk process

$$y_t = y_{t-1} + \eta_t \tag{3.12}$$

and assume that a realized measure x_t satisfies the measurement equation

$$x_t = y_t + \varepsilon_t \tag{3.13}$$

where η_t and ε_t are zero mean, orthogonal and serially uncorrelated random shocks with finite second moments. Taken together, equations (3.12) and (3.13) constitute the local level model. In case of Gaussian shocks it can be estimated with the Kalman filter and latent variance process y_t can be extracted. It was shown (see Muth (1960), Durbin and Koopman (2012)) that the exponentially weighted moving average (EWMA) recursion

$$h_t = \lambda x_{t-1} + (1-\lambda)h_{t-1}, \quad \lambda \in (0,1) \tag{3.14}$$

is able to provide an optimal (in sense of minimal mean squared error) forecast of x_t for a properly selected parameter λ. In case of constant volatility shocks, the corresponding value for λ is constant.

Equation (3.14) can be viewed as a simple example of a forecast combination. Indeed, both x_{t-1} and h_{t-1} provide the relevant signals about the future level of the latent variance y_t. A larger λ allocates more weight to the current volatility signal, whereas a smaller λ attaches more importance to the exponentially weighted sum of the past signals.

Suppose for a moment that the variances of η_t and ε_t are non-constant. Then parameter λ naturally becomes time-varying (this can be easily demonstrated by means of the Kalman filter recursions, assuming the shocks are Gaussian). As long as the measurement error ε_t is getting more volatile, λ_t decreases and assigns a lower weight to x_t reflecting the diminishing signal-to-noise ratio of x_t, and vice-versa. If the shock η_t to the latent variance process y_t is getting more disperse, then λ_t increases and puts more weight on x_t responding to the rising model uncertainty, and conversely.

In case we knew the volatilities of shocks at t, we would be able to specify the dynamic coefficient λ_t that would adapt the EWMA forecast h_t to the changing regimes of uncertainty. However, as we mentioned above, the values of such nuisance parameters are usually difficult to obtain in practice. Below we suggest an alternative data-driven way of making the EWMA model adaptive to the changing environment.

3.5.3 Score-driven EWMA with Dynamic Weighting Coefficients

Assume that the conditional distribution of daily asset return r_t is given by $p\big(r_t|h_t, \mathcal{F}_{t-1}; \theta\big)$, where h_t is a dynamic state variable that corresponds to the latent conditional variance, $\mathcal{F}_t$ is an information set available at

t and θ is a vector of constant parameters. Note, that we allow h_t to be biased with respect to the conditional return variance y_t. The dynamics of h_t is specified according to the EWMA structure

$$h_t = h_{t-1} + \lambda_t(x_{t-1} - h_{t-1}) = \lambda_t x_{t-1} + (1 - \lambda_t)h_{t-1} \tag{3.15}$$

where x_t is an ex-post realized measure of return variance and $\lambda_t \in (0, 1)$ is the time-varying weighting parameter. Note that the EWMA in (3.15) exploits realized measures to predict the volatility state h_t, but not squared returns r_t^2 as in the standard RiskMetrics™ model.

Time-varying parameter λ_t is a key element that introduces an adaptive mechanism to the model. We reparametrize λ_t using logistic transform as $\tilde{\lambda}_t = \log\big(\lambda_t/(1-\lambda_t)\big)$ in order to avoid a problem with bounded support when specifying the dynamics for the time-varying parameter. We use dynamic conditional score (DCS) approach (Creal et al. (2013), Harvey (2013)) to model the dynamics of $\tilde{\lambda}_t$

$$\tilde{\lambda}_t = \tilde{\lambda}_{t-1} + \gamma s_{t-1} \tag{3.16}$$

where $s_t = S_t^{-1}\nabla_t$ is a scaled score function and γ is a constant parameter. Score function is defined as $\nabla_t = \partial \log p\big(r_t|h_t, \mathcal{F}_{t-1}; \theta\big)/\partial\tilde{\lambda}_t$ and S_t is an appropriately chosen $\mathcal{F}_t$-measurable scaling function.[22] The idea behind such dynamic mechanism is to drive the time-varying parameter along the steepest direction of improving the local likelihood fit. By construction, scaled score process s_t is a martingale difference sequence, therefore dynamics in (3.16) is essentially a random walk. We select such parsimonious specification to simplify an exposition, though a more general ARMA process for $\tilde{\lambda}_t$ can be assumed. The score-driven updating term s_t in (3.16) is supposed to incorporate relevant information about the relative importance of lagged volatility signals and extracts it from the slope and curvature of the likelihood function.

We choose a heavy-tailed Student's t conditional density for modeling daily returns. Such choice is appropriate since it not only captures a well documented excess kurtosis in return densities, but also increases a general robustness of the model dynamics to the presence of outlying return observations. Thus, the conditional density reads

$$p\big(r_t|h_t, \mathcal{F}_{t-1}; \theta\big) = \frac{\Gamma(\frac{\nu+1}{2})}{\sqrt{\pi}\Gamma(\frac{\nu}{2})}\frac{1}{\sqrt{y(h_t)}}\left(1 + \frac{r_t^2}{y(h_t)}\right)^{-\frac{\nu+1}{2}} \tag{3.17}$$

where we set the degrees of freedom to $\nu = 5$.[23] Variable $y(h_t)$ is a latent conditional return variance that is a deterministic function of the volatility state variable h_t. In particular, we assume a presence of a multiplicative measurement bias, so $y_t = y(h_t) = (\nu - 2)bh_t$ with a bias parameter $b > 0$.[24] The conditional score is then given by

$$\nabla_t = \frac{b(\nu-2)}{2}\frac{\nu r_t^2 - y(h_t)}{y(h_t)\cdot(y(h_t) + r_t^2)}(x_{t-1} - h_{t-1})\lambda_t(1-\lambda_t) \tag{3.18}$$

The conditional score (3.18) drives the dynamics of λ_t and determines how the model manages the relative weighting of recent and old signals in order to produce an update of a volatility state variable h_t in (3.15).

As it can be seen from expression (3.18), if both terms, $\nu r_t^2 - y(h_t)$ (A) and $x_{t-1} - h_{t-1}$ (B) are of the same sign, the score ∇_t is positive and leads to an increase in λ_{t+1}. Thus, it suggests to shift the weight onto the

[22] The purpose of a scaling matrix process S_t is to standardize (or just to make more regular) the conditional score vector process ∇_t. The choice of S_t is not restricted to a particular specification and each particular form of S_t implies potentially different dynamics for the modeled parameters. A typical choice is $S_t = \mathcal{I}_t^{1/2}$, where $\mathcal{I}_t = E[\nabla_t\nabla_t'|\mathcal{F}_{t-1}]$ is a conditional Fisher information matrix. Such choice makes the scaled score variable s_t conditionally standardized with unit variance. Moreover, as it was demonstrated in Koopman et al. (2015), such specification of a scaling matrix positively affects a predictive accuracy of some score-driven models. Several other possibilities of scaling are considered in Creal et al. (2013). For example, $S_t = \mathcal{I}_t$, $S_t = diag\big(\mathcal{I}_t\big)$ or simply $S_t = I$ can be used instead, where I is the identity matrix.

[23] It is also possible to estimate ν as an additional static model parameter.

[24] We introduce bias parameter b to account for a possible bias of the realized variance measures. Such bias correction would be especially useful if we specify daily return r_t as a close-to-close return. Then we expect that $b > 1$ since the realized variance measure x_t is obtained at the open-to-close period and does not account, thus, for the close-to-open portion of variance.

current realized measure x_t when producing a forecast for h_{t+1}. Intuitively, when both terms, A and B, are of the same positive (negative) sign, it likely implies that the previous upward (downward) step $x_{t-1} - h_{t-1}$ of the variance process h_t in (3.15) resulted in insufficiently high (low) level of the state variable $y(h_t)$ compared to the latent variance proxied by the squared return. Thus, the conditional score ∇_t suggests to increase the weight on $x_t - h_t$ further in order to approach the latent volatility level more closely in a following step and get an increase in likelihood, as a consequence.

A similar intuition holds for the case when (A) and (B) are of different signs, ∇_t is negative and cause λ_{t+1} to decrease. In such case it is likely that the preceding upward (downward) step $x_{t-1} - h_{t-1}$ in (3.15) resulted in overshooting the targeted latent return variance. The model responds to this by reducing the relative weight for the newcoming signals and by increasing the persistence of the state variable h_t in order to avoid the growing discrepancy between y_t and the latent target variable.

Although these effects arise mechanically from introducing the score-driven updating term s_t into the dynamics of λ_t, they provide a nice illustration of the adaptive approach to the modeling. We can also notice an analogy with the optimal forecast combinations from the previous subsection. In particular, the case when (A) and (B) are of different signs ($\nabla_t < 0$) can be interpreted as a possible increase in the measurement error uncertainty in (3.13) (so the measurement x_t is getting more volatile than the target variable y_t). On the contrary, if (A) and (B) are of the same sign ($\nabla_t > 0$), it may be interpreted as a rise of uncertainty in the dynamics of the latent variable y_t. This prompts the model to be more adaptive in the face of a possibly changing environment and reduce the relative importance of the past observations when producing forecasts.

The suggested model is closely related to the recent developments in the time series models with score-driven dynamic parameters. The EWMA model with score-driven volatility was suggested in Lucas and Zhang (2016). In contrast to that model, we introduce the time-varying weighting parameter λ_t and use the score to govern its dynamics. In addition, we employ realized measures instead of squared returns in order to update the conditional return variance. Our model is similar to Blasques et al. (2014) and Petrella and Delle Monache (2016) who consider autoregressive adaptive models with time-varying score-driven coefficients. Apart from them, we introduce an adaptive autoregressive mechanism into the conditional variance, but not into the conditional mean of an observed process.

Another parallel can be drawn with Bauwens and Storti (2009) and their component GARCH model, where the conditional return variance is a time-varying weighted combination of two distinct variance processes. Such specification allows to capture non-stationarity and structural break effects. A somewhat similar approach to realized volatility modeling is suggested in Gallo and Otranto (2015), Gallo and Otranto (2016), where the volatility specification is exposed to the regime switching and/or smooth transition dynamics. In contrast, in the EWMA framework we alter the volatility process by dynamically changing the weights for the recent and old measures of the latent variance.

3.5.4 Empirical Application

We estimate the model on the full sample (2517 business days) for each of 12 considered stocks. In addition, we estimate the restricted model with $\gamma = 0$ in (3.16), so the weighting parameter becomes a constant (and is equal to the initial value λ_1 that is treated as a constant parameter in the estimation).

In Figure 3.3 we plot the weighting coefficients λ_t filtered from both main and restricted models for GE stock (bottom panel). We also include the plots with the annualized realized volatility (top panel) and the log-realized quarticity (medium panel) estimated using intra-daily data for the same period. The restricted model finds $\lambda = 0.344$. This value is much higher than the coefficients typically suggested in the RiskMetrix™ models. This is quite natural, since a signal-to-noise ratio of realized measures is much higher than that of squared returns. The main (unrestricted) model indicates that λ_t exhibits a considerable extent of variation. As we can see from the plot, λ_t rises significantly above the static level in 2008-09 and in 2011. It means that the model prefers to put a higher weight on the recent realized measures than on the old ones in the specification of the latent volatility dynamics. A possible interpretation of this phenomenon is following. When the stock

returns shift from a less to a more volatile regime (e.g., in the beginning of the Great Recession), the model attaches more weight to the current and recent realized measures which contain more relevant information about the new volatility regime than the old measures.

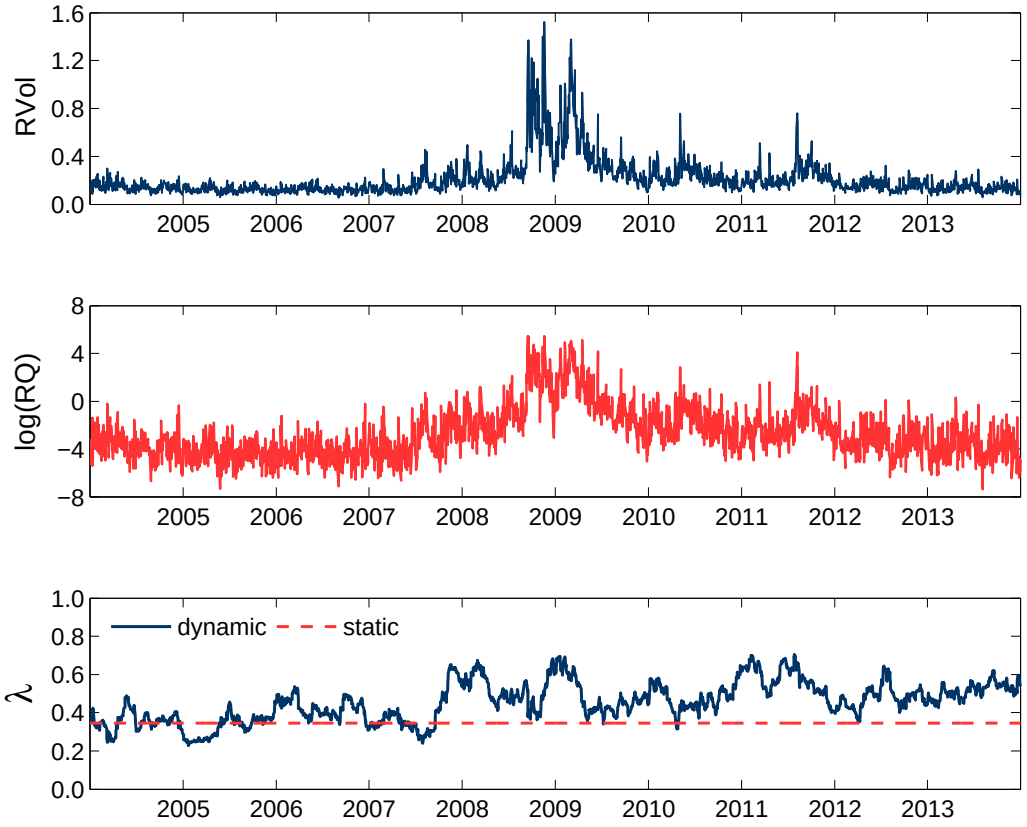

Figure 3.3: The series are obtained for the General Electric (GE) asset. The top panel shows the annualized volatility estimates obtained with the RV measures. The middle panel contains the logarithm of the realized quarticity estimates. The bottom panel contains the estimated weighting coefficients λ_t from the static (restricted) and the dynamic (unrestricted) coefficient EWMA models (from Section 3.5.3). The estimation period is January, 2004 - December, 2013.

In order to examine whether the model likelihood gets a statistically significant advantage from allowing the weighting parameter λ_t to vary, we compute the likelihood ratio statistics using the fact that the restricted model is nested in the unrestricted one. Table 3.3 contains the values of λ from the restricted models (upper numbers) for 12 stocks and 4 realized estimators. We notice that realized measures based on the transaction data (RK and MC) have on average higher weights λ as compared to the "sparse" measures (RV and TV). The middle numbers in cells provide p-values of the likelihood ratio test statistics for the null hypothesis $\gamma = 0$ (restricted model) and the bottom lines indicate a model preferred by BIC. In absolute majority of cases, the null is strongly rejected in favour of the dynamic specification of the weighting parameter.

3.5.5 Discussion and Extensions

To establish the link between the optimal combination of measures in the static measurement framework (Section 3.5.1) and the example of the dynamic model with the optimal combination of predictive signals (Sections 3.5.2-3.5.3) we discuss the trade-off that is relevant in both scenarios. On one hand, an ex-post realized measure x_t can be a noisy signal of the latent volatility level y_t, so an ex-ante forecast h_t might contain more precise information about y_t provided that the latent volatility is well described by the DGP of the forecasting model that generates h_t. Therefore, it might be preferable to increase the weight of h_t in the dynamic measurement/forecast. On the other hand, if there is a risk of changing environment (e.g., an economy moves into a more volatile regime) and the model may provide inaccurate forecasts h_t under the new conditions,

		AXP	BA	BAC	CSCO	CVX	GE	INTC	JPM	PFE	UTX	WMT	XOM
Realized variance	static λ	0.310	0.259	0.463	0.219	0.204	0.344	0.219	0.173	0.322	0.317	0.197	0.292
	LR (p-values)	0.085	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
	BIC prefers	Stat	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn
Threshold estimator	static λ	0.356	0.304	0.455	0.220	0.223	0.351	0.255	0.199	0.382	0.332	0.236	0.295
	LR (p-values)	0.076	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	0.105	<0.01
	BIC prefers	Stat	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Stat	Dyn
Realized kernel	static λ	0.380	0.344	0.501	0.271	0.238	0.390	0.257	0.218	0.389	0.589	0.210	0.331
	LR (p-values)	0.188	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
	BIC prefers	Stat	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn
Markov chain estimator	static λ	0.385	0.463	0.564	0.259	0.263	0.469	0.280	0.217	0.403	0.510	0.216	0.364
	LR (p-values)	0.243	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01	<0.01
	BIC prefers	Stat	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn	Dyn

Table 3.3: Table contains results of the statistical comparison of the in-sample fit between the static and the dynamic coefficient EWMA models (from Section 3.5.3). Rows correspond to the realized estimators used for the model estimation, columns correspond to the considered stocks. An intersection of a row and a column is a cell and each cell contains 3 entries. The top numbers correspond to the estimated λ from the static models. The middle numbers contain p-values of the likelihood ratio tests for the null of the static specification λ against the alternative of the dynamic specification λ_t. The bottom entries indicate a model (dynamic or static) preferred by BIC. The estimation period is January, 2004 - December, 2013.

ex-post measures x_t may contain a more reliable signal about latent y_t. Thus, putting more weight on x_t may improve the dynamic measurement/forecast. An adaptive mechanism that allows to balance this trade-off responding to the changing environment may provide an improvement in measuring, modeling, and forecasting of latent variables.

It is also interesting to discuss how this intuition relates to the models with a dynamic attenuation. The idea behind the attenuation[25] essentially corresponds to the first part of the trade-off from the above paragraph. Namely, more noisy signals should have lower importance (be attenuated) compared to less noisy signals when an unobserved variable is estimated. The class of models for the realized variance with the dynamic attenuation effect was recently proposed in Bollerslev et al. (2016). The authors introduce a time-varying autoregressive coefficient into the HAR model for a realized volatility. The dynamics of this coefficient is determined by the temporal variation of an attenuation bias which is approximated by means of the realized quarticity (RQ) measures. In result, the model takes a conservative stance when a realized measure has a large measurement error (large RQ) and "attenuates" the effect from such observation. This model, however, does not embed the second part of the trade-off described in the above paragraph. In particular, if we suppose that the latent state of volatility has shifted to the higher permanent level, then it will likely be reflected in the rise of RQ as well (see Figure 3.3). An attenuation model will respond in reducing the weights imposed on the new information. As a consequence, the transition to a newly established higher volatility level will be slow and the forecasting performance will be deteriorated during the transition period. In contrast, adaptive models are supposed to take into account both sides of the trade-off and shift to the new regimes faster.

Simple adaptive EWMA model from Section 3.5.3 can be extended along the several directions. So far a realized measure x_t appears in the model as a purely exogenous variable. It would be possible, however, to "close" the model by introducing a separate measurement equation for x_t in the spirit of Hansen et al. (2012). Such modification can potentially lead to a more precise signal extraction and a better estimation stability. Furthermore, the EWMA specification for volatility dynamics can be replaced by a more general autoregressive structure.

Another promising direction is to incorporate the proxies for latent measurement errors of x_t (like RQ for RV or analogous measures for other estimators) into the model. Being introduced explicitly, such information might

[25] See instructive discussion of the attenuation effect in the context of realized variance modeling in Hansen and Lunde (2014).

be helpful in solving the aforementioned trade-off by better identifying the risk of measurement inaccuracy and the risk of inconsistency between the model and the true DGP of the latent variable.

3.6 Intermediate Conclusion and Outlook

In this paper we empirically investigate how fast the information about the daily asset return variance arrives in the course of a trading day and how much an ex-ante variance forecast can contribute to this information flow. We find that the first few minutes of trading after market opening provide a relatively precise signal about the daily volatility level measured ex-post. We qualitatively confirm these findings by intra-daily forecasting of the daily return density. However, a series of additional robustness checks is needed to make more firm conclusions (e.g., use outlier-robust or non-linear forecast projections).

We document that the predictive content of realized measures that employ high frequency data at transaction level is strongly superior compared to estimators that use sparsely sampled data, but the difference is getting negligible closer to the end of the trading day as more observations are used to construct the measure. Therefore, our empirical design may serve as a natural ground for a more elaborate comparison of the small sample properties of different types of estimators as well as their (very) short-run forecasting performance.

Furthermore, we suggest a simple score-driven adaptive EWMA model that regulates in a dynamic manner the relative impact of old and recent measures on the variance forecasts. The adaptive mechanism provides an insight about a choice of the optimal combination of an ex-ante forecast and newly appeared signals (such as intra-daily/daily realized measures). Applying the model to data we find that dynamic weighting leads to statistically significant improvements. The idea can be developed further and adjusted with the intra-daily forecasting analysis. In particular, there can be suggested an unobservable component model that treats conditional variance of daily returns as a latent factor which underlies intra-daily measures. A principle of the dynamic signal combination can then be applied both to forecasting and intra-daily updating of the latent variance in a single modeling framework.

This paper is a step towards a more extensive and careful empirical analysis of the information content arriving with the intra-daily transaction flow. Possible extensions can be made in several directions. First, the predictive information contained in intra-daily volatility measures can be examined in the presence of additional control variables, such as days of the week, scheduled announcements, liquidity, market conditions, etc. It can allow us to identify a set of important alternative ex-ante predictors of the daily volatility level. Second, it would be interesting to investigate how much we can learn from the short intra-daily periods about correlations, market betas and other relevant quantities. Thus, we will be able to explore an intra-daily discovery rate of a covariance between assets that is of considerable interest for risk management and hedging purposes. Third, in order to make more substantive empirical conclusions we would like to extend the number of considered assets (to provide more variation in stock characteristics across the sample) as well as the set of realized measures in use (to better examine the extent of heterogeneity in small sample performance across different estimators).

	9:30	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	0.000	0.177	0.250	0.309	0.352	0.389	0.422	0.455	0.485	0.510	0.534	0.556	0.576
	0.000	0.524	0.617	0.667	0.708	0.745	0.765	0.785	0.811	0.821	0.817	0.826	0.832
	0.613	0.755	0.762	0.769	0.777	0.794	0.802	0.811	0.828	0.835	0.833	0.839	0.843
BA	0.000	0.193	0.276	0.346	0.391	0.435	0.469	0.503	0.533	0.559	0.584	0.605	0.629
	0.000	0.425	0.559	0.656	0.699	0.757	0.778	0.797	0.822	0.837	0.853	0.863	0.894
	0.652	0.763	0.798	0.823	0.834	0.855	0.861	0.867	0.876	0.884	0.891	0.895	0.920
BAC	0.000	0.182	0.254	0.314	0.357	0.391	0.421	0.454	0.479	0.508	0.531	0.552	0.574
	0.000	0.631	0.735	0.818	0.839	0.844	0.851	0.868	0.879	0.874	0.880	0.885	0.890
	0.571	0.784	0.817	0.852	0.863	0.864	0.866	0.880	0.888	0.887	0.890	0.893	0.896
CSCO	0.000	0.158	0.233	0.297	0.345	0.389	0.425	0.458	0.486	0.513	0.537	0.561	0.585
	0.000	0.442	0.612	0.665	0.726	0.774	0.799	0.816	0.832	0.843	0.852	0.858	0.872
	0.616	0.707	0.752	0.761	0.792	0.822	0.838	0.844	0.853	0.862	0.867	0.875	0.885
CVX	0.000	0.118	0.177	0.234	0.278	0.328	0.364	0.397	0.428	0.456	0.482	0.505	0.529
	0.000	0.375	0.551	0.648	0.717	0.787	0.822	0.840	0.851	0.863	0.864	0.869	0.880
	0.691	0.751	0.776	0.794	0.816	0.854	0.866	0.873	0.877	0.886	0.886	0.889	0.897
GE	0.000	0.160	0.240	0.306	0.351	0.391	0.427	0.459	0.492	0.519	0.543	0.565	0.588
	0.000	0.612	0.668	0.724	0.750	0.778	0.804	0.820	0.834	0.842	0.850	0.859	0.870
	0.552	0.763	0.781	0.803	0.819	0.835	0.848	0.854	0.864	0.869	0.873	0.878	0.885
INTC	0.000	0.168	0.248	0.314	0.362	0.406	0.442	0.476	0.504	0.530	0.554	0.578	0.600
	0.000	0.516	0.663	0.711	0.752	0.779	0.807	0.826	0.834	0.845	0.853	0.865	0.873
	0.639	0.764	0.814	0.835	0.851	0.861	0.873	0.880	0.883	0.887	0.890	0.899	0.901
JPM	0.000	0.161	0.237	0.302	0.347	0.385	0.420	0.452	0.480	0.506	0.527	0.547	0.568
	0.000	0.616	0.719	0.778	0.800	0.829	0.848	0.858	0.861	0.868	0.870	0.873	0.876
	0.585	0.756	0.787	0.817	0.823	0.848	0.862	0.868	0.873	0.878	0.878	0.880	0.884
PFE	0.000	0.179	0.252	0.313	0.357	0.397	0.433	0.466	0.494	0.521	0.546	0.570	0.592
	0.000	0.392	0.526	0.585	0.614	0.653	0.694	0.726	0.737	0.747	0.759	0.770	0.784
	0.635	0.710	0.755	0.761	0.766	0.779	0.798	0.810	0.812	0.815	0.819	0.823	0.830
UTX	0.000	0.185	0.261	0.329	0.377	0.420	0.454	0.487	0.514	0.540	0.563	0.585	0.605
	0.000	0.460	0.619	0.696	0.726	0.770	0.795	0.824	0.839	0.856	0.867	0.876	0.883
	0.649	0.771	0.816	0.836	0.848	0.865	0.871	0.885	0.891	0.901	0.904	0.907	0.911
WMT	0.000	0.194	0.266	0.330	0.381	0.424	0.461	0.495	0.524	0.551	0.574	0.595	0.617
	0.000	0.492	0.634	0.686	0.728	0.780	0.802	0.825	0.839	0.851	0.857	0.863	0.876
	0.606	0.765	0.805	0.822	0.843	0.864	0.873	0.881	0.885	0.891	0.894	0.896	0.903
XOM	0.000	0.118	0.176	0.228	0.270	0.316	0.351	0.384	0.413	0.440	0.465	0.488	0.512
	0.000	0.482	0.620	0.705	0.730	0.793	0.815	0.832	0.844	0.852	0.860	0.864	0.873
	0.690	0.772	0.795	0.819	0.827	0.854	0.857	0.865	0.868	0.873	0.877	0.879	0.885
Average	0.000	0.166	0.239	0.302	0.347	0.389	0.424	0.457	0.486	0.513	0.537	0.559	0.581
	0.000	0.497	0.627	0.695	0.732	0.774	0.798	0.818	0.832	0.842	0.849	0.856	0.867
	0.625	0.755	0.788	0.808	0.822	0.841	0.851	0.860	0.867	0.872	0.875	0.879	0.887

Realized Variance Estimator.

	9:30	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	0.000	0.187	0.261	0.322	0.363	0.393	0.419	0.448	0.473	0.498	0.519	0.541	0.557
	0.000	0.365	0.461	0.544	0.594	0.638	0.716	0.758	0.825	0.836	0.841	0.853	0.863
	0.670	0.722	0.728	0.723	0.728	0.743	0.758	0.780	0.828	0.837	0.845	0.856	0.863
BA	0.000	0.200	0.276	0.341	0.383	0.422	0.452	0.485	0.511	0.536	0.558	0.580	0.598
	0.000	0.265	0.380	0.504	0.597	0.650	0.707	0.740	0.775	0.803	0.829	0.845	0.849
	0.729	0.758	0.787	0.805	0.822	0.832	0.849	0.857	0.868	0.880	0.888	0.895	0.898
BAC	0.000	0.193	0.272	0.324	0.362	0.392	0.417	0.447	0.470	0.493	0.513	0.532	0.554
	0.000	0.453	0.462	0.611	0.635	0.654	0.679	0.719	0.726	0.670	0.684	0.710	0.762
	0.599	0.710	0.663	0.714	0.731	0.739	0.753	0.780	0.784	0.759	0.766	0.770	0.799
CSCO	0.000	0.170	0.244	0.310	0.352	0.390	0.422	0.453	0.476	0.498	0.520	0.540	0.561
	0.000	0.289	0.464	0.541	0.648	0.734	0.772	0.799	0.827	0.846	0.841	0.862	0.879
	0.614	0.628	0.669	0.685	0.731	0.786	0.813	0.826	0.847	0.862	0.855	0.880	0.888
CVX	0.000	0.124	0.185	0.243	0.285	0.332	0.368	0.402	0.430	0.456	0.480	0.499	0.524
	0.000	0.298	0.511	0.600	0.698	0.786	0.824	0.841	0.850	0.868	0.841	0.853	0.881
	0.664	0.695	0.724	0.749	0.782	0.838	0.854	0.862	0.869	0.880	0.858	0.867	0.884
GE	0.000	0.164	0.241	0.303	0.342	0.376	0.409	0.440	0.469	0.496	0.519	0.542	0.565
	0.000	0.508	0.584	0.648	0.691	0.732	0.763	0.791	0.802	0.808	0.816	0.824	0.848
	0.582	0.739	0.749	0.768	0.793	0.823	0.835	0.852	0.858	0.858	0.861	0.864	0.882
INTC	0.000	0.178	0.261	0.325	0.365	0.406	0.436	0.466	0.491	0.515	0.536	0.556	0.577
	0.000	0.290	0.463	0.582	0.641	0.670	0.704	0.743	0.772	0.796	0.827	0.835	0.850
	0.678	0.694	0.731	0.762	0.782	0.791	0.810	0.826	0.838	0.852	0.869	0.875	0.880
JPM	0.000	0.170	0.254	0.311	0.353	0.389	0.421	0.453	0.480	0.504	0.521	0.540	0.559
	0.000	0.490	0.582	0.557	0.632	0.767	0.793	0.792	0.818	0.845	0.867	0.876	0.884
	0.633	0.702	0.729	0.702	0.709	0.793	0.813	0.799	0.821	0.845	0.867	0.875	0.886
PFE	0.000	0.177	0.244	0.304	0.343	0.382	0.417	0.447	0.475	0.500	0.524	0.547	0.569
	0.000	0.250	0.380	0.477	0.512	0.589	0.672	0.719	0.737	0.755	0.774	0.792	0.817
	0.696	0.713	0.723	0.734	0.748	0.768	0.801	0.819	0.827	0.834	0.844	0.852	0.866
UTX	0.000	0.189	0.259	0.320	0.363	0.406	0.432	0.459	0.485	0.509	0.530	0.550	0.569
	0.000	0.236	0.483	0.594	0.660	0.723	0.750	0.779	0.807	0.836	0.860	0.875	0.886
	0.671	0.677	0.720	0.751	0.778	0.805	0.827	0.846	0.857	0.876	0.887	0.894	0.903
WMT	0.000	0.198	0.267	0.320	0.372	0.405	0.435	0.467	0.494	0.518	0.541	0.561	0.582
	0.000	0.308	0.438	0.517	0.656	0.717	0.715	0.744	0.768	0.796	0.783	0.819	0.832
	0.658	0.711	0.737	0.759	0.785	0.807	0.817	0.826	0.833	0.846	0.835	0.854	0.860
XOM	0.000	0.127	0.187	0.239	0.278	0.318	0.355	0.386	0.416	0.442	0.465	0.487	0.509
	0.000	0.390	0.555	0.694	0.711	0.808	0.829	0.837	0.849	0.864	0.883	0.894	0.902
	0.696	0.735	0.759	0.811	0.807	0.864	0.870	0.877	0.878	0.885	0.888	0.895	0.902
Average	0.000	0.173	0.246	0.305	0.347	0.384	0.415	0.446	0.472	0.497	0.519	0.540	0.560
	0.000	0.345	0.480	0.572	0.640	0.706	0.744	0.772	0.796	0.810	0.821	0.836	0.854
	0.658	0.707	0.727	0.747	0.766	0.799	0.817	0.829	0.842	0.851	0.855	0.865	0.876

Threshold Variance Estimator.

Table 3.4: Table contains the estimates of intra-daily variance fractions and the projection results for open-to-close realized measures. Rows correspond to analyzed stocks, columns correspond to the intra-daily intervals. In particular, column with header A corresponds to the intra-daily interval [9:30a.m. - A]. An intersection of a row and a column is a cell and each cell contains 3 numbers. The top number is an average intra-daily variance fraction. The middle number is the Mincer-Zarnowitz R^2 for the regressions that forecast end-of-day realized measures from (3.5). The bottom number is the Mincer-Zarnowitz R^2 for the regressions that forecast end-of-day realized measures from (3.7). The out-of-sample period is January, 2007 - December, 2013.

	9:30	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	0.000	0.185	0.262	0.325	0.369	0.408	0.442	0.473	0.501	0.527	0.550	0.571	0.592
	0.000	0.650	0.742	0.787	0.810	0.833	0.856	0.871	0.883	0.891	0.892	0.896	0.902
	0.715	0.815	0.830	0.837	0.846	0.858	0.872	0.880	0.888	0.893	0.894	0.897	0.902
BA	0.000	0.190	0.274	0.345	0.392	0.435	0.470	0.503	0.532	0.559	0.583	0.605	0.627
	0.000	0.633	0.724	0.771	0.805	0.830	0.850	0.863	0.872	0.883	0.891	0.897	0.916
	0.727	0.819	0.840	0.852	0.867	0.880	0.888	0.895	0.900	0.906	0.910	0.914	0.928
BAC	0.000	0.167	0.246	0.311	0.355	0.393	0.425	0.457	0.484	0.511	0.533	0.555	0.575
	0.000	0.730	0.783	0.830	0.842	0.835	0.832	0.834	0.835	0.828	0.825	0.827	0.828
	0.537	0.757	0.788	0.829	0.839	0.831	0.829	0.832	0.833	0.827	0.824	0.825	0.826
CSCO	0.000	0.149	0.225	0.292	0.339	0.382	0.417	0.450	0.478	0.505	0.530	0.554	0.576
	0.000	0.659	0.761	0.796	0.825	0.851	0.870	0.884	0.891	0.895	0.900	0.902	0.911
	0.648	0.766	0.817	0.830	0.849	0.867	0.883	0.894	0.900	0.903	0.908	0.912	0.918
CVX	0.000	0.118	0.180	0.242	0.285	0.334	0.371	0.405	0.436	0.464	0.490	0.515	0.538
	0.000	0.589	0.713	0.772	0.812	0.859	0.885	0.903	0.904	0.910	0.916	0.920	0.926
	0.720	0.781	0.813	0.836	0.853	0.881	0.897	0.909	0.911	0.916	0.921	0.925	0.930
GE	0.000	0.147	0.224	0.294	0.343	0.384	0.420	0.453	0.484	0.512	0.535	0.556	0.577
	0.000	0.676	0.762	0.774	0.802	0.811	0.826	0.833	0.851	0.860	0.867	0.872	0.880
	0.625	0.769	0.804	0.807	0.824	0.833	0.846	0.852	0.867	0.874	0.878	0.882	0.888
INTC	0.000	0.159	0.238	0.309	0.359	0.402	0.437	0.471	0.500	0.527	0.551	0.574	0.596
	0.000	0.701	0.793	0.836	0.856	0.871	0.890	0.901	0.907	0.911	0.914	0.918	0.924
	0.710	0.812	0.850	0.874	0.891	0.900	0.912	0.921	0.925	0.929	0.930	0.933	0.937
JPM	0.000	0.158	0.238	0.306	0.353	0.394	0.429	0.461	0.489	0.515	0.537	0.558	0.580
	0.000	0.714	0.794	0.848	0.871	0.878	0.884	0.889	0.891	0.895	0.898	0.904	0.905
	0.650	0.751	0.800	0.846	0.868	0.879	0.887	0.892	0.895	0.899	0.901	0.906	0.907
PFE	0.000	0.173	0.247	0.310	0.356	0.395	0.431	0.463	0.492	0.520	0.544	0.568	0.589
	0.000	0.549	0.665	0.722	0.754	0.778	0.803	0.821	0.833	0.843	0.850	0.858	0.867
	0.716	0.785	0.823	0.834	0.845	0.853	0.864	0.873	0.878	0.882	0.886	0.889	0.894
UTX	0.000	0.176	0.252	0.321	0.367	0.409	0.443	0.476	0.504	0.531	0.555	0.578	0.598
	0.000	0.656	0.759	0.795	0.817	0.842	0.860	0.876	0.886	0.894	0.902	0.910	0.916
	0.693	0.805	0.841	0.849	0.863	0.876	0.886	0.897	0.901	0.906	0.910	0.915	0.920
WMT	0.000	0.175	0.251	0.318	0.367	0.411	0.448	0.481	0.510	0.537	0.561	0.584	0.606
	0.000	0.629	0.729	0.782	0.810	0.835	0.853	0.870	0.880	0.886	0.893	0.898	0.906
	0.698	0.793	0.824	0.841	0.858	0.871	0.881	0.891	0.898	0.902	0.906	0.909	0.915
XOM	0.000	0.111	0.171	0.230	0.271	0.317	0.352	0.386	0.415	0.443	0.468	0.493	0.516
	0.000	0.650	0.767	0.809	0.833	0.856	0.874	0.890	0.894	0.900	0.905	0.909	0.915
	0.735	0.791	0.829	0.847	0.858	0.873	0.887	0.898	0.900	0.905	0.909	0.913	0.917
Average	0.000	0.159	0.234	0.300	0.346	0.389	0.424	0.457	0.485	0.513	0.536	0.559	0.581
	0.000	0.653	0.749	0.793	0.820	0.840	0.857	0.869	0.877	0.883	0.888	0.893	0.899
	0.681	0.787	0.822	0.840	0.855	0.867	0.878	0.886	0.891	0.895	0.898	0.902	0.907

Realized Kernel Estimator.

	9:30	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	0.000	0.173	0.248	0.314	0.360	0.400	0.435	0.466	0.496	0.521	0.545	0.568	0.589
	0.000	0.675	0.770	0.820	0.839	0.865	0.885	0.901	0.912	0.920	0.922	0.927	0.934
	0.817	0.843	0.858	0.868	0.874	0.886	0.898	0.908	0.916	0.921	0.923	0.927	0.934
BA	0.000	0.162	0.246	0.317	0.365	0.408	0.445	0.477	0.507	0.534	0.559	0.581	0.603
	0.000	0.569	0.712	0.786	0.832	0.858	0.875	0.887	0.896	0.905	0.913	0.918	0.929
	0.763	0.807	0.840	0.861	0.878	0.893	0.902	0.910	0.916	0.922	0.927	0.930	0.939
BAC	0.000	0.154	0.228	0.291	0.337	0.375	0.408	0.439	0.468	0.494	0.517	0.540	0.561
	0.000	0.674	0.786	0.852	0.869	0.862	0.865	0.863	0.864	0.868	0.869	0.871	0.873
	0.595	0.737	0.803	0.851	0.865	0.857	0.861	0.860	0.861	0.865	0.866	0.868	0.870
CSCO	0.000	0.151	0.224	0.290	0.335	0.376	0.412	0.443	0.472	0.498	0.522	0.545	0.568
	0.000	0.649	0.738	0.780	0.810	0.838	0.864	0.878	0.886	0.891	0.898	0.902	0.912
	0.691	0.789	0.818	0.830	0.846	0.863	0.882	0.892	0.898	0.901	0.907	0.911	0.919
CVX	0.000	0.098	0.161	0.224	0.270	0.319	0.358	0.392	0.423	0.452	0.479	0.505	0.530
	0.000	0.680	0.802	0.832	0.862	0.888	0.906	0.914	0.920	0.926	0.929	0.933	0.939
	0.769	0.812	0.850	0.861	0.875	0.894	0.906	0.913	0.918	0.922	0.925	0.929	0.935
GE	0.000	0.142	0.216	0.285	0.334	0.378	0.414	0.445	0.476	0.504	0.527	0.549	0.570
	0.000	0.735	0.797	0.819	0.829	0.827	0.847	0.853	0.868	0.877	0.885	0.897	0.907
	0.640	0.786	0.832	0.851	0.857	0.858	0.874	0.878	0.891	0.898	0.902	0.910	0.917
INTC	0.000	0.158	0.235	0.302	0.352	0.392	0.428	0.461	0.490	0.517	0.540	0.563	0.585
	0.000	0.683	0.782	0.834	0.855	0.871	0.888	0.900	0.907	0.914	0.917	0.924	0.928
	0.724	0.802	0.838	0.863	0.879	0.889	0.901	0.911	0.916	0.921	0.923	0.929	0.933
JPM	0.000	0.144	0.220	0.287	0.335	0.377	0.413	0.445	0.473	0.500	0.524	0.547	0.568
	0.000	0.781	0.862	0.894	0.914	0.918	0.924	0.931	0.936	0.938	0.942	0.946	0.947
	0.732	0.812	0.862	0.893	0.914	0.923	0.929	0.936	0.941	0.944	0.946	0.950	0.952
PFE	0.000	0.178	0.251	0.313	0.359	0.398	0.432	0.464	0.491	0.518	0.541	0.564	0.585
	0.000	0.575	0.696	0.744	0.766	0.784	0.799	0.820	0.831	0.841	0.849	0.857	0.865
	0.731	0.794	0.821	0.833	0.843	0.849	0.856	0.866	0.871	0.875	0.879	0.883	0.888
UTX	0.000	0.152	0.228	0.295	0.342	0.384	0.421	0.453	0.481	0.508	0.533	0.556	0.577
	0.000	0.574	0.774	0.821	0.847	0.863	0.882	0.891	0.901	0.907	0.914	0.920	0.925
	0.723	0.790	0.859	0.872	0.884	0.893	0.904	0.910	0.915	0.917	0.922	0.926	0.930
WMT	0.000	0.164	0.241	0.309	0.358	0.401	0.439	0.472	0.500	0.527	0.552	0.575	0.596
	0.000	0.589	0.708	0.774	0.812	0.834	0.855	0.868	0.879	0.888	0.897	0.902	0.909
	0.741	0.790	0.820	0.843	0.862	0.873	0.885	0.893	0.899	0.904	0.910	0.913	0.918
XOM	0.000	0.099	0.158	0.218	0.261	0.309	0.346	0.380	0.409	0.437	0.464	0.490	0.514
	0.000	0.718	0.818	0.851	0.870	0.889	0.902	0.911	0.915	0.920	0.924	0.928	0.934
	0.771	0.823	0.853	0.873	0.885	0.899	0.910	0.916	0.919	0.923	0.926	0.930	0.935
Average	0.000	0.148	0.221	0.287	0.334	0.376	0.412	0.445	0.474	0.501	0.525	0.549	0.571
	0.000	0.659	0.770	0.817	0.842	0.858	0.874	0.885	0.893	0.899	0.905	0.910	0.917
	0.725	0.799	0.838	0.858	0.872	0.881	0.892	0.899	0.905	0.909	0.913	0.917	0.922

Markov Chain Estimator.

Table 3.5: Table contains the estimates of intra-daily variance fractions and the projection results for open-to-close realized measures. Rows correspond to analyzed stocks, columns correspond to the intra-daily intervals. In particular, column with header A corresponds to the intra-daily interval [9:30a.m. - A]. An intersection of a row and a column is a cell and each cell contains 3 numbers. The top number is an average intra-daily variance fraction. The middle number is the Mincer-Zarnowitz R^2 for the regressions that forecast end-of-day realized measures from (3.5). The bottom number is the Mincer-Zarnowitz R^2 for the regressions that forecast end-of-day realized measures from (3.7). The out-of-sample period is January, 2007 - December, 2013.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30	12:45	1:00	1:15	1:30	1:45	2:00	2:15	2:30	2:45	3:00	3:15	3:30	3:45	4:00
AXP	*-2.213*	0.154	0.219	0.238	0.260	0.279	0.287	0.296	0.298	0.297	0.301	0.304	0.309	0.312	0.315	0.318	0.321	0.323	0.327	0.332	0.333	0.336	0.338	0.339	0.341	0.345
	-1.930	0.002	0.004	0.006	0.008	0.010	0.012	0.014	0.014	0.016	0.017	0.018	0.020	0.021	0.023	0.024	0.029	0.031	0.034	0.039	0.036	0.040	0.043	0.043	0.046	0.048
BA	*-2.029*	0.195	0.246	0.271	0.278	0.292	0.298	0.304	0.312	0.319	0.326	0.334	0.336	0.337	0.338	0.339	0.339	0.339	0.342	0.345	0.349	0.351	0.352	0.352	0.354	0.356
	-1.729	0.004	0.003	0.010	0.010	0.012	0.016	0.018	0.021	0.024	0.028	0.036	0.036	0.038	0.038	0.040	0.040	0.040	0.042	0.044	0.047	0.048	0.049	0.050	0.051	0.052
BAC	*-2.366*	0.152	0.197	0.211	0.221	0.227	0.237	0.246	0.249	0.249	0.251	0.261	0.264	0.266	0.268	0.272	0.278	0.279	0.284	0.291	0.293	0.293	0.297	0.301	0.304	0.306
	-2.157	0.009	0.015	0.016	0.011	0.020	0.023	0.025	0.025	0.028	0.029	0.034	0.035	0.040	0.040	0.049	0.053	0.051	0.059	0.066	0.052	0.067	0.058	0.067	0.074	0.076
CSCO	*-1.940*	0.124	0.181	0.206	0.213	0.219	0.223	0.228	0.230	0.234	0.238	0.242	0.245	0.249	0.251	0.253	0.255	0.255	0.256	0.259	0.263	0.265	0.268	0.269	0.272	0.275
	-1.719	0.004	0.004	0.008	0.009	0.010	0.011	0.009	0.003	0.005	0.009	0.010	0.011	0.012	0.020	0.015	0.015	0.013	0.016	0.022	0.024	0.025	0.027	0.027	0.030	0.032
CVX	*-1.873*	0.191	0.242	0.274	0.294	0.307	0.315	0.318	0.320	0.322	0.322	0.323	0.325	0.327	0.328	0.329	0.332	0.334	0.338	0.344	0.347	0.349	0.352	0.354	0.360	0.368
	-1.564	0.003	0.005	0.008	0.009	0.010	0.012	0.012	0.013	0.012	0.012	0.013	0.012	0.010	0.008	0.009	0.009	0.015	0.011	0.020	0.022	0.022	0.019	0.022	0.034	0.050
GE	*-2.027*	0.188	0.254	0.274	0.294	0.307	0.321	0.331	0.337	0.348	0.353	0.358	0.361	0.363	0.364	0.366	0.370	0.372	0.379	0.384	0.387	0.388	0.390	0.393	0.395	0.397
	-1.700	0.002	-0.003	0.005	0.003	0.012	0.016	0.018	0.020	0.023	0.024	0.025	0.032	0.029	0.033	0.034	0.035	0.035	0.039	0.042	0.045	0.045	0.045	0.048	0.049	0.052
INTC	*-2.047*	0.184	0.242	0.256	0.274	0.290	0.299	0.303	0.309	0.314	0.317	0.320	0.323	0.325	0.327	0.328	0.331	0.333	0.334	0.339	0.342	0.344	0.346	0.347	0.349	0.352
	-1.742	0.004	0.004	0.004	0.004	0.007	0.008	0.009	0.011	0.013	0.014	0.015	0.016	0.018	0.019	0.019	0.021	0.023	0.023	0.026	0.029	0.031	0.032	0.034	0.035	0.040
JPM	*-2.204*	0.143	0.189	0.196	0.206	0.216	0.222	0.225	0.226	0.227	0.230	0.234	0.238	0.241	0.244	0.244	0.249	0.251	0.257	0.264	0.268	0.270	0.271	0.277	0.282	0.288
	-1.988	0.001	0.005	0.004	0.005	0.004	0.017	0.018	0.018	0.019	0.020	0.022	0.024	0.026	0.027	0.027	0.030	0.030	0.034	0.036	0.035	0.044	0.043	0.047	0.052	0.058
PFE	*-1.818*	0.112	0.212	0.232	0.258	0.268	0.277	0.284	0.288	0.293	0.297	0.299	0.301	0.304	0.306	0.307	0.309	0.310	0.314	0.317	0.318	0.318	0.323	0.325	0.326	0.331
	-1.542	0.002	0.002	0.002	0.006	0.010	0.014	0.015	0.011	0.011	0.017	0.016	0.009	0.018	0.021	0.009	0.024	0.025	0.019	0.026	0.022	0.025	0.028	0.031	0.029	0.027
UTX	*-1.850*	0.135	0.225	0.249	0.261	0.273	0.286	0.292	0.297	0.300	0.303	0.304	0.307	0.309	0.310	0.312	0.313	0.315	0.318	0.324	0.326	0.329	0.331	0.333	0.336	0.340
	-1.553	0.005	0.004	0.006	0.006	0.009	0.010	0.011	0.012	0.012	0.012	0.012	0.012	0.014	0.012	0.014	0.014	0.015	0.015	0.019	0.020	0.022	0.022	0.021	0.023	0.029
WMT	*-1.638*	0.218	0.234	0.276	0.290	0.306	0.325	0.331	0.334	0.339	0.343	0.345	0.348	0.351	0.352	0.353	0.356	0.358	0.361	0.364	0.367	0.369	0.371	0.372	0.374	0.377
	-1.330	0.005	0.011	0.019	0.018	0.020	0.028	0.029	0.030	0.029	0.029	0.030	0.032	0.032	0.032	0.033	0.034	0.035	0.038	0.040	0.042	0.046	0.048	0.048	0.047	0.052
XOM	*-1.747*	0.171	0.217	0.235	0.254	0.260	0.267	0.272	0.277	0.280	0.281	0.276	0.280	0.283	0.287	0.289	0.292	0.293	0.296	0.300	0.304	0.306	0.311	0.313	0.315	0.322
	-1.478	0.004	0.003	0.004	0.006	0.006	0.006	0.007	0.008	0.002	0.008	0.010	0.011	0.013	0.015	0.016	0.018	0.017	0.019	0.022	0.025	0.025	0.030	0.031	0.041	0.047
Average	*-1.979*	0.164	0.221	0.243	0.258	0.270	0.280	0.286	0.290	0.294	0.297	0.300	0.303	0.305	0.307	0.309	0.312	0.314	0.317	0.322	0.325	0.327	0.329	0.331	0.334	0.338
	-1.703	0.004	0.005	0.008	0.008	0.011	0.014	0.016	0.016	0.016	0.018	0.020	0.021	0.023	0.024	0.024	0.027	0.028	0.029	0.033	0.033	0.037	0.037	0.039	0.043	0.047

Table 3.6: **Realized Variance Estimator.** Table contains the forecasting results of the daily return density. Analysis was made using the realized variance (RV) estimator. Rows correspond to analyzed stocks, columns correspond to the intra-daily intervals. In particular, column with header A corresponds to the intra-daily interval [9:30a.m. - A]. The numbers in the first column indicate the mean log-predictive likelihood (LPL) values, whereas the numbers in other column are changes in LPL relative to the first column. An intersection of a row and a column is a cell and each cell contains 2 numbers. The top numbers correspond to the variance specification given by (3.8). The bottom numbers correspond to the variance specification given by (3.9). The out-of-sample period is January, 2007 - December, 2013.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30	12:45	1:00	1:15	1:30	1:45	2:00	2:15	2:30	2:45	3:00	3:15	3:30	3:45	4:00
AXP	*-2.317*	0.189	0.305	0.342	0.347	0.374	0.381	0.390	0.382	0.398	0.396	0.402	0.410	0.413	0.415	0.417	0.420	0.423	0.426	0.431	0.432	0.435	0.437	0.441	0.442	0.445
	-1.925	-0.003	-0.000	0.004	0.008	0.010	0.010	0.010	0.010	0.012	0.013	0.015	0.019	0.020	0.021	0.022	0.023	0.025	0.029	0.032	0.034	0.035	0.037	0.039	0.040	0.042
BA	*-2.118*	0.200	0.304	0.326	0.325	0.352	0.363	0.372	0.383	0.390	0.394	0.402	0.411	0.413	0.416	0.416	0.421	0.421	0.423	0.428	0.429	0.431	0.432	0.433	0.434	0.436
	-1.728	0.002	-0.002	0.002	-0.001	0.002	0.005	0.006	0.011	0.014	0.016	0.020	0.026	0.027	0.028	0.029	0.033	0.033	0.034	0.038	0.039	0.040	0.040	0.041	0.040	0.042
BAC	*-2.425*	0.182	0.231	0.250	0.255	0.256	0.273	0.278	0.283	0.286	0.293	0.300	0.307	0.311	0.310	0.318	0.322	0.325	0.326	0.333	0.335	0.337	0.342	0.346	0.349	0.353
	-2.165	0.003	-0.000	0.002	0.002	-0.000	0.011	0.017	0.011	0.003	0.022	0.013	0.018	0.015	0.006	0.008	0.010	0.017	0.020	0.019	0.015	0.024	0.037	0.043	0.044	0.054
CSCO	*-1.990*	0.134	0.212	0.228	0.233	0.245	0.251	0.253	0.262	0.264	0.272	0.276	0.281	0.287	0.290	0.293	0.291	0.291	0.292	0.294	0.298	0.303	0.305	0.305	0.309	0.312
	-1.721	0.003	0.006	0.005	0.003	0.006	0.005	0.003	0.008	0.009	0.003	0.007	0.012	0.015	0.016	0.017	0.015	0.014	0.013	0.016	0.018	0.020	0.022	0.022	0.026	0.026
CVX	*-1.941*	0.223	0.279	0.314	0.348	0.364	0.376	0.379	0.380	0.385	0.386	0.388	0.389	0.392	0.391	0.393	0.395	0.399	0.402	0.405	0.407	0.410	0.410	0.413	0.417	0.424
	-1.567	0.010	0.010	0.011	0.013	0.015	0.018	0.017	0.019	0.017	0.017	0.015	0.011	0.012	-0.003	0.003	0.008	0.011	0.013	0.024	0.026	0.027	0.028	0.031	0.025	0.039
GE	*-2.080*	0.210	0.287	0.321	0.337	0.351	0.361	0.371	0.381	0.389	0.395	0.404	0.408	0.408	0.408	0.407	0.411	0.414	0.416	0.422	0.426	0.430	0.432	0.434	0.436	0.437
	-1.710	0.004	0.009	0.013	0.013	0.015	0.017	0.020	0.022	0.024	0.027	0.031	0.037	0.036	0.035	0.036	0.036	0.037	0.039	0.041	0.044	0.046	0.048	0.048	0.049	0.051
INTC	*-2.165*	0.254	0.327	0.348	0.371	0.388	0.401	0.405	0.409	0.412	0.413	0.416	0.421	0.425	0.428	0.428	0.433	0.433	0.436	0.444	0.448	0.451	0.454	0.456	0.458	0.463
	-1.739	0.003	0.002	-0.001	-0.003	0.001	0.004	0.004	0.004	0.004	0.005	0.006	0.007	0.009	0.009	0.009	0.011	0.011	0.013	0.016	0.018	0.021	0.022	0.026	0.026	0.029
JPM	*-2.229*	0.138	0.164	0.186	0.203	0.223	0.227	0.232	0.233	0.230	0.233	0.246	0.248	0.253	0.257	0.259	0.265	0.269	0.274	0.268	0.273	0.276	0.280	0.287	0.295	0.301
	-1.984	0.008	-0.005	0.001	0.001	-0.008	0.004	0.005	-0.010	-0.004	-0.004	0.001	0.004	0.007	0.010	0.008	0.021	0.010	0.021	0.017	0.019	0.020	0.023	0.030	0.036	0.042
PFE	*-1.889*	0.171	0.239	0.267	0.305	0.322	0.336	0.342	0.352	0.353	0.356	0.361	0.361	0.361	0.365	0.366	0.370	0.372	0.373	0.375	0.377	0.377	0.379	0.382	0.384	0.386
	-1.558	0.013	0.005	0.016	0.020	0.024	0.027	0.026	0.027	0.022	0.021	0.015	0.021	0.020	0.030	0.029	0.015	0.026	0.034	0.029	0.020	0.029	0.034	0.031	0.037	0.042
UTX	*-1.900*	0.130	0.240	0.285	0.299	0.315	0.323	0.332	0.337	0.343	0.348	0.351	0.352	0.354	0.357	0.360	0.360	0.363	0.363	0.369	0.371	0.373	0.376	0.378	0.380	0.385
	-1.551	-0.003	-0.003	-0.003	-0.004	0.002	0.003	-0.002	0.001	-0.002	0.000	0.005	0.006	0.002	0.007	0.012	0.009	0.012	0.008	0.010	0.018	0.011	0.011	0.014	0.008	0.025
WMT	*-1.667*	0.210	0.225	0.293	0.307	0.324	0.338	0.343	0.349	0.355	0.358	0.359	0.364	0.366	0.364	0.365	0.367	0.370	0.375	0.379	0.380	0.382	0.386	0.387	0.389	0.390
	-1.321	-0.001	0.001	0.012	0.015	0.012	0.015	0.016	0.020	0.020	0.016	0.007	0.010	0.011	0.009	0.010	0.003	0.014	0.013	0.014	0.013	0.014	0.008	0.005	0.022	0.019
XOM	*-1.808*	0.196	0.259	0.284	0.301	0.305	0.309	0.322	0.330	0.331	0.332	0.328	0.334	0.337	0.341	0.346	0.347	0.349	0.351	0.356	0.360	0.361	0.364	0.367	0.370	0.373
	-1.475	0.000	-0.004	0.000	0.002	0.003	0.003	0.005	0.005	0.007	0.007	0.010	0.010	0.011	0.011	0.014	0.014	0.014	0.017	0.021	0.023	0.023	0.026	0.027	0.029	0.034
Average	*-2.044*	0.186	0.256	0.287	0.303	0.318	0.328	0.335	0.340	0.345	0.348	0.353	0.357	0.360	0.362	0.364	0.367	0.369	0.371	0.375	0.378	0.381	0.383	0.386	0.389	0.392
	-1.704	0.003	0.001	0.005	0.006	0.007	0.010	0.011	0.011	0.011	0.012	0.012	0.015	0.015	0.015	0.016	0.016	0.019	0.021	0.023	0.024	0.026	0.028	0.030	0.032	0.037

Table 3.7: **Threshold Variance Estimator.** Table contains the forecasting results of the daily return density. Analysis was made using the threshold variance (TV) estimator. Rows correspond to analyzed stocks, columns correspond to the intra-daily intervals. In particular, column with header A corresponds to the intra-daily interval [9:30a.m. - A]. The numbers in the first column indicate the mean log-predictive likelihood (LPL) values, whereas the numbers in other column are changes in LPL relative to the first column. An intersection of a row and a column is a cell and each cell contains 2 numbers. The top numbers correspond to the variance specification given by (3.8). The bottom numbers correspond to the variance specification given by (3.9). The out-of-sample period is January, 2007 - December, 2013.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30	12:45	1:00	1:15	1:30	1:45	2:00	2:15	2:30	2:45	3:00	3:15	3:30	3:45	4:00
AXP	*-1.997*	0.035	0.054	0.050	0.062	0.067	0.072	0.076	0.078	0.081	0.083	0.085	0.089	0.091	0.094	0.097	0.098	0.100	0.103	0.106	0.108	0.109	0.111	0.112	0.115	0.117
	-1.934	0.001	0.004	0.006	0.010	0.012	0.014	0.016	0.016	0.017	0.018	0.019	0.022	0.023	0.025	0.026	0.028	0.032	0.035	0.037	0.038	0.038	0.040	0.041	0.043	0.044
BA	*-1.811*	0.043	0.060	0.073	0.080	0.083	0.092	0.095	0.099	0.102	0.094	0.106	0.108	0.109	0.110	0.110	0.111	0.112	0.113	0.116	0.119	0.120	0.121	0.122	0.123	0.124
	-1.739	0.001	0.004	0.009	0.021	0.022	0.026	0.028	0.030	0.032	0.023	0.034	0.035	0.037	0.037	0.038	0.039	0.040	0.041	0.044	0.047	0.048	0.049	0.050	0.050	0.051
BAC	*-2.166*	0.019	0.024	0.026	0.031	0.033	0.037	0.040	0.042	0.044	0.044	0.049	0.052	0.054	0.055	0.059	0.062	0.065	0.068	0.075	0.076	0.078	0.082	0.086	0.089	0.091
	-2.159	0.004	0.005	0.006	0.014	0.015	0.017	0.026	0.027	0.029	0.028	0.033	0.029	0.043	0.045	0.044	0.044	0.044	0.035	0.023	0.037	0.035	0.024	0.045	0.049	0.038
CSCO	*-1.786*	0.034	0.046	0.056	0.066	0.069	0.071	0.075	0.076	0.079	0.081	0.082	0.085	0.088	0.089	0.092	0.093	0.093	0.096	0.099	0.102	0.105	0.107	0.108	0.110	0.113
	-1.735	0.006	0.011	0.012	0.010	0.020	0.021	0.021	0.022	0.023	0.024	0.024	0.025	0.027	0.026	0.027	0.027	0.027	0.026	0.030	0.033	0.035	0.036	0.035	0.039	0.038
CVX	*-1.627*	0.043	0.045	0.052	0.062	0.068	0.072	0.074	0.075	0.076	0.077	0.078	0.081	0.081	0.082	0.083	0.085	0.087	0.091	0.094	0.096	0.097	0.099	0.101	0.104	0.108
	-1.563	0.007	0.006	0.006	0.008	0.008	0.011	0.011	0.011	0.010	0.008	0.007	0.007	0.007	0.007	0.007	0.007	0.005	0.018	0.018	0.018	0.019	0.021	0.022	0.026	0.033
GE	*-1.769*	0.038	0.051	0.063	0.070	0.077	0.082	0.087	0.090	0.095	0.098	0.102	0.105	0.106	0.107	0.109	0.111	0.113	0.119	0.124	0.125	0.127	0.129	0.131	0.132	0.135
	-1.704	0.007	0.008	0.014	0.015	0.018	0.020	0.023	0.025	0.027	0.030	0.033	0.035	0.036	0.037	0.038	0.039	0.039	0.043	0.047	0.049	0.050	0.052	0.053	0.054	0.056
INTC	*-1.771*	0.019	0.022	0.025	0.026	0.031	0.035	0.036	0.039	0.043	0.045	0.047	0.048	0.051	0.052	0.053	0.055	0.057	0.058	0.063	0.064	0.067	0.068	0.069	0.071	0.074
	-1.738	0.003	0.002	0.002	0.001	0.005	0.006	0.006	0.007	0.009	0.010	0.010	0.011	0.013	0.014	0.015	0.016	0.017	0.019	0.021	0.024	0.025	0.027	0.028	0.030	0.035
JPM	*-2.013*	0.021	0.026	0.026	0.030	0.033	0.034	0.035	0.036	0.036	0.037	0.039	0.042	0.043	0.044	0.045	0.047	0.049	0.053	0.058	0.061	0.063	0.070	0.074	0.078	0.082
	-1.998	0.004	0.010	0.004	0.016	0.005	0.017	0.016	0.020	0.019	0.023	0.023	0.014	0.018	0.018	0.030	0.031	0.031	0.036	0.028	0.032	0.033	0.043	0.047	0.050	0.056
PFE	*-1.589*	0.028	0.044	0.052	0.058	0.061	0.066	0.068	0.070	0.072	0.074	0.076	0.076	0.077	0.078	0.079	0.081	0.083	0.085	0.086	0.087	0.087	0.089	0.091	0.092	0.095
	-1.536	0.001	0.005	0.005	0.006	0.007	0.008	0.007	0.010	0.011	0.010	0.011	0.011	0.012	0.012	0.015	0.014	0.014	0.015	0.010	0.014	0.010	0.019	0.021	0.021	0.024
UTX	*-1.618*	0.016	0.035	0.046	0.052	0.059	0.064	0.067	0.069	0.073	0.075	0.076	0.079	0.080	0.080	0.082	0.082	0.084	0.086	0.092	0.094	0.096	0.098	0.100	0.102	0.104
	-1.556	-0.010	-0.006	0.002	0.000	0.001	0.001	0.002	-0.002	0.001	0.008	-0.000	0.007	0.010	0.009	0.010	0.009	0.011	0.010	0.012	0.009	0.011	0.015	0.017	0.027	0.026
WMT	*-1.410*	0.067	0.079	0.092	0.096	0.100	0.109	0.112	0.114	0.115	0.116	0.117	0.119	0.120	0.120	0.121	0.123	0.124	0.126	0.129	0.131	0.132	0.134	0.135	0.136	0.137
	-1.335	-0.002	0.013	0.017	0.019	0.020	0.026	0.025	0.022	0.020	0.031	0.028	0.030	0.030	0.030	0.031	0.033	0.034	0.038	0.040	0.042	0.043	0.043	0.044	0.045	0.044
XOM	*-1.537*	0.026	0.033	0.041	0.049	0.052	0.056	0.058	0.062	0.064	0.064	0.068	0.070	0.071	0.073	0.075	0.077	0.079	0.082	0.085	0.086	0.087	0.091	0.093	0.095	0.099
	-1.479	0.003	0.004	0.005	0.007	0.007	0.007	0.006	0.008	0.008	0.009	0.010	0.011	0.006	0.015	0.017	0.017	0.018	0.021	0.021	0.022	0.022	0.025	0.026	0.028	0.028
Average	*-1.758*	0.032	0.043	0.050	0.057	0.061	0.066	0.069	0.071	0.073	0.074	0.077	0.079	0.081	0.082	0.084	0.086	0.087	0.090	0.094	0.096	0.097	0.100	0.102	0.104	0.107
	-1.706	0.002	0.005	0.007	0.011	0.012	0.014	0.016	0.016	0.017	0.019	0.019	0.020	0.022	0.023	0.025	0.026	0.026	0.028	0.028	0.030	0.031	0.033	0.036	0.039	0.040

Table 3.8: **Realized Kernel Estimator.** Table contains the forecasting results of the daily return density. Analysis was made using the realized kernel (RK) estimator. Rows correspond to analyzed stocks, columns correspond to the intra-daily intervals. In particular, column with header A corresponds to the intra-daily interval [9:30a.m. - A]. The numbers in the first column indicate the mean log-predictive likelihood (LPL) values, whereas the numbers in other column are changes in LPL relative to the first column. An intersection of a row and a column is a cell and each cell contains 2 numbers. The top numbers correspond to the variance specification given by (3.8). The bottom numbers correspond to the variance specification given by (3.9). The out-of-sample period is January, 2007 - December, 2013.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30	12:45	1:00	1:15	1:30	1:45	2:00	2:15	2:30	2:45	3:00	3:15	3:30	3:45	4:00
AXP	*-1.998*	0.044	0.046	0.051	0.058	0.063	0.066	0.069	0.070	0.073	0.075	0.077	0.080	0.082	0.085	0.087	0.088	0.090	0.094	0.097	0.098	0.099	0.100	0.101	0.103	0.105
	-1.936	0.003	0.004	0.006	0.008	0.009	0.011	0.012	0.012	0.013	0.015	0.016	0.018	0.020	0.022	0.024	0.025	0.028	0.030	0.032	0.032	0.033	0.035	0.035	0.036	0.037
BA	*-1.850*	0.083	0.094	0.106	0.113	0.117	0.123	0.125	0.128	0.130	0.121	0.133	0.134	0.135	0.136	0.137	0.138	0.139	0.141	0.143	0.146	0.146	0.148	0.148	0.149	0.149
	-1.751	0.005	0.007	0.010	0.012	0.013	0.027	0.028	0.030	0.031	0.023	0.033	0.035	0.036	0.037	0.037	0.039	0.039	0.041	0.043	0.045	0.046	0.047	0.047	0.047	0.048
BAC	*-2.184*	0.023	0.029	0.034	0.039	0.043	0.047	0.052	0.054	0.055	0.057	0.061	0.064	0.066	0.068	0.071	0.074	0.077	0.080	0.087	0.089	0.090	0.093	0.096	0.099	0.101
	-2.167	0.002	0.007	0.004	0.007	0.006	0.007	0.019	0.015	0.021	0.033	0.032	0.040	0.036	0.040	0.037	0.031	0.045	0.043	0.045	0.044	0.029	0.032	0.034	0.033	0.024
CSCO	*-1.782*	0.027	0.037	0.049	0.054	0.058	0.061	0.065	0.067	0.069	0.070	0.071	0.074	0.076	0.078	0.080	0.081	0.082	0.085	0.087	0.090	0.093	0.095	0.096	0.098	0.099
	-1.735	-0.003	0.007	0.008	0.008	0.012	0.015	0.016	0.016	0.017	0.018	0.018	0.018	0.010	0.019	0.013	0.014	0.016	0.020	0.019	0.023	0.026	0.028	0.030	0.032	0.033
CVX	*-1.636*	0.029	0.039	0.047	0.055	0.063	0.067	0.069	0.072	0.074	0.075	0.077	0.080	0.081	0.082	0.083	0.085	0.087	0.090	0.093	0.095	0.096	0.098	0.099	0.101	0.103
	-1.615	0.044	0.051	0.053	0.054	0.056	0.056	0.056	0.058	0.058	0.059	0.059	0.060	0.059	0.058	0.055	0.061	0.063	0.065	0.066	0.067	0.068	0.070	0.071	0.073	0.075
GE	*-1.758*	0.041	0.047	0.053	0.057	0.063	0.067	0.071	0.073	0.077	0.080	0.085	0.087	0.089	0.090	0.091	0.093	0.096	0.100	0.105	0.107	0.108	0.110	0.111	0.112	0.113
	-1.706	0.006	0.007	0.009	0.011	0.013	0.014	0.016	0.018	0.020	0.021	0.024	0.026	0.027	0.028	0.029	0.029	0.031	0.034	0.036	0.038	0.039	0.040	0.041	0.041	0.043
INTC	*-1.758*	0.009	0.011	0.013	0.015	0.018	0.021	0.022	0.025	0.028	0.031	0.033	0.034	0.036	0.037	0.038	0.039	0.041	0.043	0.047	0.049	0.051	0.052	0.054	0.056	0.058
	-1.741	0.003	0.002	0.002	0.002	0.005	0.006	0.005	0.007	0.009	0.011	0.012	0.014	0.015	0.016	0.017	0.018	0.019	0.019	0.021	0.022	0.023	0.025	0.027	0.029	0.033
JPM	*-2.020*	0.016	0.023	0.024	0.029	0.031	0.033	0.034	0.035	0.035	0.036	0.038	0.042	0.043	0.044	0.045	0.048	0.049	0.053	0.058	0.060	0.062	0.067	0.072	0.075	0.079
	-2.008	0.003	0.012	0.010	0.013	0.014	0.015	0.015	0.018	0.019	0.021	0.023	0.025	0.026	0.028	0.028	0.029	0.029	0.025	0.032	0.033	0.033	0.037	0.045	0.044	0.046
PFE	*-1.598*	0.029	0.043	0.047	0.054	0.059	0.062	0.064	0.066	0.068	0.069	0.071	0.072	0.073	0.073	0.075	0.077	0.079	0.082	0.083	0.083	0.084	0.085	0.086	0.087	0.088
	-1.547	0.003	0.007	0.007	0.008	0.010	0.012	0.013	0.015	0.016	0.016	0.018	0.019	0.020	0.020	0.022	0.023	0.025	0.027	0.027	0.027	0.028	0.028	0.028	0.028	0.031
UTX	*-1.651*	0.048	0.065	0.066	0.076	0.082	0.085	0.087	0.091	0.094	0.095	0.097	0.099	0.101	0.102	0.103	0.105	0.105	0.108	0.115	0.117	0.118	0.120	0.121	0.123	0.125
	-1.559	-0.003	-0.001	-0.001	-0.000	-0.000	0.001	0.001	0.001	0.003	0.002	0.001	0.005	0.005	0.002	0.004	0.003	0.004	0.005	0.004	0.006	0.004	0.008	0.008	0.012	0.008
WMT	*-1.397*	0.036	0.051	0.063	0.069	0.073	0.081	0.084	0.086	0.087	0.088	0.089	0.090	0.091	0.092	0.092	0.094	0.096	0.099	0.102	0.103	0.104	0.105	0.106	0.107	0.107
	-1.353	0.013	0.026	0.028	0.032	0.033	0.037	0.039	0.041	0.041	0.041	0.041	0.042	0.042	0.043	0.043	0.045	0.047	0.050	0.051	0.052	0.053	0.054	0.055	0.054	0.053
XOM	*-1.539*	0.022	0.025	0.031	0.037	0.043	0.046	0.049	0.053	0.056	0.057	0.060	0.063	0.064	0.066	0.068	0.069	0.071	0.074	0.077	0.079	0.080	0.082	0.084	0.086	0.088
	-1.487	-0.003	-0.002	-0.000	0.003	0.004	0.005	0.006	0.008	0.009	0.010	0.011	0.012	0.012	0.012	0.010	0.019	0.020	0.023	0.024	0.025	0.026	0.028	0.030	0.032	0.034
Average	*-1.764*	0.034	0.042	0.049	0.055	0.059	0.063	0.066	0.068	0.070	0.071	0.074	0.077	0.078	0.079	0.081	0.083	0.084	0.087	0.091	0.093	0.094	0.096	0.098	0.100	0.101
	-1.717	0.006	0.011	0.011	0.013	0.015	0.017	0.019	0.020	0.021	0.022	0.024	0.026	0.026	0.027	0.027	0.028	0.030	0.032	0.033	0.035	0.034	0.036	0.037	0.038	0.039

Table 3.9: **Markov Chain Estimator.** Table contains the forecasting results of the daily return density. Analysis was made using the Markov chain (MC) estimator. Rows correspond to analyzed stocks, columns correspond to the intra-daily intervals. In particular, column with header A corresponds to the intra-daily interval [9:30a.m. - A]. The numbers in the first column indicate the mean log-predictive likelihood (LPL) values, whereas the numbers in other column are changes in LPL relative to the first column. An intersection of a row and a column is a cell and each cell contains 2 numbers. The top numbers correspond to the variance specification given by (3.8). The bottom numbers correspond to the variance specification given by (3.9). The out-of-sample period is January, 2007 - December, 2013.

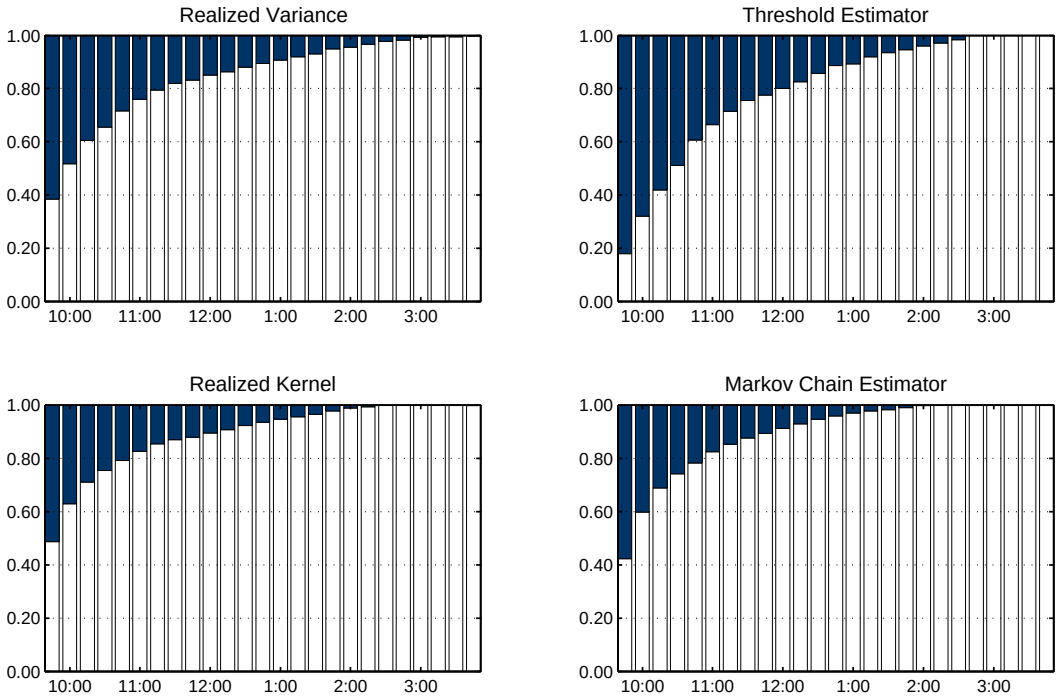

(a) The target variable is the realized measure of variance computed with the corresponding realized estimator at the end of the day.

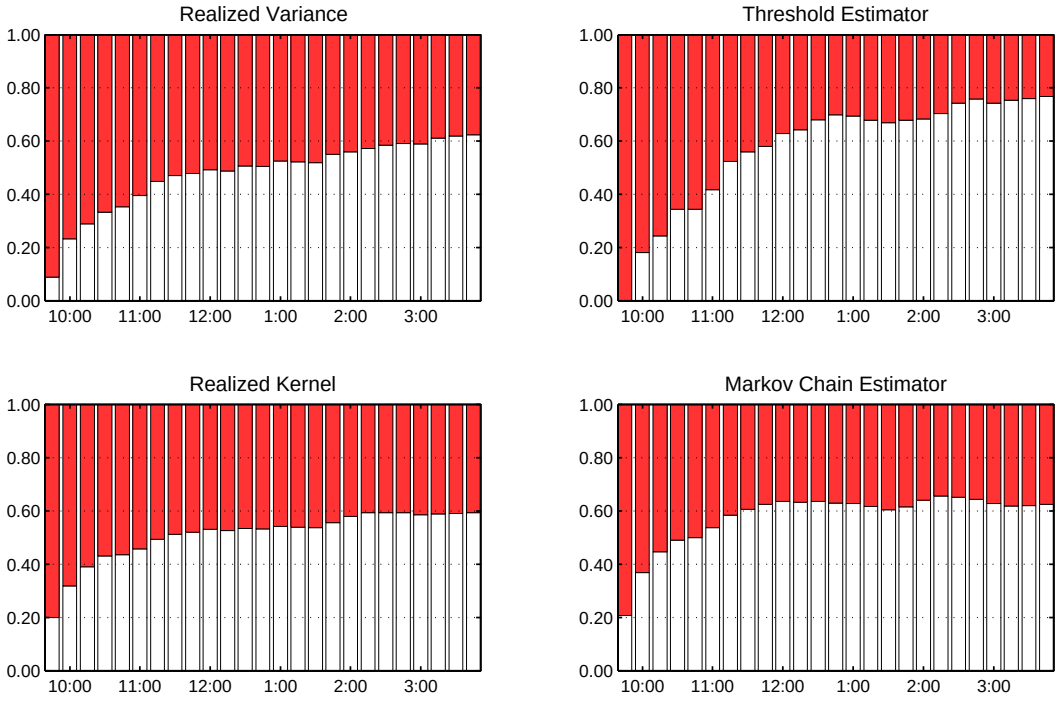

(b) The target variable is the squared open-to-close return (the same for all 4 cases).

Figure 3.4: Subplots contain the averaged (across all considered stocks) optimal forecast weights for 4 estimators. The lower (white) parts of the bar-plots correspond to weights for the forecasts based on current intra-daily measures. The upper (colored) parts of the bar-plots correspond to weights for the ex-ante model-based forecasts. The optimal static weights are constructed as in Bates and Granger (1969) with an additionally imposed constraint, $\omega \in [0, 1]$. The target variables are specified below the plots. The out-of-sample period is January, 2007 - December, 2013.

Appendix C

Supplement to Chapter III

C.1 Estimating Intra-daily Variance Fractions

The usual practice in the literature is to estimate a ("U-shape") periodic intra-daily volatility pattern that is oftenly considered as a deterministic seasonal component of the intra-daily return variance (e.g., Andersen and Bollerslev (1997), Engle and Sokalska (2012)). We slightly depart from this practice and instead estimate average volatility fractions or, more precisely, cumulative shares of volatility that realize before the certain intra-daily moments of time (knots). We investigate the fractions of those volatility quantities, that correspond to the estimator in use. For example, if we use TV, the latent quantity that is measured is supposed to be the integrated variance (IV). In contrast, if we use RK, the latent quantity is the quadratic variation (QV), that potentially includes the jump volatility component as well. As a result, the variance distributions computed by means of distinct realized measures can be different.

We now describe a simple method for estimation of the cumulative variance fractions generated during trading days using possibly biased intra-daily measures. Suppose that we have realized measures $x_{t,i}$ (as specified above) for days $t = 1, ..., T$ and a set of intra-daily knots s_i, $i = 1, ..., N$. The most straightforward way to estimate the average variance fraction q_i that arrives between the market opening and a time point s_i is to compute a ratio

$$q_i = \sum_{t=1}^{T} x_{t,i} \Big/ \sum_{t=1}^{T} x_t \tag{C.1.1}$$

that would correspond to the method of moments solution based on the unconditional moment $\mathrm{E}[q_i x_t - x_{t,i}] = 0$. A similar approach was used, for instance, in Koopman et al. (2005) and Hansen and Lunde (2005) for scaling up an open-to-close realized variance measure to the latent close-to-close variance[1].

However, this approach is not perfect as long as x_t and $x_{i,t}$ are possibly biased. Then, the true moment of interest is $\mathrm{E}[q_i V_t - V_{t,i}] = 0$, where V_t and $V_{t,i}$ are unobservable variance quantities (that depend on the estimator choice). Thus, as far as we assume a multiplicative form of the bias, the moment condition can be expressed as $\mathrm{E}[q_i(1+\tilde{b})x_t - (1+\tilde{b}_i)x_{t,i}] = 0$, where $\tilde{b}$ and $\tilde{b}_i$ are bias components (> -1).

We suggest to use the additional information from the remaining part of the day. It can be associated with the following moment, $\mathrm{E}[(1-q_i)(1+\tilde{b})x_t - (1+\tilde{b}_{-i})x_{t,-i}] = 0$. Here $x_{t,-i}$ denotes the realized volatility measure computed over the intra-daily period that starts at s_i and lasts until the market closure and $\tilde{b}_{-i}$ is the corresponding bias component. After few manipulations we get

$$\mathrm{E}[(1+\tilde{B}_i)x_{t,i} + (1+\tilde{B}_{-i})x_{t,-i} - x_t] = 0$$

which implies that generically $\mathrm{E}[x_{t,i} + x_{t,-i} - x_t] \neq 0$. Note, that we expect both $\tilde{B}_i$ and $\tilde{B}_{-i}$ be sufficiently

[1] In these papers, the latent close-to-close variance was approximated by squared demeaned close-to-close returns.

close to zero[2].

We impose a further simplifying assumption on the structure of the bias terms. Namely, we assume that $\tilde{B}_i = (1 - s_i)B_i$ and $\tilde{B}_{-i} = s_i B_i$. Intuitively, this means that the relative bias $\tilde{B}_i$ ($\tilde{B}_{-i}$) disappears as long as the intra-daily period $[0, s_i]$ ($[s_i, 1]$) approaches the length of the whole trading day, $[0, 1]$. Based on this assumption, we can specify the following moment conditions

$$g(\theta_i) = \begin{pmatrix} q_i x_t^* - (1 + (1 - s_i)B_i)x_{t,i} \\ (1 - q_i)x_t^* - (1 + s_i B_i)x_{t,-i} \end{pmatrix}; \quad \forall i = 1, ..., N \tag{C.1.2}$$

where $\theta_i = (q_i, B_i)'$ is a parameter vector. Instead of x_t, we use a variable x_t^* that is supposed to be highly correlated with the daily realized measure x_t. It is introduced to break the excessive correlation between the two considered moments. We exploit the strong persistence in daily realized volatility measures and choose $x_t^* = \frac{1}{2}(x_{t-1} + x_{t+1})$ to instrumentize x_t. Then we use the iterated GMM approach to find the parameter vector that solves

$$\hat{\theta}_i = arg\min_{\theta_i} \left(\frac{1}{T} \sum_{t=1}^{T} g(\theta_i) \right)' \hat{W} \left(\frac{1}{T} \sum_{t=1}^{T} g(\theta_i) \right); \quad \forall i = 1, ..., N$$

for the moment conditions specified in (C.1.2). Since the weighting matrix $\hat{W}$ converges in probability to the inverse of the long-run covariance (Newey and West (1987)) of the moments, the procedure automatically puts more weight on the more informative (less noisy) moments and increases the estimation efficiency relative to (C.1.1).

A proper distribution curve $\{q_i\}_{i=1}^N$, though, requires monotonicity[3], i.e. for two intra-daily knots s_i and s_j such that $i < j$, we should require that $q_i \leq q_j$. Although in our empirical application we never encountered the violation of monotonicity, if one decides to increase the number of intra-daily knots to enhance the precision of the extracted distribution curve, this requirement may be violated. In such cases, we recommend to use constrained sequential estimation with the restrictions $\hat{q}_i \geq \hat{q}_{i-1}$, $i = 1, ..., N$.

As long as the cumulative variance distribution q_i is estimated at sufficiently large number of intra-daily time points, a more standard intra-daily periodic volatility pattern can be restored from q_s. However, such pattern will appear as a set of mass points, possibly erratic. The techniques suggested in Boudt et al. (2011) or Bos et al. (2012) can be applied to regularize such pattern and to filter outliers.

Finally, the method suggested above relies on the unconditional moments which do not impose much structure on the intra-daily volatility distribution. In particular, the periodic pattern may be non-deterministic and change significantly over time. Thus, it is interesting to use conditional moments of type $\mathrm{E}[q_i V_t - V_{t,i} | V_t] = 0$ with a properly selected vector of instruments to estimate q_i under more strong assumptions. Such moments would imply a presence of a deterministic component in the intra-daily variance distribution conditionally on the daily variance level.

Figure C.1.1 contains 95% confidence intervals for the estimated cumulative distributions of intra-daily variance measured with the MC estimator for the whole sample period (10 years). One curve is obtained using the unconditional moments, the other is based on the conditional moments with the lagged realized volatility measure, x_{t-1}, selected as an instrument. A comparison of these two curves can be considered as an informal check for the assumption that variance distributions are conditionally deterministic. As we can see from the plots, there is a noticeable discrepancy between the curves obtained using unconditional and conditional moments for most of the examined assets. To a certain extent, this evidence invalidates the assumption that periodic intra-daily volatility distribution is static (conditionally on the latent variance level) throughout the whole sample period. The validity of this hypothesis can further be additionally investigated by using another instruments or considering the sub-samples.

[2]For example, in case of the classical RV estimator, we have $x_{t,i} + x_{t,-i} = x_t$ by construction. Therefore, in that case, $\tilde{B}_i = \tilde{B}_{-i} = 0$. Though, this is not necessarily true for the parametric MC estimator, for example, where the obtained measure heavily depends on the quality of the estimated state transition matrix.

[3]Also, we should impose border conditions. Namely, $q_0 = 0$ and $q_1 = 1$.

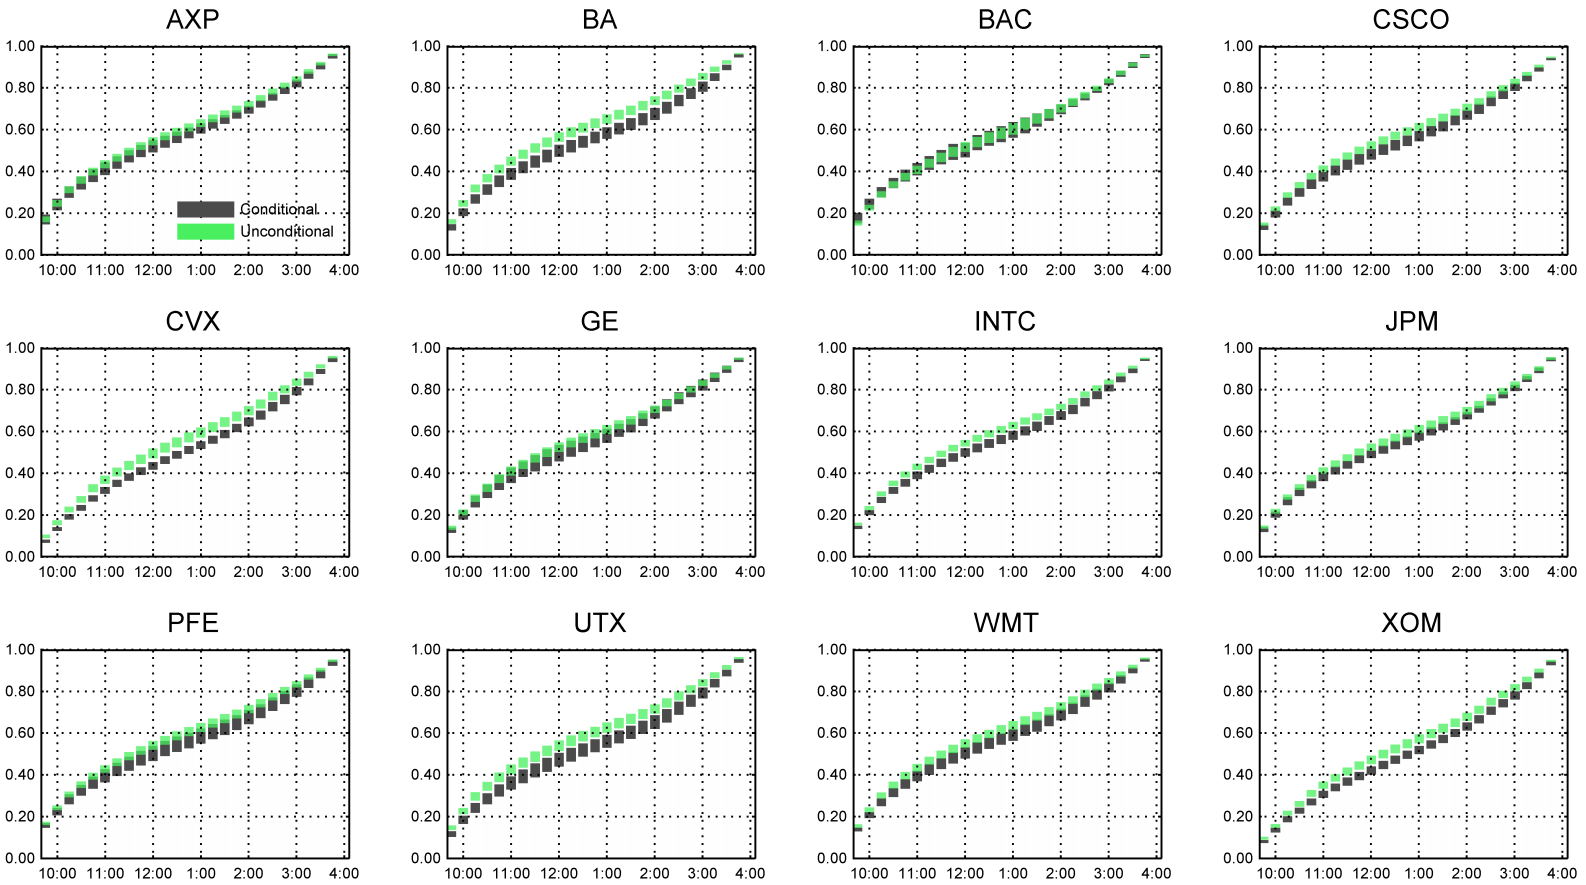

Figure C.1.1: Estimated cumulative intra-daily variance fractions for 12 considered assets by the Markov chain estimator. Green bars correspond to the 95% confidence intervals for the estimates obtained with the unconditional moments. Black bars correspond to the 95% confidence intervals for the estimates obtained with the conditional moments.

C.2 Subsample Analysis

We investigate the sensitiveness of our intra-daily forecasting results with respect to the subsample selection. For this purpose we repeat the analysis held in Section 3.3.2 and compute the predictive root of mean squared error (RMSE) for different sets (subsamples) of trading days. We use the RMSE criterion instead of the Mincer-Zarnowitz R^2 since heterogeneous variation across distinct subsamples would complicate the R^2 comparison. We consider the following subsets: days of the week, pre-crisis, crisis, post-crisis periods (partitioned in accordance with NBER dates), days with low volatility (bottom quintile), days with high volatility (top quintile), days with no jumps, days with significant jumps, days of scheduled FOMC statements, days of different months. The results are provided in Table C.2.1.

As it can be seen from the results, the quality of the out-of-sample fit is indeed sensitive to the choice of a subsample. Figure C.2.1 contains several interesting comparisons between the out-of-sample RMSE for the selected subsamples. We may note that before the crisis the intra-daily variance predictability is on average lower than after the crisis. On low volatile days the intra-daily measures predict daily variance enormously better than on highly volatile days. This is not surprising since the mispredictions on low volatile days are not that large in absolute terms. In addition, the sensitivity of the criterion to outliers lead to very high RMSE values on highly volatile days when abnormally large mispredictions appear quite often.

Comparison between the fit on Mondays (start of the week) and Thursdays (middle of the week) indicate that the intra-daily predictability of a daily variance is slightly better after a weekend. An interesting evidence is captured by the comparison of the intra-daily predictability on jump days and the Federal Open Market Committee (FOMC) scheduled statement days. Since the FOMC releases usually cause sharp price changes or jumps (see Andersen et al. (2007), Zebedee et al. (2008), etc.), the FOMC days can be treated as a subset of jump days. It might be the case then, that the scheduled FOMC announcements (that traders are aware about in advance) lead to a considerable increase in the intra-daily predictability of a daily variance.

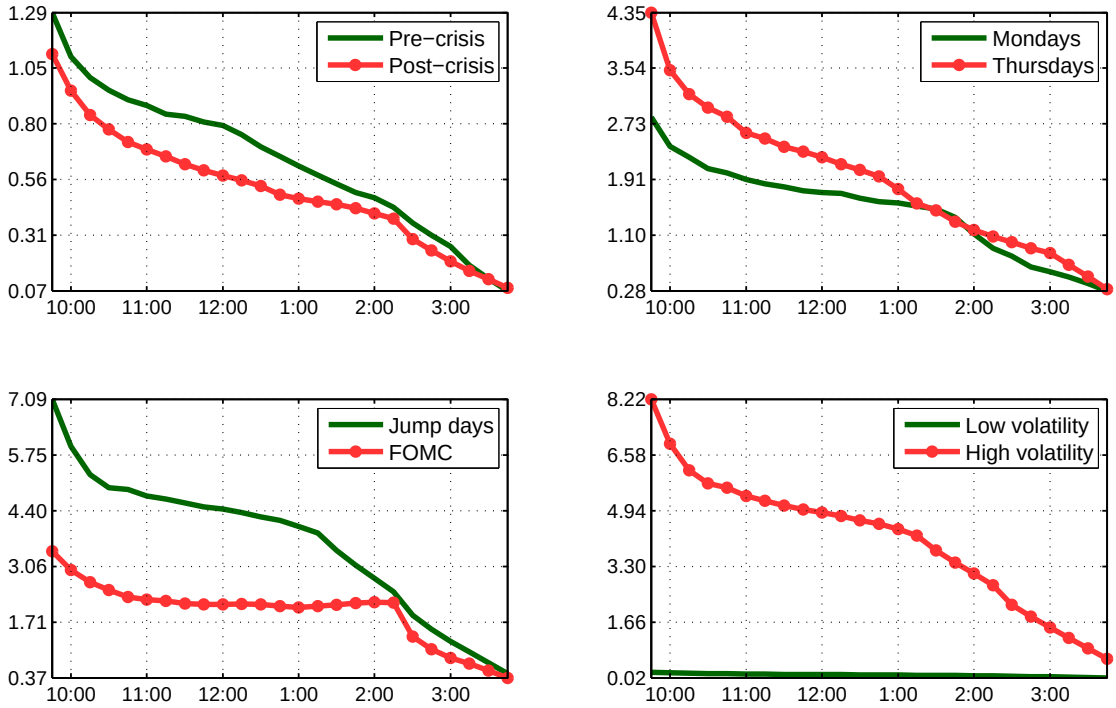

Figure C.2.1: Subsample comparison of the predictive RMSE for intra-daily forecasting with the linear specification (3.5). The realized volatility measures are obtained using the Markov chain estimator. Plots show aggregated results over 12 considered stocks. The out-of-sample period is January, 2007 - December, 2013.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30	12:45	1:00	1:15	1:30	1:45	2:00	2:15	2:30	2:45	3:00	3:15	3:30	3:45	N
Full sample	4.007	3.298	2.900	2.712	2.642	2.522	2.451	2.380	2.318	2.275	2.219	2.156	2.102	2.026	1.940	1.745	1.579	1.433	1.276	1.012	0.852	0.708	0.569	0.422	0.274	1745
Monday	2.821	2.395	2.233	2.072	2.008	1.913	1.849	1.805	1.748	1.718	1.705	1.635	1.589	1.566	1.524	1.478	1.355	1.115	0.902	0.792	0.633	0.561	0.487	0.390	0.284	330
Tuesday	3.518	2.833	2.402	2.160	2.079	1.945	1.766	1.669	1.603	1.536	1.492	1.425	1.396	1.367	1.318	1.283	1.236	1.159	1.061	0.863	0.701	0.594	0.480	0.351	0.211	356
Wednesday	2.904	2.400	2.137	1.987	1.912	1.867	1.842	1.763	1.707	1.662	1.604	1.559	1.476	1.436	1.420	1.399	1.373	1.320	1.227	0.936	0.820	0.671	0.533	0.387	0.259	361
Thursday	4.354	3.508	3.157	2.960	2.827	2.593	2.508	2.387	2.318	2.237	2.134	2.052	1.957	1.769	1.558	1.456	1.293	1.174	1.074	0.997	0.903	0.832	0.664	0.494	0.304	351
Friday	4.513	3.878	3.360	3.137	3.075	3.040	2.974	2.958	2.901	2.892	2.843	2.801	2.785	2.742	2.666	2.216	1.870	1.674	1.501	1.150	0.955	0.723	0.537	0.380	0.219	347
Pre-crisis	1.290	1.097	1.005	0.950	0.907	0.882	0.844	0.834	0.811	0.795	0.756	0.702	0.659	0.617	0.577	0.539	0.502	0.479	0.436	0.367	0.313	0.264	0.184	0.124	0.071	229
Crisis	7.724	6.423	5.660	5.297	5.162	4.925	4.793	4.660	4.546	4.469	4.377	4.265	4.171	4.022	3.849	3.447	3.109	2.806	2.487	1.965	1.659	1.375	1.105	0.824	0.539	391
Post-crisis	1.108	0.948	0.841	0.777	0.722	0.690	0.659	0.625	0.598	0.576	0.554	0.531	0.491	0.474	0.460	0.449	0.432	0.410	0.386	0.298	0.247	0.200	0.159	0.121	0.083	1125
Low variance	0.188	0.170	0.156	0.147	0.139	0.133	0.128	0.124	0.121	0.117	0.113	0.109	0.105	0.102	0.098	0.093	0.088	0.084	0.075	0.063	0.056	0.048	0.040	0.032	0.022	350
High variance	8.219	6.902	6.125	5.743	5.613	5.372	5.231	5.092	4.966	4.885	4.780	4.652	4.549	4.390	4.204	3.774	3.408	3.086	2.744	2.168	1.822	1.506	1.200	0.892	0.580	349
No jumps	3.217	2.586	2.226	2.051	1.902	1.756	1.644	1.541	1.457	1.395	1.344	1.269	1.227	1.169	1.108	1.044	0.987	0.927	0.860	0.771	0.701	0.614	0.488	0.376	0.251	738
Jumps	7.090	5.946	5.271	4.951	4.907	4.754	4.674	4.580	4.494	4.436	4.350	4.251	4.159	4.020	3.855	3.443	3.080	2.769	2.442	1.869	1.533	1.243	0.995	0.733	0.468	349
FOMC	3.415	2.962	2.673	2.478	2.311	2.250	2.221	2.157	2.135	2.132	2.144	2.134	2.089	2.066	2.091	2.130	2.163	2.190	2.176	1.362	1.050	0.839	0.709	0.534	0.365	61
January	3.642	2.765	2.363	2.103	1.959	1.835	1.772	1.692	1.633	1.585	1.517	1.397	1.334	1.285	1.237	1.173	1.097	1.034	0.972	0.858	0.791	0.709	0.581	0.421	0.287	140
February	4.540	4.060	3.469	3.429	3.498	3.486	3.504	3.533	3.494	3.506	3.528	3.520	3.516	3.480	3.468	2.878	2.424	2.191	1.878	1.263	1.011	0.621	0.452	0.315	0.187	135
March	3.364	2.717	2.429	2.293	2.290	2.199	2.060	1.920	1.817	1.762	1.702	1.651	1.592	1.567	1.527	1.460	1.433	1.376	1.317	1.082	0.763	0.552	0.420	0.318	0.187	152
April	1.590	1.430	1.256	1.185	1.122	1.079	1.057	1.027	1.006	0.982	0.956	0.928	0.910	0.878	0.832	0.792	0.740	0.689	0.501	0.408	0.368	0.331	0.296	0.232	0.135	146
May	1.609	1.317	1.071	0.974	0.870	0.817	0.745	0.698	0.679	0.660	0.640	0.622	0.588	0.564	0.534	0.509	0.479	0.437	0.396	0.356	0.308	0.273	0.231	0.158	0.102	147
June	0.896	0.776	0.680	0.638	0.591	0.564	0.532	0.507	0.488	0.472	0.461	0.451	0.436	0.422	0.407	0.391	0.378	0.363	0.326	0.272	0.229	0.195	0.158	0.124	0.085	149
July	1.769	1.591	1.424	1.323	1.267	1.212	1.121	1.038	0.988	0.925	0.841	0.754	0.685	0.639	0.599	0.561	0.528	0.494	0.442	0.402	0.365	0.276	0.182	0.140	0.085	145
August	1.983	1.691	1.601	1.501	1.440	1.398	1.351	1.283	1.228	1.186	1.128	1.080	1.042	1.020	0.995	0.980	0.952	0.918	0.877	0.659	0.528	0.397	0.300	0.211	0.136	155
September	3.740	3.084	2.967	2.727	2.632	2.588	2.505	2.515	2.476	2.425	2.322	2.228	2.092	1.993	1.882	1.787	1.537	1.147	0.827	0.657	0.596	0.559	0.453	0.322	0.157	141
October	5.102	4.168	3.765	3.545	3.254	3.017	2.952	2.867	2.791	2.765	2.767	2.685	2.602	2.505	2.343	2.205	2.019	1.845	1.729	1.382	1.180	1.033	0.832	0.628	0.434	152
November	5.677	4.392	3.691	3.215	2.971	2.790	2.613	2.372	2.280	2.163	2.084	2.004	1.981	1.797	1.680	1.613	1.527	1.382	1.241	1.113	1.010	0.859	0.645	0.540	0.403	139
December	2.392	1.859	1.610	1.456	1.313	1.222	1.159	1.098	1.060	1.010	0.975	0.939	0.924	0.915	0.907	0.899	0.879	0.858	0.817	0.689	0.618	0.549	0.447	0.326	0.190	144

Table C.2.1: **Subsample Analysis.** Subsample comparison of the predictive RMSE for intra-daily forecasting with the linear specification (3.5). Realized volatility measures are obtained using the Markov chain estimator. Table contains aggregated results over 12 considered stocks. Rows correspond to the subsample categories, columns correspond to the intra-daily intervals. In particular, column with header A corresponds to intra-daily interval [9:30a.m. - A] that is used to compute a realized mesure. The number of days in a given subsample is provided in the right column. The out-of-sample period is January, 2007 - December, 2013.

C.3 Additional Tables. Forecasting Realized Log-measures

Tables contain the Mincer-Zarnowitz R^2 for the regressions that forecast the end-of-day log realized measures with the intra-daily log-measures alone (top numbers) and in conjunction with the ex-ante forecasts (bottom numbers). The out-of-sample period is January, 2007 - December, 2013.

	9:30	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	0.000	0.546	0.730	0.806	0.848	0.875	0.897	0.912	0.925	0.931	0.939	0.945	0.950
	0.817	0.874	0.895	0.909	0.919	0.926	0.934	0.941	0.947	0.950	0.954	0.958	0.960
BA	0.000	0.423	0.603	0.711	0.762	0.802	0.828	0.854	0.871	0.886	0.899	0.908	0.926
	0.649	0.757	0.799	0.837	0.859	0.873	0.885	0.899	0.908	0.916	0.923	0.929	0.944
BAC	0.000	0.611	0.759	0.836	0.865	0.885	0.900	0.916	0.927	0.933	0.938	0.943	0.949
	0.811	0.861	0.887	0.908	0.917	0.924	0.931	0.938	0.944	0.948	0.950	0.953	0.957
CSCO	0.000	0.411	0.619	0.716	0.767	0.804	0.829	0.848	0.863	0.878	0.890	0.900	0.911
	0.681	0.765	0.812	0.841	0.860	0.876	0.887	0.896	0.903	0.911	0.918	0.924	0.931
CVX	0.000	0.424	0.589	0.683	0.741	0.799	0.833	0.854	0.872	0.887	0.898	0.908	0.918
	0.750	0.797	0.822	0.845	0.861	0.883	0.896	0.905	0.913	0.920	0.926	0.931	0.937
GE	0.000	0.490	0.669	0.762	0.812	0.842	0.863	0.880	0.894	0.905	0.916	0.924	0.931
	0.774	0.824	0.850	0.874	0.891	0.901	0.909	0.916	0.924	0.930	0.936	0.941	0.945
INTC	0.000	0.420	0.611	0.705	0.757	0.803	0.834	0.855	0.872	0.889	0.901	0.911	0.919
	0.703	0.772	0.813	0.842	0.862	0.880	0.893	0.903	0.911	0.921	0.928	0.934	0.938
JPM	0.000	0.517	0.724	0.808	0.843	0.872	0.889	0.903	0.914	0.923	0.929	0.934	0.940
	0.804	0.847	0.876	0.898	0.909	0.920	0.927	0.934	0.939	0.944	0.946	0.949	0.952
PFE	0.000	0.451	0.608	0.707	0.754	0.790	0.818	0.839	0.857	0.870	0.881	0.893	0.904
	0.677	0.772	0.812	0.841	0.858	0.872	0.882	0.892	0.899	0.905	0.911	0.918	0.925
UTX	0.000	0.433	0.599	0.713	0.771	0.814	0.841	0.864	0.878	0.893	0.904	0.915	0.924
	0.689	0.783	0.819	0.854	0.876	0.891	0.903	0.914	0.921	0.929	0.934	0.940	0.945
WMT	0.000	0.481	0.632	0.717	0.775	0.816	0.842	0.870	0.886	0.899	0.909	0.917	0.925
	0.677	0.777	0.817	0.847	0.871	0.888	0.898	0.913	0.922	0.929	0.934	0.938	0.943
XOM	0.000	0.450	0.615	0.703	0.748	0.798	0.823	0.845	0.863	0.877	0.889	0.899	0.909
	0.748	0.792	0.821	0.844	0.857	0.877	0.887	0.895	0.904	0.911	0.917	0.922	0.929

Realized Variance Estimator.

	9:30	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	0.000	0.382	0.594	0.697	0.758	0.793	0.829	0.858	0.879	0.892	0.904	0.917	0.926
	0.845	0.859	0.870	0.880	0.889	0.896	0.905	0.914	0.922	0.928	0.933	0.939	0.944
BA	0.000	0.252	0.411	0.532	0.591	0.652	0.698	0.739	0.767	0.798	0.824	0.842	0.859
	0.709	0.736	0.757	0.779	0.797	0.808	0.824	0.840	0.851	0.866	0.879	0.887	0.898
BAC	0.000	0.480	0.660	0.746	0.788	0.818	0.844	0.871	0.887	0.899	0.910	0.920	0.929
	0.823	0.843	0.860	0.873	0.882	0.892	0.902	0.912	0.919	0.926	0.931	0.937	0.944
CSCO	0.000	0.243	0.458	0.575	0.635	0.689	0.728	0.762	0.785	0.809	0.831	0.853	0.871
	0.702	0.722	0.751	0.778	0.795	0.814	0.830	0.844	0.855	0.866	0.877	0.890	0.900
CVX	0.000	0.314	0.480	0.579	0.653	0.725	0.774	0.806	0.832	0.851	0.868	0.884	0.898
	0.782	0.792	0.805	0.819	0.832	0.852	0.867	0.879	0.890	0.898	0.907	0.915	0.923
GE	0.000	0.331	0.570	0.671	0.738	0.776	0.805	0.829	0.849	0.869	0.886	0.901	0.914
	0.794	0.809	0.827	0.845	0.861	0.872	0.880	0.890	0.899	0.910	0.918	0.927	0.934
INTC	0.000	0.271	0.449	0.566	0.627	0.683	0.728	0.763	0.794	0.818	0.841	0.859	0.874
	0.738	0.753	0.768	0.791	0.806	0.822	0.839	0.851	0.863	0.875	0.887	0.898	0.906
JPM	0.000	0.408	0.635	0.718	0.770	0.810	0.835	0.859	0.875	0.890	0.901	0.910	0.923
	0.819	0.833	0.852	0.866	0.877	0.889	0.898	0.907	0.915	0.922	0.927	0.933	0.940
PFE	0.000	0.259	0.430	0.554	0.621	0.680	0.724	0.762	0.792	0.816	0.832	0.849	0.868
	0.725	0.745	0.767	0.792	0.809	0.824	0.837	0.850	0.862	0.874	0.882	0.891	0.902
UTX	0.000	0.235	0.397	0.531	0.614	0.670	0.714	0.753	0.782	0.811	0.833	0.856	0.876
	0.750	0.765	0.781	0.800	0.817	0.831	0.844	0.859	0.869	0.882	0.891	0.902	0.913
WMT	0.000	0.335	0.474	0.568	0.643	0.693	0.738	0.772	0.805	0.832	0.850	0.869	0.885
	0.736	0.756	0.774	0.793	0.815	0.829	0.844	0.855	0.869	0.883	0.892	0.903	0.912
XOM	0.000	0.349	0.525	0.627	0.680	0.732	0.773	0.802	0.827	0.849	0.866	0.878	0.892
	0.772	0.786	0.803	0.820	0.830	0.845	0.860	0.870	0.881	0.892	0.901	0.908	0.916

Threshold Variance Estimator.

	9:30	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	0.000	0.769	0.852	0.892	0.912	0.927	0.938	0.946	0.952	0.956	0.960	0.964	0.967
	0.865	0.913	0.926	0.936	0.943	0.949	0.954	0.958	0.961	0.964	0.966	0.969	0.971
BA	0.000	0.663	0.758	0.816	0.845	0.867	0.881	0.895	0.903	0.914	0.921	0.927	0.943
	0.731	0.823	0.850	0.876	0.890	0.901	0.909	0.918	0.922	0.929	0.934	0.938	0.952
BAC	0.000	0.830	0.881	0.909	0.921	0.932	0.939	0.946	0.951	0.955	0.957	0.960	0.963
	0.847	0.900	0.917	0.929	0.936	0.943	0.947	0.952	0.956	0.959	0.961	0.963	0.966
CSCO	0.000	0.643	0.756	0.810	0.838	0.862	0.877	0.891	0.899	0.908	0.916	0.924	0.931
	0.733	0.815	0.850	0.870	0.883	0.896	0.905	0.914	0.919	0.924	0.930	0.936	0.941
CVX	0.000	0.653	0.752	0.803	0.834	0.874	0.892	0.906	0.916	0.924	0.931	0.937	0.943
	0.807	0.851	0.868	0.883	0.894	0.912	0.922	0.930	0.935	0.940	0.944	0.948	0.952
GE	0.000	0.738	0.818	0.861	0.884	0.900	0.912	0.923	0.930	0.936	0.941	0.946	0.951
	0.819	0.868	0.888	0.903	0.914	0.922	0.929	0.935	0.941	0.945	0.948	0.952	0.956
INTC	0.000	0.666	0.778	0.829	0.862	0.884	0.899	0.911	0.919	0.928	0.935	0.941	0.946
	0.759	0.835	0.865	0.885	0.903	0.914	0.922	0.929	0.935	0.941	0.946	0.950	0.954
JPM	0.000	0.793	0.866	0.901	0.916	0.929	0.937	0.944	0.949	0.953	0.955	0.958	0.961
	0.853	0.900	0.916	0.929	0.936	0.943	0.948	0.952	0.956	0.959	0.961	0.963	0.965
PFE	0.000	0.644	0.753	0.812	0.842	0.860	0.878	0.891	0.901	0.910	0.916	0.924	0.932
	0.755	0.834	0.862	0.882	0.894	0.903	0.911	0.918	0.923	0.928	0.932	0.937	0.943
UTX	0.000	0.619	0.736	0.804	0.838	0.864	0.882	0.898	0.909	0.920	0.927	0.935	0.941
	0.757	0.836	0.863	0.886	0.900	0.911	0.919	0.928	0.933	0.940	0.944	0.949	0.953
WMT	0.000	0.650	0.763	0.822	0.855	0.878	0.893	0.909	0.918	0.925	0.932	0.938	0.943
	0.756	0.833	0.863	0.884	0.899	0.911	0.918	0.929	0.934	0.939	0.943	0.947	0.951
XOM	0.000	0.665	0.759	0.809	0.834	0.871	0.885	0.899	0.908	0.917	0.924	0.930	0.937
	0.805	0.847	0.865	0.881	0.889	0.906	0.913	0.921	0.927	0.932	0.936	0.940	0.945

Realized Kernel Estimator.

	9:30	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	0.000	0.667	0.852	0.893	0.912	0.929	0.941	0.948	0.954	0.958	0.962	0.965	0.968
	0.882	0.910	0.929	0.940	0.945	0.952	0.956	0.960	0.963	0.965	0.968	0.970	0.973
BA	0.000	0.559	0.724	0.803	0.840	0.865	0.881	0.895	0.903	0.915	0.922	0.928	0.942
	0.768	0.822	0.853	0.878	0.892	0.904	0.912	0.920	0.924	0.932	0.937	0.941	0.951
BAC	0.000	0.824	0.880	0.909	0.924	0.935	0.942	0.948	0.954	0.957	0.960	0.963	0.966
	0.869	0.906	0.921	0.932	0.939	0.946	0.950	0.954	0.959	0.961	0.963	0.966	0.968
CSCO	0.000	0.640	0.743	0.796	0.829	0.854	0.871	0.886	0.894	0.902	0.911	0.919	0.927
	0.742	0.817	0.846	0.865	0.878	0.891	0.901	0.909	0.914	0.919	0.926	0.932	0.938
CVX	0.000	0.640	0.757	0.815	0.847	0.881	0.899	0.912	0.921	0.929	0.935	0.941	0.947
	0.832	0.860	0.875	0.890	0.901	0.917	0.926	0.933	0.938	0.943	0.947	0.951	0.955
GE	0.000	0.741	0.823	0.865	0.887	0.905	0.918	0.927	0.935	0.940	0.945	0.950	0.954
	0.832	0.871	0.892	0.908	0.917	0.926	0.934	0.939	0.944	0.948	0.952	0.956	0.959
INTC	0.000	0.668	0.769	0.822	0.855	0.877	0.891	0.903	0.913	0.922	0.929	0.937	0.942
	0.777	0.839	0.861	0.881	0.896	0.907	0.914	0.921	0.929	0.935	0.940	0.946	0.950
JPM	0.000	0.782	0.868	0.904	0.920	0.933	0.941	0.949	0.953	0.957	0.959	0.962	0.964
	0.869	0.905	0.922	0.934	0.941	0.947	0.952	0.957	0.960	0.963	0.965	0.966	0.968
PFE	0.000	0.644	0.757	0.812	0.844	0.864	0.881	0.894	0.902	0.910	0.916	0.923	0.930
	0.773	0.840	0.865	0.884	0.896	0.905	0.913	0.919	0.924	0.928	0.932	0.937	0.941
UTX	0.000	0.354	0.709	0.788	0.829	0.860	0.879	0.896	0.906	0.917	0.924	0.932	0.939
	0.787	0.816	0.861	0.884	0.898	0.910	0.919	0.927	0.932	0.938	0.942	0.947	0.951
WMT	0.000	0.622	0.750	0.814	0.849	0.874	0.891	0.908	0.918	0.925	0.932	0.937	0.942
	0.765	0.831	0.864	0.886	0.901	0.912	0.920	0.931	0.936	0.941	0.945	0.948	0.951
XOM	0.000	0.655	0.768	0.823	0.851	0.884	0.897	0.909	0.918	0.925	0.931	0.937	0.943
	0.826	0.858	0.876	0.890	0.899	0.915	0.922	0.929	0.935	0.939	0.943	0.946	0.950

Markov Chain Estimator.

C.4 Additional Tables. Forecasting Return Density Assuming Heavy-tailed Distribution

Tables contain mean log-pedictive likelihoods of the density forecasting exercise where daily returns are assumed to have a conditional Student's t-distribution. Top numbers correspond to the variance specification (3.8). Bottom numbers correspond to the variance specification (3.9). The out-of-sample period is January, 2007 - December, 2013.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	*-0.743*	0.133	0.174	0.190	0.205	0.215	0.221	0.227	0.230	0.234	0.237	0.240
	-0.508	0.001	0.003	0.003	0.005	0.007	0.008	0.009	0.010	0.011	0.013	0.014
BA	*-0.526*	0.142	0.174	0.192	0.203	0.210	0.219	0.224	0.229	0.232	0.235	0.238
	-0.305	0.003	0.003	0.007	0.008	0.009	0.012	0.013	0.015	0.017	0.018	0.021
BAC	*-0.912*	0.134	0.172	0.184	0.192	0.197	0.205	0.211	0.213	0.214	0.215	0.219
	-0.729	0.009	0.015	0.018	0.020	0.022	0.026	0.030	0.032	0.032	0.033	0.036
CSCO	*-0.533*	0.156	0.189	0.207	0.217	0.222	0.225	0.229	0.231	0.234	0.237	0.240
	-0.301	0.003	0.005	0.008	0.010	0.011	0.011	0.012	0.012	0.013	0.014	0.015
CVX	*-0.361*	0.144	0.175	0.192	0.205	0.213	0.219	0.221	0.224	0.225	0.226	0.227
	-0.140	0.002	0.004	0.005	0.006	0.007	0.009	0.010	0.011	0.012	0.012	0.013
GE	*-0.498*	0.130	0.169	0.190	0.200	0.208	0.216	0.222	0.225	0.231	0.234	0.237
	-0.274	0.003	0.006	0.008	0.009	0.010	0.012	0.014	0.015	0.017	0.018	0.019
INTC	*-0.546*	0.131	0.171	0.184	0.194	0.204	0.209	0.212	0.217	0.220	0.221	0.223
	-0.326	0.003	0.005	0.005	0.005	0.006	0.008	0.008	0.009	0.010	0.010	0.011
JPM	*-0.752*	0.132	0.162	0.170	0.177	0.184	0.188	0.191	0.193	0.193	0.195	0.197
	-0.561	0.003	0.005	0.006	0.008	0.010	0.011	0.012	0.013	0.013	0.013	0.014
PFE	*-0.322*	0.104	0.147	0.164	0.179	0.188	0.194	0.199	0.203	0.207	0.209	0.211
	-0.115	0.002	0.003	0.003	0.005	0.006	0.007	0.009	0.009	0.011	0.012	0.013
UTX	*-0.367*	0.136	0.176	0.197	0.205	0.213	0.219	0.223	0.227	0.229	0.231	0.233
	-0.131	0.004	0.004	0.005	0.005	0.007	0.008	0.008	0.009	0.010	0.010	0.010
WMT	*-0.153*	0.158	0.198	0.221	0.230	0.239	0.249	0.253	0.257	0.260	0.262	0.263
	0.101	0.004	0.005	0.009	0.009	0.010	0.013	0.015	0.016	0.016	0.017	0.017
XOM	*-0.242*	0.127	0.149	0.159	0.168	0.172	0.177	0.181	0.185	0.187	0.187	0.190
	-0.052	0.004	0.004	0.004	0.004	0.005	0.006	0.006	0.007	0.008	0.008	0.009

Realized Variance Estimator.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	*-0.831*	0.178	0.245	0.267	0.283	0.298	0.307	0.312	0.314	0.319	0.322	0.326
	-0.504	-0.001	0.000	-0.000	0.002	0.004	0.005	0.006	0.005	0.007	0.008	0.009
BA	*-0.621*	0.202	0.246	0.260	0.270	0.282	0.291	0.297	0.304	0.307	0.310	0.314
	-0.305	-0.000	-0.004	-0.002	-0.004	-0.002	-0.002	-0.001	0.003	0.004	0.006	0.007
BAC	*-0.980*	0.178	0.219	0.240	0.244	0.252	0.260	0.266	0.268	0.271	0.274	0.277
	-0.732	0.008	0.010	0.014	0.015	0.016	0.020	0.023	0.024	0.026	0.028	0.030
CSCO	*-0.641*	0.236	0.283	0.298	0.309	0.317	0.321	0.324	0.327	0.331	0.335	0.338
	-0.304	0.003	0.006	0.006	0.008	0.009	0.009	0.009	0.009	0.010	0.011	0.012
CVX	*-0.458*	0.208	0.250	0.274	0.291	0.299	0.308	0.311	0.313	0.316	0.316	0.318
	-0.141	0.003	0.003	0.004	0.006	0.007	0.009	0.009	0.011	0.012	0.012	0.012
GE	*-0.593*	0.204	0.253	0.276	0.281	0.290	0.296	0.303	0.309	0.314	0.316	0.322
	-0.277	0.002	0.005	0.007	0.005	0.005	0.006	0.008	0.010	0.011	0.012	0.014
INTC	*-0.628*	0.181	0.230	0.242	0.255	0.268	0.277	0.280	0.284	0.286	0.287	0.289
	-0.326	0.002	0.003	0.001	0.001	0.002	0.004	0.004	0.004	0.004	0.004	0.005
JPM	*-0.795*	0.149	0.184	0.199	0.206	0.216	0.221	0.224	0.226	0.224	0.227	0.231
	-0.558	0.001	0.001	0.001	0.001	0.003	0.005	0.006	0.005	0.005	0.005	0.007
PFE	*-0.417*	0.165	0.215	0.238	0.255	0.267	0.277	0.282	0.289	0.291	0.293	0.296
	-0.118	0.002	-0.001	-0.000	0.000	0.001	0.002	0.003	0.006	0.006	0.006	0.008
UTX	*-0.446*	0.185	0.242	0.265	0.275	0.284	0.287	0.294	0.297	0.300	0.304	0.306
	-0.130	0.001	0.002	0.003	0.003	0.005	0.004	0.004	0.004	0.005	0.006	0.006
WMT	*-0.213*	0.181	0.235	0.272	0.283	0.295	0.304	0.308	0.312	0.315	0.316	0.316
	0.102	0.001	0.001	0.006	0.004	0.006	0.009	0.009	0.011	0.011	0.011	0.010
XOM	*-0.289*	0.145	0.178	0.193	0.202	0.206	0.208	0.214	0.221	0.222	0.223	0.226
	-0.051	0.001	0.000	0.001	-0.000	0.000	-0.000	0.001	0.002	0.002	0.002	0.004

Threshold Variance Estimator.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	*-0.571*	0.027	0.039	0.044	0.051	0.055	0.058	0.061	0.063	0.065	0.066	0.067
	-0.507	0.001	0.002	0.002	0.004	0.005	0.006	0.007	0.008	0.009	0.009	0.010
BA	*-0.370*	0.029	0.041	0.049	0.054	0.056	0.062	0.063	0.067	0.068	0.070	0.074
	-0.310	0.002	0.003	0.005	0.007	0.008	0.009	0.010	0.012	0.013	0.014	0.017
BAC	*-0.745*	0.015	0.021	0.025	0.029	0.031	0.034	0.037	0.039	0.040	0.041	0.043
	-0.731	0.006	0.009	0.012	0.015	0.017	0.019	0.022	0.024	0.025	0.026	0.028
CSCO	*-0.350*	0.025	0.033	0.040	0.045	0.048	0.050	0.051	0.052	0.054	0.055	0.056
	-0.303	0.002	0.003	0.006	0.007	0.009	0.010	0.010	0.010	0.011	0.012	0.012
CVX	*-0.186*	0.027	0.031	0.034	0.039	0.043	0.046	0.047	0.048	0.049	0.049	0.050
	-0.140	0.002	0.002	0.003	0.003	0.004	0.005	0.006	0.007	0.007	0.007	0.008
GE	*-0.321*	0.029	0.036	0.041	0.046	0.049	0.052	0.055	0.057	0.059	0.061	0.063
	-0.275	0.005	0.006	0.008	0.009	0.010	0.011	0.013	0.013	0.015	0.015	0.017
INTC	*-0.352*	0.013	0.018	0.020	0.022	0.025	0.027	0.028	0.030	0.032	0.032	0.033
	-0.323	0.001	0.001	0.001	0.002	0.003	0.004	0.004	0.005	0.006	0.006	0.007
JPM	*-0.587*	0.016	0.021	0.022	0.025	0.027	0.029	0.030	0.031	0.031	0.032	0.032
	-0.565	0.004	0.006	0.006	0.008	0.010	0.011	0.012	0.013	0.013	0.014	0.014
PFE	*-0.162*	0.019	0.031	0.035	0.038	0.042	0.045	0.047	0.049	0.050	0.051	0.053
	-0.116	0.001	0.003	0.004	0.005	0.006	0.007	0.008	0.009	0.010	0.010	0.011
UTX	*-0.188*	0.023	0.034	0.038	0.043	0.047	0.050	0.052	0.053	0.056	0.057	0.058
	-0.133	0.002	0.003	0.004	0.004	0.006	0.007	0.007	0.007	0.008	0.009	0.009
WMT	*0.047*	0.030	0.042	0.050	0.053	0.056	0.060	0.062	0.064	0.064	0.065	0.066
	0.099	0.002	0.004	0.007	0.008	0.009	0.011	0.013	0.014	0.014	0.014	0.015
XOM	*-0.098*	0.022	0.027	0.032	0.034	0.037	0.039	0.041	0.043	0.044	0.044	0.046
	-0.052	0.002	0.002	0.002	0.002	0.003	0.003	0.004	0.004	0.005	0.005	0.006

Realized Kernel Estimator.

	9:45	10:00	10:15	10:30	10:45	11:00	11:15	11:30	11:45	12:00	12:15	12:30
AXP	*-0.559*	0.029	0.035	0.036	0.041	0.044	0.046	0.048	0.049	0.051	0.052	0.053
	-0.511	0.002	0.003	0.003	0.005	0.005	0.006	0.007	0.007	0.008	0.009	0.010
BA	*-0.390*	0.046	0.057	0.065	0.070	0.073	0.076	0.077	0.080	0.081	0.082	0.085
	-0.319	0.004	0.005	0.007	0.008	0.009	0.011	0.012	0.013	0.013	0.014	0.016
BAC	*-0.749*	0.014	0.019	0.024	0.029	0.032	0.034	0.038	0.040	0.041	0.042	0.044
	-0.736	0.005	0.008	0.011	0.015	0.017	0.020	0.024	0.025	0.027	0.028	0.030
CSCO	*-0.350*	0.022	0.030	0.037	0.040	0.043	0.046	0.047	0.048	0.049	0.050	0.050
	-0.306	0.002	0.002	0.005	0.005	0.007	0.008	0.008	0.009	0.009	0.010	0.009
CVX	*-0.191*	0.020	0.027	0.032	0.036	0.040	0.043	0.045	0.046	0.048	0.048	0.049
	-0.145	0.000	0.001	0.002	0.003	0.004	0.005	0.005	0.006	0.007	0.007	0.008
GE	*-0.312*	0.024	0.027	0.030	0.033	0.037	0.039	0.041	0.042	0.044	0.045	0.048
	-0.279	0.006	0.006	0.007	0.008	0.009	0.010	0.012	0.012	0.013	0.013	0.015
INTC	*-0.344*	0.008	0.011	0.014	0.016	0.018	0.020	0.021	0.023	0.024	0.025	0.025
	-0.324	0.002	0.001	0.002	0.002	0.004	0.004	0.005	0.006	0.006	0.007	0.007
JPM	*-0.590*	0.015	0.020	0.022	0.025	0.026	0.028	0.029	0.030	0.030	0.030	0.031
	-0.571	0.005	0.007	0.008	0.010	0.011	0.012	0.013	0.014	0.014	0.014	0.015
PFE	*-0.162*	0.016	0.025	0.028	0.032	0.036	0.039	0.041	0.042	0.044	0.044	0.045
	-0.123	0.001	0.003	0.003	0.004	0.006	0.007	0.008	0.009	0.009	0.010	0.010
UTX	*-0.208*	0.037	0.046	0.053	0.058	0.061	0.064	0.065	0.068	0.070	0.071	0.072
	-0.135	0.000	0.000	0.001	0.002	0.003	0.004	0.004	0.005	0.006	0.006	0.007
WMT	*0.050*	0.023	0.030	0.038	0.042	0.044	0.047	0.050	0.052	0.052	0.053	0.053
	0.093	0.002	0.003	0.005	0.006	0.007	0.008	0.010	0.011	0.011	0.011	0.012
XOM	*-0.096*	0.017	0.019	0.024	0.026	0.029	0.031	0.033	0.035	0.037	0.037	0.039
	-0.056	0.000	0.000	0.001	0.001	0.001	0.002	0.002	0.003	0.004	0.004	0.005

Markov Chain Estimator.

Chapter 4

A Markov Chain Estimator of Multivariate Volatility

Joint with Peter Hansen, Guillaume Horel and Asger Lunde

4.1 Introduction

This paper introduces the Markov chain estimator of multivariate volatility. Our analysis builds on the results by Hansen and Horel (2009) who proposed the univariate Markov chain estimator. The multivariate extension poses new challenges related to asynchronicity and the potential need to enforce the estimator to be positive semidefinite.

The availability of high-frequency financial data has made it possible to estimate volatility over relatively short periods of time, such as an hour or a day. The main obstacle in obtaining precise estimators is the fact that high-frequency returns do not conform with conventional no-arbitrage models. The reason is that there is a great deal of autocorrelation in tick-by-tick returns. The apparent contradiction can be explained by market microstructure noise, which gives rise to the notion that the observed price is a noisy measure of the efficient price. In this paper, we introduce a multivariate volatility estimator that is built on the theory of Markov chains. The estimator utilizes the discreteness of high-frequency data, and the framework implicitly permits a high degree of serial dependence in the noise as well as dependence between the efficient price and the noise.

The use of high-frequency data for volatility estimation has been very active over the past two decades, since Andersen and Bollerslev (1998b) used the realized variance to evaluate GARCH models. The realized variance is simply the sum of squared intraday returns, and its properties were detailed in Barndorff-Nielsen and Shephard (2002a), for the case where the semimartingale is observed without noise, which was extended to the multivariate context in Barndorff-Nielsen and Shephard (2004). The noise in high-frequency returns motivated a number of robust estimators, including the two-scale estimator by Zhang et al. (2005), the realized kernels by Barndorff-Nielsen et al. (2008), and the pre-average estimator by Jacod et al. (2009). Empirical features of the market microstructure noise were detailed in Hansen and Lunde (2006b), which documented that the noise is both serially dependent and endogenous, in the sense that there is dependence between the underlying semimartingale and the noise. These empirical features motivated the development of the multivariate realized kernel in Barndorff-Nielsen et al. (2011a), which is an estimator that permits the noise to have both of these features.

An attractive feature of the Markov framework is that serially dependent and endogenous noise is a natural part of the framework. Moreover, the Markov chain estimator is simple to compute and the same is the case for the estimator of its asymptotic variance. It only takes basic matrix operations to compute the estimator and its confidence intervals.

To illustrate our estimator consider the case with two assets. The bivariate sequence of high-frequency returns is denoted by $\{\Delta X_t\}_{i=1}^n$, and we define the $S \times 2$ matrix $\mathbf{x}$, where S is the number of states for ΔX_t, and each row of $\mathbf{x}$ corresponds to a possible realization of ΔX_t. For instance, the s-th row of $\mathbf{x}$ may equal $x_{s,\cdot} = (2, -1)$ that is the state where the first asset increased by 2 units, while the second asset went down by one unit. The $S \times S$ transition matrix, P, for a Markov chain of order $k = 1$ is given by

$$P_{r,s} = \Pr\left(\Delta X_{t+1} = x_{s,\cdot} | \Delta X_t = x_{r,\cdot}\right), \qquad r, s = 1, \ldots, S,$$

and its stationary distribution, $\pi = (\pi_1, \ldots, \pi_S)'$, is characterized by $\pi' P = \pi'$. We define $\Lambda_\pi = \text{diag}(\pi_1, \ldots, \pi_S)$ and the *fundamental matrix* $Z = (I - P + \Pi)^{-1}$ where $\Pi = \iota\pi'$ with $\iota = (1, \ldots, 1)' \in \mathbb{R}^S$. From the maximum likelihood estimator of P we deduced estimates of π and Z, denoted $\hat{\pi}$ and $\hat{Z}$, see Section 3 for details. The multivariate Markov chain estimator is given by

$$\text{MC} = nD^{-1}\left\{\mathbf{x}'(\Lambda_{\hat{\pi}}\hat{Z} + \hat{Z}'\Lambda_{\hat{\pi}} - \hat{\pi}\hat{\pi}' - \Lambda_{\hat{\pi}})\mathbf{x}\right\}D^{-1},$$

where $D = \text{diag}\left(\delta_1, \delta_2\right)$ and $\delta_j^2 = n^{-1}\sum_{t=1}^n X_{j,t}^2$ is the sample average of the squared price of the j-th asset, $j = 1, 2$. The expression inside the curly brackets is the estimator of the long-run variance of a finite Markov chain, see Hansen and Horel (2014). The scaling involving D, is a transformation needed for the estimator to be an estimator of volatility of logarithmic prices. The scaling with the sample size, n, relates to the local-to-zero asymptotic scheme that arises under in-fill asymptotics.

Hansen and Horel (2009) showed that filtering can resolve the problems caused by market microstructure noise under weak assumptions that essentially amounts to the noise process to be ergodic with finite first moment. This result is theoretical in nature, because the ideal filter requires knowledge about the data generating process. In order to turn the theoretical filtering result into an actual estimator, one needs to adopt a statistical model, and our approach is to model the increments of the process with a Markov chain model, which is a natural starting point given the discrete nature of high-frequency data.

The discreteness of financial data is a product of the so-called *tick size*, which defines the coarseness of the grid that prices are confined to. For example, the tick-size is currently 1 cent for most of the stocks that are listed on the New York Stock Exchange. The implication is that all transaction and quoted prices are in whole cents. The Markov estimator can also be applied to time series that do not live on a grid, by forcing the process onto a grid. While this will introduce rounding error, it will not affect the long-run variance of the process. Delattre and Jacod (1997) studied the effect of rounding on realized variances for a standard Brownian motion, and Li and Mykland (2006) extended this analysis to log-normal diffusions.

The present paper adds to a growing literature on volatility estimation using high-frequency data, dating back to Zhou (1996), Zhou (1998). Well known estimators include the realized variance, see Andersen et al. (2001b) and Barndorff-Nielsen et al. (2008); the two-scale and multi-scale estimators, see Zhang et al. (2005) and Zhang (2006); the realized kernels, see Barndorff-Nielsen et al. (2008), Barndorff-Nielsen et al. (2011a). The finite sample properties of these estimators are analyzed in Bandi and Russell (2005), Bandi and Russell (2008), and the close relation between multi-scale estimators and realized kernels is established in Barndorff-Nielsen et al. (2011b). Other estimators include those based on moving average filtered returns, see Andersen et al. (2001a), Maheu and McCurdy (2002), and Hansen et al. (2008); the range-based estimator, see Christensen and Podolskij (2007); the pre-averaging estimator, see Jacod et al. (2009); the quantile-based estimator Christensen et al. (2008); and the duration-based estimator, see Andersen et al. (2008).

The stochastic properties of market microstructure noise are very important in this context. Estimators that are robust to iid noise can be adversely affected by dependent noise. Hansen and Lunde (2006b) analyzed the empirical features of market microstructure noise and showed that serial dependence and endogenous noise are pronounced in high-frequency stock prices. Endogenous noise refers to the dependence between the noise and the efficient price. A major advantage of the Markov chain estimator is that dependent and endogenous noise is permitted in the framework. In fact, dependent and endogenous noise arises naturally in this context,

see Hansen (2015). Thus estimation and inference are done under a realistic set of assumptions in regard to the noise.

The present paper is an extension of Hansen and Horel (2009) to the multivariate context. This extension posed new challenges that are specific to the multivariate context. For instance, different assets are typically not traded at synchronous times. This non-synchronicity leads to the so-called Epps effect, which manifests itself by a bias towards zero for the realized covariance as the sampling frequency increases. See Renò (2003) for a study of the determinants of the Epps effect. Another issue that may arise in the multivariate context is a need for the estimator to be positive semidefinite, which is not guaranteed by all multivariate estimators. The asynchronicity poses few obstacles for the Markov chain estimator, albeit a large order of the Markov chain, or another remedy, may be needed if an illiquid asset is paired with a liquid asset.

The outline of this paper is as follows. The Markov chain framework is presented in Section 2, and the estimator in Section 3. In Section 4 we present two composite estimators that estimate every element of the matrix separately. In Section 5 we propose a novel projection methods that may be needed to ensure that the composite estimators are positive semidefinite. The properties of the estimators are evaluated in Section 6 with a simulation study, and an empirical application to commodity prices is presented in Section 7.

4.2 The Markov Chain Framework

Let $\{X_t\}$ denote a d-dimensional process, whose returns ΔX_t can take S distinct values in $\mathbb{R}^d$. For notational convenience we take ΔX_t to be a row-vector. The possible states for the k-tuple, $\Delta\mathcal{X}_t = \{\Delta X_{t-k+1}, \ldots, \Delta X_t\}$, are indexed by $s = 1, \ldots, S^k$, where the s-th state corresponds to the case where $\Delta\mathcal{X}_t = \mathbf{x}_s$, which is an $1 \times kd$ vector. See the example below.

We make the following assumption about the increments of the process.

Assumption 1. *The increments $\{\Delta X_t\}_{t=1}^n$ are ergodic and distributed as a homogeneous Markov chain of order $k < \infty$, with $S < \infty$ states.*

The homogeneity assumption is unlikely to be valid in the context of high-frequency data. Fortunately the assumption is not critical for our results, because by increasing the order, k, of the homogeneous Markov chain that is imposed on the high-frequency returns, the resulting estimator becomes robust to inhomogeneity, see Hansen and Horel (2009). This feature of the Markov chain estimator is demonstrated in our simulation study in Section 6.

The transition matrix, P, is given by

$$P_{r,s} = \Pr(\Delta\mathcal{X}_{t+1} = \mathbf{x}_s | \Delta\mathcal{X}_t = \mathbf{x}_r), \qquad \text{for } r, s = 1, \ldots, S^k,$$

and the corresponding *stationary distribution,* $\pi \in \mathbb{R}^{S^k}$, which is unique under Assumption 1, is defined by $\pi' P = \pi'$. The *fundamental matrix* by Kemeny and Snell (1976) is defined by

$$Z = (I - P + \Pi)^{-1},$$

where $\Pi = \iota\pi'$ with $\iota = (1, \ldots, 1)' \in \mathbb{R}^{S^k}$ so that each row of Π is simply π'.

The $S^k \times d$ matrix, f, is defined to be the last d columns of $\mathbf{x}$. So f_s is the value that (the latest observation of) ΔX_t has in state s. (Recall that a state represents a realization of k consecutive returns.) Finally, we define the diagonal matrix $\Lambda_\pi = \operatorname{diag}(\pi_1, \ldots, \pi_{S^k})$.

The following example illustrates the multivariate Markov chain estimation in the case where $d = 2$ and $S = 2$, and $k = 1, 2$.

Example 1. Suppose that we have two assets and that all price changes are up or down by a single unit. If the order of the Markov chain is $k = 1$, then the transition matrix, P, is a 4×4 matrix, and we can define the

state matrix as

$$\mathbf{x} = f = \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ -1 & 1 \\ -1 & -1 \end{pmatrix}.$$

If, instead, the order is $k = 2$, then we have $S^2 = 16$ states, and consequently P will be an 16×16 matrix and f an 16×2 matrix. For instance, we may order the states as below, so that a row of $\mathbf{x}$ corresponds to a state value for $(\Delta X_{t-1}, \Delta X_t)$ and the corresponding row of f will have just the state value for ΔX_t:

$$\mathbf{x} = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & -1 \\ 1 & 1 & -1 & 1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & -1 & -1 & 1 \\ 1 & -1 & -1 & -1 \\ -1 & 1 & 1 & 1 \\ -1 & 1 & 1 & -1 \\ -1 & 1 & -1 & 1 \\ -1 & 1 & -1 & -1 \\ -1 & -1 & 1 & 1 \\ -1 & -1 & 1 & -1 \\ -1 & -1 & -1 & 1 \\ -1 & -1 & -1 & -1 \end{pmatrix} \qquad f = \begin{pmatrix} 1 & 1 \\ 1 & -1 \\ -1 & 1 \\ -1 & -1 \\ 1 & 1 \\ 1 & -1 \\ -1 & 1 \\ -1 & -1 \\ 1 & 1 \\ 1 & -1 \\ -1 & 1 \\ -1 & -1 \\ 1 & 1 \\ 1 & -1 \\ -1 & 1 \\ -1 & -1 \end{pmatrix}.$$

Although the transition matrix is a 16×16 matrix, it has at most four non-zero elements in each row. The reason is that many transitions are impossible. For, instance if $\Delta\mathcal{X}_t = \{(1,1),(1,1)\}$ then the next state will have to be $\{(1,1),(*,*)\}$, and a transition to, $\{(-1,-1),(1,1)\}$ say, is impossible, and thus have probability zero. So the transition matrix, P, will be increasingly sparse as k increases.

The underlying idea of the Markov chain estimator is a martingale decomposition of

$$X_t = Y_t + \mu_t + U_t,$$

where $\{Y_t, \mathcal{F}_t\}$ is a martingale with increments $\Delta Y_t' = e_{s_t}' Zf - e_{s_{t-1}}' PZf$, $\mu_t = t\mu$ with $\mu = \mathrm{E}(\Delta X_t)$, and U_t is a stationary ergodic bounded process.

The asymptotic scheme that will be used in the present context is the following:

$$f = n^{-1/2}\xi, \qquad \text{with} \quad \xi \in \mathbb{R}^{S^k \times d} \quad \text{fixed.} \tag{4.1}$$

This local-to-zero asymptotic scheme is similar to those used in Delattre and Jacod (1997) and Li and Mykland (2006), and is natural under in-fill asymptotics. In the present context, it guarantees almost sure convergence of the estimator.

Under this scheme, it follows from Hansen and Horel (2014) (and the ergodic theorem) that

Proposition 1. *Suppose that Assumption 1 holds, then under the asymptotic scheme (4.1), we have*

$$\sum_{t=1}^{n} \Delta Y_t \Delta Y_t' \xrightarrow{a.s.} \xi' Z'(\Lambda_\pi - P'\Lambda_\pi P)Z\xi = \xi'(\Lambda_\pi Z + Z'\Lambda_\pi - \pi\pi' - \Lambda_\pi)\xi,$$

as $n \to \infty$.

Proof. By Assumption 1 it follows that $\vartheta_t' = e_{s_t}' Z\xi - e_{s_{t-1}}' PZ\xi$ is an ergodic Markov chain (of order $k+1$) with

$$\mathrm{E}\vartheta_t\vartheta_t' = \xi' Z'(\Lambda_\pi - P'\Lambda_\pi P)Z\xi = \xi'(\Lambda_\pi Z + Z'\Lambda_\pi - \pi\pi' - \Lambda_\pi)\xi,$$

where the first identity follows from Hansen (2015, theorem 2) and the second from Hansen and Horel (2014, lemma 1). By the ergodic theorem it follows that $\frac{1}{n}\sum_{t=1}^{n}\vartheta_t\vartheta_t' = \sum_{t=1}^{n}\Delta Y_t\Delta Y_t'$ converges to $\xi'(\Lambda_\pi Z + Z'\Lambda_\pi - \pi\pi' - \Lambda_\pi)\xi$ almost surely (and in mean). □

An implication of Proposition 1 is that

$$\Sigma^{\#} = \xi'(\Lambda_\pi Z + Z'\Lambda_\pi - \pi\pi' - \Lambda_\pi)\xi,$$

is the quadratic variation of the martingale component. The same quantity is also the long-run variance of $\Delta X_t'$ in the sense that

$$\mathrm{var}(X_n' - X_0') = \mathrm{var}(\sum_{t=1}^{n}\Delta X_t') \to \Sigma^{\#}, \qquad \text{as} \quad n \to \infty.$$

There are different ways to construct a multivariate volatility estimator using Markov chain methods, and we shall present three distinct estimators and highlight each of their advantages.

4.3 The Markov Estimator

Let $\hat{P}$ be the maximum likelihood estimator of P and let $\hat{\pi}$ be its corresponding eigenvector, $\hat{\pi}'\hat{P} = \hat{\pi}'$. Furthermore, let $\hat{\Pi} = \iota\hat{\pi}'$ and $\hat{Z} = (I - \hat{P} + \hat{\Pi})^{-1}$. The expression for the long-run variance of the Markov chain motivates the estimator

$$\mathrm{MC}^{\#} = nf'(\Lambda_{\hat{\pi}}\hat{Z} + \hat{Z}'\Lambda_{\hat{\pi}} - \hat{\pi}\hat{\pi}' - \Lambda_{\hat{\pi}})f,$$

for which we have the following asymptotic distribution.

Proposition 2. *Suppose that Assumption 1 holds, then under the asymptotic scheme (4.1), we have*

$$n^{1/2}(\mathrm{MC}^{\#} - \Sigma^{\#}) \xrightarrow{d} N(0, \Omega),$$

where the asymptotic covariance between the (i,j)th and (l,m)th elements is

$$\Omega_{ij,kl} = \sum_{r,s,v}[V(r)]_{s,v}[\Xi(r,s)]_{i,j}[\Xi(r,v)]_{l,m}, \tag{4.2}$$

with $V(s) = \frac{1}{\pi_r}(\Lambda_{e_r'P} - P'e_re_r'P)$ and

$$\begin{aligned}\Xi(r,s) \;=\; & \pi_r\xi'Z'(\Lambda_{z_s} - P'\Lambda_{z_s}P)Z\xi \\ & +\pi_r\xi'(\pi\pi' - 2\pi z_s' - 2z_s\pi')\xi + \xi'[\Lambda_\pi Ze_rz_s' + z_se_r'Z'\Lambda_\pi]\xi,\end{aligned}$$

and where $z_s' = e_s'Z$ is the s-th row of Z.

Proof. Follows from Hansen and Horel (2014, theorem 2) by adapting their expressions (substitute ξ for f and $\xi'\pi$ for μ). □

Remark. We note that in the univariate case, $\mathrm{MC}^{\#}$ simplifies to $nf'\Lambda_{\hat{\pi}}(2\hat{Z} - \hat{\Pi} - I)f$, which was the estimator proposed by Hansen and Horel (2009).

4.3.1 Volatility of logarithmic prices

The object of interest is, typically, the volatility of log-prices, rather than $\Sigma^{\#}$, which is the volatility of the price process in levels. An exact estimator can be obtained with the Markov framework, by first extracting the

Martingale component of X_t, however, for the univariate case Hansen and Horel (2009) show that the following estimator,

$$\mathrm{MC} = \frac{\mathrm{MC}^{\#}}{\frac{1}{n}\sum_{t=1}^{n} X_{T_t}^2},$$

is virtually identical to the realized variance of the filtered logarithmic prices that are deduced from the estimated Markov chain. The resulting approximate estimator of the quadratic variation has several advantages, such as computational simplicity. In the present multivariate context, we adopt the following estimator:

$$\mathrm{MC} = D^{-1}\mathrm{MC}^{\#}D^{-1}, \tag{4.3}$$

with $D = \mathrm{diag}\,(\delta_1, \ldots, \delta_d)$ and $\delta_j^2 = n^{-1}\sum_{t=1}^{n} X_{j,t}^2$ $j = 1, \ldots, d$. Our simulation in Section 6 shows that this approximate estimator is more accurate than other realized measures.

Alternatively one could use the estimator $\mathrm{MC}_{i,j}^{\#}/\frac{1}{n}\sum_{t=1}^{n} X_{i,t}X_{j,t}$, for $i, j = 1, \ldots, d$, but we prefer (4.3) because positively definiteness of $\mathrm{MC}^{\#}$ is passed onto MC, and in practice $\delta_i\delta_j \simeq \frac{1}{n}\sum_{t=1}^{n} X_{i,t}X_{j,t}$ because the prices do not vary much in relation to their average level over the estimation window, which is typically a trading day.

4.4 Composite Markov Estimators

The number of possible states increases exponentially with the dimension of the process, d. Consequently, the dimension of P can become unmanageable even with moderate values of S, k, and d. For instance, with $d = 10$ assets, and price changes ranging from -4 to 4 cents, $S = 9$, and a Markov chain of order $k = 2$, the transition matrix would be $(S^d)^k \times (S^d)^k = 9^{20} \times 9^{20}$, which is impractical.

As an alternative, one can construct a composite estimator, that combines lower dimensional Markov estimator, which is in the spirit of Hautsch et al. (2012) and Lunde et al. (2014). In this section we consider two such estimators. The first is constructed from univariate estimators, using a simple transformation for the estimation of covariances. The second estimator is constructed from bivariate Markov estimators, which has the advantage that standard errors of each element will be readily available. We will make use of these standard errors in the next section.

4.4.1 The 1-Composite Markov Estimator

In this section we introduce a composite estimator that is based on univariate Markov estimators. The identity

$$\mathrm{cov}(X, Y) = \frac{\mathrm{var}(X + Y) - \mathrm{var}(X - Y)}{4}$$

motivates the estimator

$$\mathrm{MC}_{i,j}^{\#_1} = \tfrac{1}{4}(\mathrm{MC}_{X_i+X_j}^{\#} - \mathrm{MC}_{X_i-X_j}^{\#}),$$

where $\mathrm{MC}_{X_i+X_j}^{\#}$ and $\mathrm{MC}_{X_i-X_j}^{\#}$ are the univariate Markov chain estimator, applied to the time series $X_{i,t}+X_{j,t}$ and $X_{i,t}-X_{j,t}$, respectively. Note that the diagonal terms, $i = j$, simplifies to $\frac{1}{4}\mathrm{MC}_{2X_i}^{\#} = \mathrm{MC}_{X_i}^{\#}$. This approach to polarization-based estimation of the covariance is well known. In the context of high-frequency data it was first used in Horel (2007, section 3.6.1) who also explored related identities. More recently it has been used in Ait-Sahalia et al. (2010).

The 1-Composite estimator is mapped into estimators of the volatility of log-returns using the same diagonal matrix, D, as in (4.3), thus $\mathrm{MC}^1 = D^{-1}\mathrm{MC}^{\#_1}D^{-1}$.

4.4.1.1 Pre-Scaling

If one seeks to estimate the covariance of two assets, whose increments are on different grid sizes, it can be advantageous to use differentiated scaling of the assets, specifically

$$\operatorname{cov}(X,Y) = \frac{\operatorname{var}(aX+bY) - \operatorname{var}(aX-bY)}{4ab},$$

where a and b are constants. This can, in some cases, greatly reduce the number of states, which is computationally advantageous.

4.4.2 The 2-Composite Markov Estimator

In this subsection we introduce a composite estimator that uses bivariate $\mathrm{MC}^{\#}$ estimates. For all pairs of assets we compute the correlation along with an estimate of its asymptotic variance, which will be used in the next section.

We simply estimate the bivariate Markov process $(X_{i,t}, X_{j,t})'$, and obtain the estimator of Section 3, $\mathrm{MC}^{\#}$, which is a 2×2 matrix. The covariance terms we seek is the lower-left (or upper-right) element

$$\mathrm{MC}^{\#_2}_{i,j} = \begin{cases} \mathrm{MC}^{\#}(X_i) & \text{if} \quad i = j, \\ \mathrm{MC}^{\#}_{1,2}(X_i, X_j) & \text{if} \quad i \neq j, \end{cases}$$

where $\mathrm{MC}^{\#}_{1,2}(X_i, X_j)$ is the upper right element of the 2×2 matrix $\mathrm{MC}^{\#}$, for the bivariate process, (X_i, X_j). In contrast to the covariance estimated with the 1-composite estimator, the standard error of $\mathrm{MC}^{\#}_{1,2}(X_i, X_j)$ is readily available from (4.2).

Analogous to the other estimators, the 2-composite estimator is mapped into estimators of the volatility of log-returns with $\mathrm{MC}^2 = D^{-1}\mathrm{MC}^{\#_2}D^{-1}$.

4.4.3 Advantages and Drawbacks of Composite Estimators

The advantages of the composite estimators are threefold.

- Computational: The state space for a univariate series is smaller than that of a multivariate.
- Dimension: Enables the construction of covariance matrices of any dimension, whereas the multivariate approach is limited to relatively low dimensions.
- No need to synchronize the observation times for each of the asset, e.g. by refresh time, see Barndorff-Nielsen et al. (2011a).

The drawbacks of the composite estimators include:

- Positive semidefinite estimate is not guaranteed
- Estimate of asymptotic variance is not readily available.

The dimension of the transition matrix (and fundamental matrix) increases rapidly with the dimension of the process d, and at some point it becomes computationally impossible to manipulate the relevant expressions that are needed for the computation of the Markov estimator. In our empirical analysis with $k = 5$, the dimension of P was about 500-1000 for $d = 1$, about 3000-5000 for $d = 2$, and about 8000-10,000 for $d = 3$. The problem with non-psd appears to be relatively rare in practice when d is small. We have only seen one case where a 5×5 estimate was non-psd estimate. The occurrence is more common in higher dimensions. Of the 251 14×14 estimators we obtained for 2013, 14 of them were non-psd.

4.5 Enforcing Positivity

While $\text{MC}^{\#}$ is a quadratic form that yields a positive semidefinite estimator, there is no reason to expect that the composite estimators, $\text{MC}^{\#_1}$ and $\text{MC}^{\#_2}$, will be positive semidefinite (PSD) in finite samples. This problem is often encountered in estimation of high-dimensional variance-covariance matrices.

One can project a non-PSD estimate, by solving the following semi-definite program for the variable Σ

$$\min_{\Sigma} \|\Sigma - A\|_{\text{Fro}} \quad \text{subject to} \quad \Sigma \geq 0. \tag{4.4}$$

The solution can be found efficiently by computing the spectral decomposition of the matrix A, and drop all negative eigenvalues, i.e. map the symmetric matrix, $A = Q\text{diag}(\lambda_1, \ldots, \lambda_d)Q'$ into $Q\text{diag}(\lambda_1^+, \ldots, \lambda_d^+)Q'$, where $\lambda_1, \ldots, \lambda_d$ are the eigenvalues of A and $x^+ = \max(x, 0)$. Such an estimator will, due to the zero eigenvalues, be on the boundary of the space of psd matrices, which motivated Ledoit et al. (2003) to impose an additional constraint, $\text{diag}(\Sigma) = \text{diag}(A)$.

In this paper we propose a novel projection that takes advantage of standard errors of the individual elements of the matrix A when these are available. Thus let ω_{ij} be (an estimate of) the standard errors of A_{ij}. Then we solve the following program

$$\min_{\Sigma} \sum_{i,j=1}^{d} \left(\frac{\Sigma_{ij} - A_{ij}}{\omega_{ij}} \right)^2 \quad \text{subject to} \quad \Sigma \geq 0. \tag{4.5}$$

The solution can be obtained using semidefinite programming solvers that are readily available, including the `cvx` software for Matlab by Grant et al. (2014). The optimization problem can be supplemented with the constraint $\text{diag}(\Sigma) = \text{diag}(A)$, which would produce a constrained estimate with strictly positive eigenvalues, except in pathological cases, e.g. if A is psd with zero eigenvalues to begin with.

The projection in (4.5) is appealing because it attempts to influence accurately measured elements of A less than those that are relatively inaccurate. An even more appealing projection along these lines would also account for correlations across elements. In the present context, such cross correlations are only available for the estimator $\text{MC}^{\#}$. However, since this estimator, $\text{MC}^{\#}$, is psd per construction, there is no need for a projection of this estimator.

4.6 Simulation

In this section we compare the 1-composite Markov estimator against some benchmark. Diagonal elements are compared with the realized variance (RV) and the realized kernel (RK). Off-diagonal elements are compared with the realized covariance (RC).

4.6.1 Efficient Price

Our simulations are based on two designs for the latent price process, Y_t. In the first design, Y_t is simply sampled from a Brownian motion with constant volatility. In the second design, Y_t is drawn from a stochastic volatility model, which is known as the Dothan model in the literature on interest rates, similar to that used in Barndorff-Nielsen et al. (2008). Specifically we simulate

$$\log Y_{i,t} = \log Y_{i,t} + \sigma_{i,t} V_{i,t}, \qquad i = 1, 2,$$

where $V_{i,t} = \gamma Z_{i,t} + \sqrt{1-\gamma^2} W_{i,t}$ with $(Z_{1,t}, Z_{2,t}, W_{1,t}, W_{2,t})$ being iid Gaussian, all having unit variance and zero correlation, with the exception that $\text{cov}(W_{1,t}, W_{2,t}) = \rho$.

In the design with stochastic volatility, the volatility, $\sigma_{i,t}$, correlates with $Z_{i,t}$, so that γ controls the leverage

effect of the volatility on the stock prices. Specifically,

$$\sigma_{i,t} = \sqrt{\Delta}\left\{\exp\left(\beta_0 + \beta_1 \tau_{i,t}\right)\right\},$$

where $\tau_{i,t} = \exp(\alpha\Delta)\tau_{i,t-1} + \sqrt{\frac{\exp(2\alpha\Delta)-1}{2\alpha}} Z_{i,t}$, with $\tau_{i,1}$ drawn from its unconditional distribution, and $\Delta = \frac{1}{N}$ with $N = 23,400$. Additional details about the specification is given in the Appendix.

The values of the parameters in both designs are summarized in Table 4.1.

	β_0	β_1	α	ρ	γ
Constant Volatility	0	0	–	-0.3	0
Stochastic Volatility	-0.3125	0.125	-0.025	-0.3	0.5

Table 4.1: Parameters values for simulating the efficient price process, Y_t.

4.6.2 Noise

We will use two specifications for the noise. The first is pure rounding noise, so that

$$X_t = \delta[Y_t/\delta],$$

where $[a]$ is the rounding of a to the integers so that the parameter δ controls the coarseness of the rounding error.

The second specification has an additive noise component in addition to the rounding error, specifically

$$X_t = \delta[(Y_t + U_t)/\delta],$$

where U_t are iid and uniformly distributed. The idea is that it would more closely resemble the bid ask bounce (due to the additional jitter introduced by U_t, we will either round up or down).

In our simulation study we use $\delta = 0.01$ to emulate rounding errors to a grid, and the noise is $U_t \sim \text{iid}U[-\frac{1}{3}, \frac{1}{3}]$ which adds additional (mean-reverting) jitter to the returns.

In Figure 4.1 we show an example of a realization of the process with stochastic volatility using the design in Table 4.1. The upper panel has Y_t and X_t, where the latter is clearly identified by it being confined to the grid values. The lower panel displays the corresponding volatility process, specifically we plot $\sigma^2(t) = \sigma^2_{i,t}/\Delta$.

4.6.3 Estimators and Tuning Parameters

We consider the realized variance computed with different sampling frequencies. To imitate a 24 hour period with second-by-second price observations, we generate 23,400 noisy high-frequency returns in each simulation.

The realized variance (RV) and the realized covariance (RC) is computed using different sampling frequencies. The choice of sampling frequency entails a bias-variance trade-off, because the bias arising from the noise is most pronounced at high sampling frequencies, while the variance of the estimator increases as the sampling frequency is lowered. Thus for the RV and the RC we sample every H-th price observation where $H \in \{1, 3, 5, 10, 15, 30, 60, 120, 240\}$.

The multivariate realized kernel (MRK) follows the implementation in Barndorff-Nielsen et al. (2011a), which is based on the Parzen kernel, and an automatic selection of the bandwidth parameter. This estimator is also applied to high-frequency returns based on the various sample frequencies. The MRK should, in principle, be most accurate when based on returns sampled at the highest frequency, $H = 1$.

The tuning parameter for the Markov chain estimator is the order of the Markov chain, k, and we apply this estimator for $k \in \{1, \ldots, 5\}$.

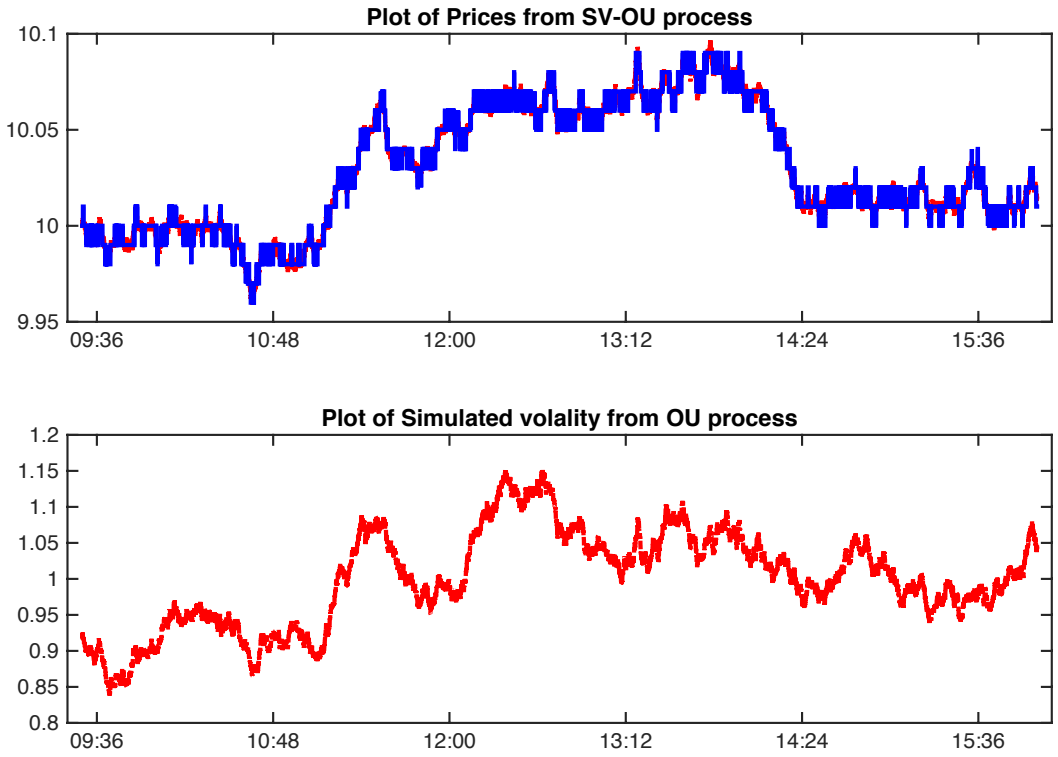

Figure 4.1: A typical sample path of the simulated stochastic volatility process. The upper panel displays the price process, Y_t, and the observed process, X_t, that is subject to noise and rounding error. The lower panel displays the corresponding volatility process, $\sigma^2(t)$.

4.6.4 Simulation Results

We report bias and the root mean squared error (RMSE) for each of the estimators using the various choices for their respective tuning parameters. The results are based on 10,000 simulations. The results are presented in Table 4.2 for the case with constant volatility and stochastic volatility.

Consider first the case with constant volatility from panel A in Table 4.2. With pure rounding error we note that the Markov chain estimator tend to outperform both the kernel estimator and the realized variance in terms of the mean squared error. Similarly for the covariance, the MC 1-composite estimator dominates the RC and performs on par with the MRK. The Markov estimator is somewhat insensitive to the choice of k, so even with a non-optimal choice for k, the Markov estimator is fairly accurate. The MRK is similarly insensitive to the choice for H. In contrast, the RV and the RC are very sensitive to the choice of H, and suffer from large biases when H is small.

Turning to the case with both additive noise and rounding error. This design generates increments with rather different features. While $k = 1$ was the optimal choice with rounding error, the best configuration is now $k = 3$ or $k = 4$. The RV performs even worse in this design, the RC just as bad as in the previous design, whereas MRK performs as well as in the previous design, and is on par with the Markov estimator. This comparison is again made with hindsight as assume that relatively good choices of tuning parameters, for k and H, respectively, are used. For the covariance, we observe that the RMSE of the Markov estimator is predominantly driven by a bias.

Next we turn to the result from panel B in Table 4.2 which is for the case with stochastic volatility. The Markov chain estimator is based on fitting a homogeneous Markov chain to the observed increment. For this reason it might be expected that the Markov estimator is not well suited for the design with time varying volatility, see Figure 4.1. However, even in the case with stochastic volatility that induced an inhomogeneous model for the increments, we see that MC performs well. The RMSEs are, as expected, a bit larger. Interestingly, it is the design with pure rounding errors that results in the largest RMSEs. Both the Markov estimator and the MRK appear to benefit from the additional layer of noise that is added prior to the rounding error.

4.7 Empirical Analysis

4.7.1 Data Description

We apply the Markov chain estimator to high-frequency commodity prices, that have previously been used in Christoffersen et al. (2014). We confine our empirical analysis to 2013 data and consider in our study high frequency data for 14 assets. The 14 assets include the exchange traded fund, SPY, that tracks the S&P 500 index, and 13 commodity futures. We refer to Christoffersen et al. (2014) for detailed information about the data, including the procedures used for cleaning the high frequency data for outliers and other anomalies.

Of the 15 commodities analyzed in Christoffersen et al. (2014), we drop two of these series for computational reasons. Specifically, we dropped "Heating Oil" (HO) because it has an unusually large number of distinct second-to-second price increments and "Feeder Cattle" (FC) because it is substantially less liquid compared with the other commodities. Thus, in addition to the SPY, we use the following 13 commodities in our empirical analysis: Crude Light (CL), Natural Gas (NG), Gold (GC), Silver (SV), Copper (HG), Live Cattle (LC), Lean Hogs (LH), Coffee (KC), Sugar (SB), Cotton (CT), Corn (CN), Soybeans (SY) and Wheat (WC).

We exclusively apply the Markov estimators to high-frequency data from the time interval 10:00-14:00 eastern standard time, because all assets are actively traded in this period. The high-frequency prices for eight selected assets for March 18th, 2013 are displayed in Figure 4.2.

4.7.2 Empirical Results

First we present detailed results for March 18th, 2013 (to celebrate the occasion for writing this paper). Daily estimates (for the 10:00-14:00 interval) for all trading days in 2013 will be presented in figures.

4.7.2.1 Daily Estimates for March 18, 2013

In Table 4.4 we present five estimators of the volatility matrix for five assets. There are relatively large discrepancies between the two realized variances, which may be due to sampling error or market microstructure noise. The Markov estimators are largely in agreement about the correlations, but the full estimator yields a smaller estimate of the diagonal elements in some cases. This may be caused by the estimator being somewhat unreliable, as it is based on $n = 8,700$ observations and the underlying transition matrix is an $8,600 \times 8,600$ matrix in this case. Further research is needed to characterize the limitations of the full estimator in practice.

In Tables 4.5 and 4.6 we present estimates of the full 14×14 matrix. The realized variances are in Table 4.5 and the two composite Markov estimators are in Table 4.6. Joint estimation of the full 14×14 covariance matrix is not expected to be precise because the number of observed states (and the dimensionality of the transition matrix) is equal to the number of observations in that case. As in the previous example, the two realized variance estimators produce quite different values whereas the composite Markov estimators produce rather similar results. In general, signs and magnitudes of the elements of Markov covariance matrices are largely in agreement with those of the realized variances.

4.7.2.2 Estimates for Pooled March Data

Finally we have pooled the high-frequency data for all of March and estimated the 14×14 matrix that reflects the volatility in March, 2013 that occurred during the 10:00AM to 2:00PM trading periods. With 20 trading days in March, 2013, this adds up to 80 hours of high-frequency data. Precision is expected to improve with the larger sample size, although the dimensions of the underlying transition matrices are expected to increase as a larger number of states and transitions will be observed in a larger sample, and the latter can potentially cause computational difficulties. For the 2-composite estimator with $k = 4$ we observed between 15000 and 30000 distinct states in the pooled data set. Another challenges for the Markov estimator in the pooled sample is that a larger degree of inhomogeneity may be expected. Hansen and Horel (2009) showed that an inhomogeneous

Markov process can be approximated by a homogeneous Markov process, by increasing the order of the Markov chain. So a larger k may be needed in the pooled data, which also poses computational challenges.

In Tables 4.7 and 4.8 we report estimates for the 14×14 covariance matrix computed with the realized variances and the two composite Markov estimators. In contrast to the data for March 18th, 2013, the realized variances are largely in agreement for the pooled data. Albeit some differences are observed between the 5- and 10-minute realized variances. The composite Markov estimators are in disagreement in some cases, which we attributed to the different order of the Markov chain that were used. The 1-composite estimator was computed with $k = 5$ whereas 2-composite was estimated with $k = 3$, for computational reasons. Naturally, one could use a higher order to compute the diagonal elements, but we used the same order for all entries of the 2-composite estimator to illustrate the differences that arise in this case. The 1-composite Markov estimator produces estimates that are generally in agreement with those of the realized variance, both in terms of magnitude and signs of covariances.

4.7.2.3 Daily Estimates for 2013

We have estimated variances and covariances for the 10:00-14:00 interval for all trading days during 2013. Some selected series are presented in Figures 4.3 and 4.4.

In Figure 4.3 we plot annualized volatilities for SPY, Crude Oil, Gold, and Wheat based on the Markov estimator with $k = 5$, and these are benchmarked with those of the realized variance with 10-minute sampling. The estimates are quite similar, both for the very liquid assets, SPY, Crude Oil, and Gold, and the relatively illiquid assets, Wheat, whose high-frequency data have pronounced bid-ask bounces.

In Figure 4.4 we plot daily estimates of the correlations for both the 1- and 2-composite Markov estimators. These are benchmarked with the realized correlations based on 10-minute sampling. We observed that Gold/Silver are highly correlated and its correlation is highly persistent over the year 2013. More moderate correlations are observed for SPY/Crude Oil and Soybeans/Wheat, and these series exhibit a higher degree of time-variation. In the case of SPY/Gold we observe a less stable correlation that changes sign several times during the year. The general patterns are successfully captured by the composite Markov estimators, and while the realized correlation is in agreement about the general trends, it exhibits far more day-to-day variation which suggests that it is less accurate. The smooth and persistent behavior of the Markov estimators may be attributed to these estimators being more accurate.

4.8 Conclusion

In this paper we have proposed a multivariate volatility estimator that is based on the theory of finite Markov chains. The Markov chain estimator takes advantage of the fact that high-frequency prices are confined to a grid. This is the first robust multivariate estimator for which standard errors are readily available. Previous estimators include the multivariate realized kernel estimator, whose standard error also requires an estimate of the long-run variance of the noise, which is difficult to estimate because the noise is, in practice, small, serially dependent, and endogenous. The multivariate kernel estimator (MRK) converges at rate $n^{1/5}$. In contrast, the Markov estimator converges at rate $n^{1/2}$ owing to the specification assumed for the high-frequency data. These rates are, however, not directly comparable for practical situation, as the order of the Markov chain may be required to increase with n, in order to accommodate inhomogeneity resulting from time-varying volatility. Our simulation design suggests that the Markov estimator and the MRK performs similarly in practice, so the major advantage of the Markov estimator is the readily available standard error.

The estimator performs well in a simulation design, and is relatively insensitive to the choice of the order of the Markov chain, k, which is the tuning parameter that must be chosen in practice.

A potential limitation of the estimator is the high-dimensional objects that the estimator is computed from. For the full estimator the dimension can be as large as $(S^d)^k$, where S is the number of primitive states for the individual series, d is the dimension of the process and k is the order of the Markov chain. The dimension will

typically be much smaller in practice because many states are not observed in a given sample, and the transition matrix will be very sparse, because most transitions between states are unobserved. So there is a need to further analyze the finite sample properties of the full Markov estimator, and to characterize its limitations.

The two composite Markov estimators alleviate the challenges with high dimensional objects, but may require a projection to guarantee a positive semidefinite estimate. For this purpose we have proposed a novel projection that makes use of the standard errors of the elements of the matrix being projected. Since these are readily available for the 2-composite estimator it is appealing to incorporate these, so that a projection leaves accurately estimated elements relatively unchanged.

The empirical analysis of commodity prices illustrated the three Markov estimators, and benchmarked them against conventional realized variances. The estimates were largely in agreement, but the Markov estimators fare particularly well with regards to estimating correlations. While the time series of daily correlation estimates based on the realized variance were somewhat erratic, those of the Markov estimators were more stable.

	Pure Rounding Error				*Noise and Rounding Error*			
	Variance		Covariance		Variance		Covariance	
	MC		MC		MC		MC	
k	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse
1	0.002	0.109	-0.008	0.075	0.281	0.302	-0.066	0.115
2	0.004	0.144	-0.053	0.093	0.140	0.173	-0.137	0.151
3	0.007	0.174	-0.030	0.098	0.056	0.131	-0.078	0.104
4	0.005	0.202	-0.020	0.107	0.025	0.130	-0.076	0.103
5	0.003	0.221	-0.010	0.120	0.011	0.141	-0.049	0.092

	RV		MRK		RC		MRK		RV		MRK		RC		MRK	
H	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse
1	3.752	3.763	0.140	0.182	-0.450	0.451	-0.008	0.074	9.697	9.708	0.112	0.159	-0.451	0.456	-0.006	0.075
3	3.445	3.455	0.058	0.136	-0.373	0.377	-0.007	0.088	8.086	8.095	0.044	0.131	-0.374	0.389	-0.005	0.090
5	3.136	3.145	0.037	0.136	-0.313	0.320	-0.007	0.097	6.829	6.838	0.027	0.135	-0.314	0.338	-0.006	0.099
10	2.486	2.494	0.017	0.146	-0.215	0.229	-0.008	0.110	4.720	4.727	0.013	0.149	-0.213	0.252	-0.007	0.113
15	2.014	2.021	0.010	0.155	-0.156	0.178	-0.009	0.119	3.488	3.495	0.009	0.160	-0.158	0.204	-0.008	0.122
30	1.222	1.229	0.002	0.175	-0.082	0.120	-0.011	0.135	1.859	1.866	0.005	0.181	-0.081	0.140	-0.011	0.138
60	0.646	0.657	-0.005	0.200	-0.042	0.100	-0.014	0.154	0.938	0.949	0.002	0.208	-0.039	0.114	-0.015	0.160
120	0.324	0.350	-0.015	0.234	-0.020	0.104	-0.018	0.182	0.467	0.491	-0.007	0.243	-0.020	0.115	-0.018	0.188
240	0.161	0.231	-0.024	0.283	-0.010	0.128	-0.019	0.224	0.233	0.293	-0.016	0.291	-0.009	0.137	-0.021	0.230

(a) Constant Volatility

	Pure Rounding Error				*Noise and Rounding Error*			
	Variance		Covariance		Variance		Covariance	
	MC		MC		MC		MC	
k	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse
1	0.005	0.125	-0.020	0.091	0.219	0.272	-0.055	0.143
2	0.003	0.165	-0.024	0.084	0.117	0.185	-0.093	0.131
3	0.003	0.214	-0.019	0.089	0.055	0.175	-0.053	0.097
4	0.002	0.260	-0.010	0.098	0.034	0.193	-0.050	0.094
5	-0.001	0.327	-0.007	0.108	0.021	0.212	-0.034	0.090

	RV		MRK		RC		MRK		RV		MRK		RC		MRK	
H	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse	Bias	rmse
1	3.070	3.325	0.132	0.196	-0.325	0.443	-0.007	0.072	9.619	9.639	0.103	0.202	-0.333	0.461	-0.007	0.073
3	2.812	3.032	0.058	0.168	-0.270	0.370	-0.006	0.086	8.010	8.028	0.042	0.170	-0.275	0.392	-0.005	0.087
5	2.555	2.741	0.037	0.180	-0.227	0.313	-0.007	0.094	6.757	6.773	0.026	0.177	-0.230	0.340	-0.006	0.095
10	2.017	2.141	0.018	0.198	-0.155	0.221	-0.008	0.105	4.656	4.668	0.013	0.197	-0.160	0.256	-0.007	0.108
15	1.636	1.721	0.011	0.222	-0.114	0.173	-0.009	0.113	3.432	3.443	0.008	0.214	-0.116	0.205	-0.007	0.117
30	1.016	1.056	0.000	0.249	-0.059	0.112	-0.011	0.129	1.825	1.836	0.002	0.253	-0.062	0.142	-0.009	0.134
60	0.571	0.622	-0.010	0.295	-0.031	0.098	-0.014	0.151	0.924	0.942	-0.004	0.311	-0.032	0.116	-0.011	0.157
120	0.304	0.371	-0.022	0.374	-0.016	0.103	-0.016	0.181	0.462	0.508	-0.013	0.389	-0.017	0.113	-0.014	0.185
240	0.157	0.314	-0.031	0.498	-0.009	0.127	-0.016	0.223	0.232	0.357	-0.023	0.488	-0.009	0.132	-0.016	0.226

(b) Stochastic Volatility

Table 4.2: Simulation results. Top panel contains the simulation results where the underlying volatility is constant. The bottom panel contains the corresponding results for the stochastic volatility. For both cases we consider two situations - the observed process is only subject to rounding errors (left panels) and to both noise and rounding errors (right panels).

Data Summary Statistics

	SPY	CL	NG	GC	SV	HG	LC	LH	KC	SB	CT	CN	SY	WC
Transactions per day (full trading day)	19178	21993	9761	21671	9607	9738	2835	2775	2922	4013	3325	7058	8882	4684
Transactions per day: n (10:00am - 2:00pm)	10117	8634	4495	6320	3024	2706	2315	2213	1613	2125	1940	4362	5011	2895
Primitive states: (10:00am - 2:00pm)	11	19	11	17	11	8	8	8	11	7	23	11	13	9
Tick size	0.01	0.01	0.001	0.1	0.5	0.05	0.025	0.025	0.05	0.01	0.01	0.25	0.25	0.25
Volatility share: κ (10am-2pm)/24h	0.31	0.36	0.40	0.24	0.20	0.20	0.47	0.25	0.44	0.37	0.46	0.24	0.25	0.52
Annual volatility (2013)	11%	19%	31%	20%	32%	19%	12%	24%	27%	18%	21%	33%	25%	22%

Table 4.3: Summary statistics for the 251 trading days in 2013. The average number of price observations per day, and within the 10:00 AM to 2:00 PM window (n) are reported in the first two rows, followed by the average number of primitive states (S) and the tick size for each of the assets. The second last row reports κ – the fraction of daily volatility that occurs during the 10:00 AM to 2:00 PM window on average. The average volatility for 2013 is reported in the last row.

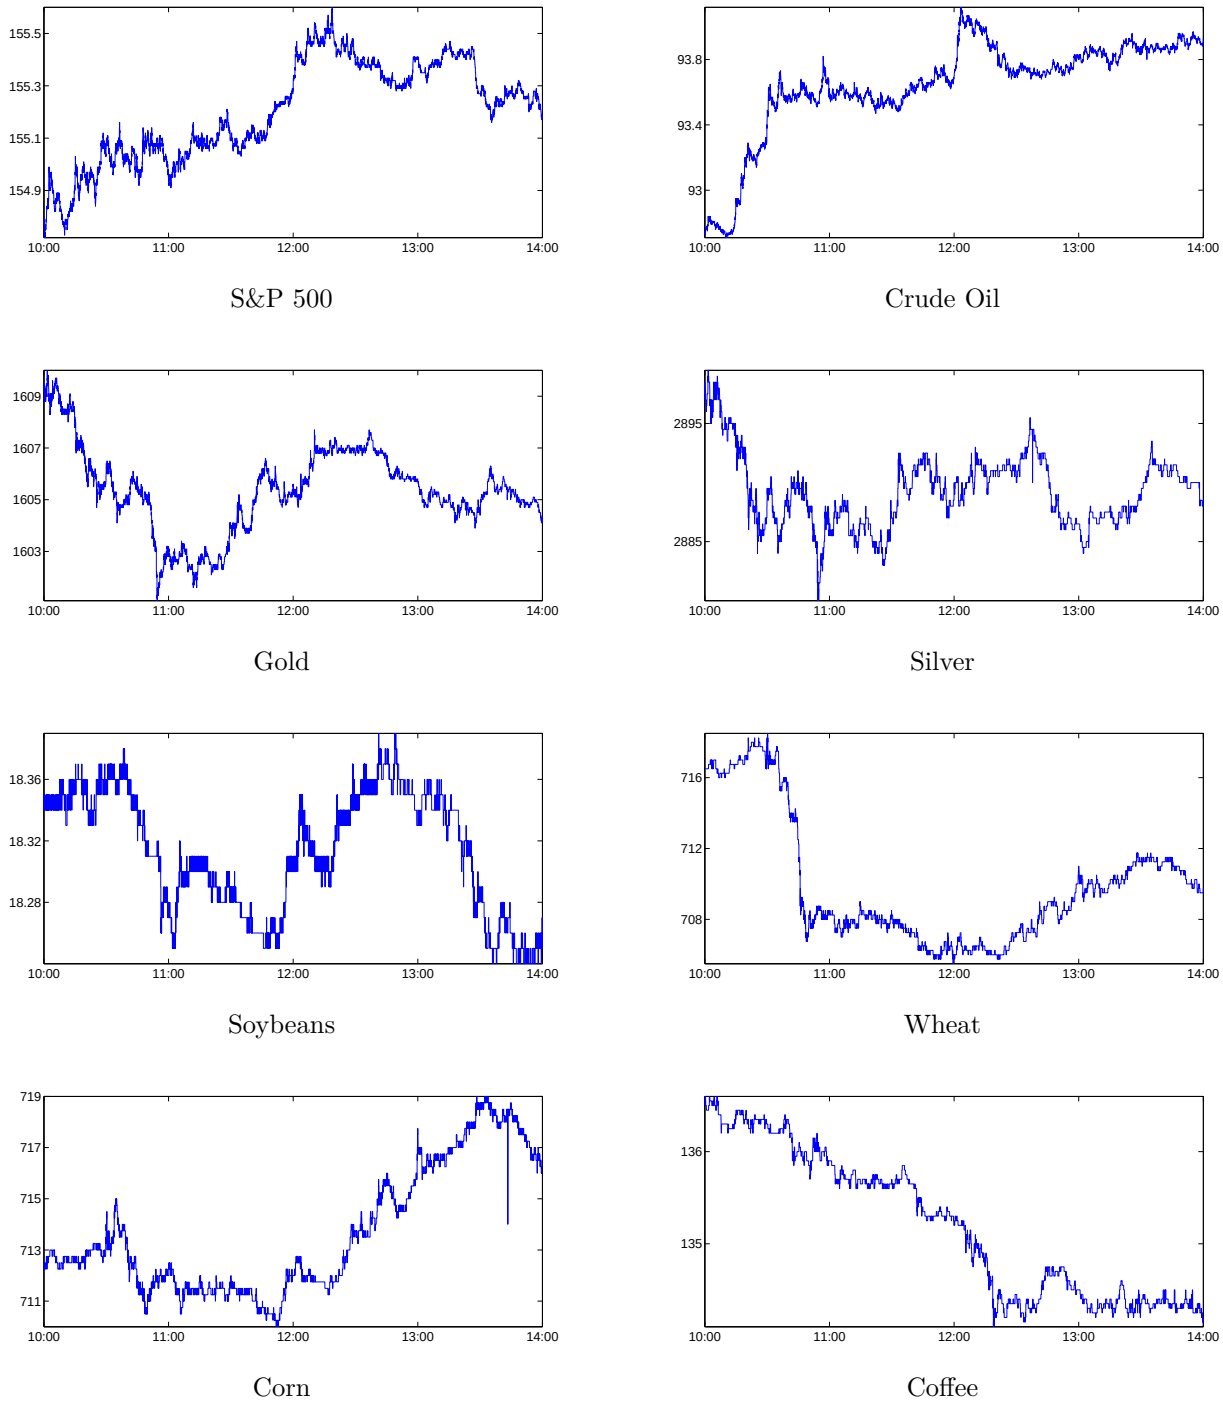

Figure 4.2: High frequency prices for eight selected commodities on March 18th, 2013 during the period from 10:00AM to 2:00PM.

	SPY	CL	GC	SV	KC
S&P 500	**81.16**	79.39	-8.26	-20.03	-15.24
Light Crude	*0.48*	**344.08**	-36.52	-1.39	-2.29
Gold	*-0.08*	*-0.16*	**146.11**	164.45	-47.74
Silver	*-0.13*	*-0.00*	*0.77*	**316.22**	-36.77
Coffee "C"	*-0.08*	*-0.01*	*-0.20*	*-0.10*	**407.69**

(a) RV_{5min}: *Realized variance with 5-minute sampling*

	SPY	CL	GC	SV	KC
S&P 500	**68.92**	90.57	1.98	-20.04	-40.79
Light Crude	*0.59*	**342.05**	-69.66	-57.07	-11.96
Gold	*0.03*	*-0.43*	**78.43**	50.73	-35.05
Silver	*-0.20*	*-0.26*	*0.48*	**139.93**	-17.48
Coffee "C"	*-0.23*	*-0.03*	*-0.19*	*-0.07*	**446.70**

(b) RV_{10min}: *Realized variance with 10-minute sampling*

	SPY	CL	GC	SV	KC
S&P 500	**80.65**	88.17	-24.05	-77.56	-1.89
Light Crude	*0.49*	**407.24**	-75.01	-124.63	49.36
Gold	*-0.26*	*-0.35*	**109.99**	155.36	-32.78
Silver	*-0.45*	*-0.32*	*0.77*	**370.69**	-12.07
Coffee "C"	*-0.01*	*0.15*	*-0.19*	*-0.04*	**257.15**

(c) MC: *Markov Chain Estimator (Full)*

	SPY	CL	GC	SV	KC
S&P 500	**116.84**	75.77	-10.69	4.58	21.01
Light Crude	*0.35*	**391.86**	14.34	61.58	41.60
Gold	*-0.09*	*0.07*	**116.87**	151.25	-16.83
Silver	*0.02*	*0.16*	*0.72*	**380.17**	-0.71
Coffee "C"	*0.09*	*0.10*	*-0.08*	*-0.00*	**421.94**

(d) MC^1: *Markov Chain Estimator 1-Composite*

	SPY	CL	GC	SV	KC
S&P 500	**116.84**	80.35	-12.92	-7.44	13.76
Light Crude	*0.38*	**391.86**	-1.45	39.30	29.37
Gold	*-0.11*	*-0.01*	**116.87**	149.72	-27.00
Silver	*-0.04*	*0.10*	*0.71*	**380.17**	-59.90
Coffee "C"	*0.06*	*0.07*	*-0.12*	*-0.15*	**421.94**

(e) MC^2: *Markov Chain 2-Composite*

Table 4.4: Two realized variances (5-min and 10-min) and three versions of the Markov chain estimator (full, 1-composite and 2-composite) are presented. The estimates are for March 18, 2013 (10:00AM-2:00PM) for five selected commodities. Variances and covariances are annualized and further scaled by 10^4. Correlations are in the lower triangle in italic font.

RV_{5min}	**SPY**	**CL**	**NG**	**GC**	**SV**	**HG**	**LC**	**LH**	**KC**	**SB**	**CT**	**CN**	**SY**	**WC**
S&P 500	**81.16**	79.39	18.76	-8.26	-20.03	44.26	-23.18	-1.03	-15.24	23.70	46.58	12.24	16.52	-21.65
Light Crude	*0.48*	**344.08**	80.92	-36.52	-1.39	73.65	-7.29	-69.62	-2.29	-20.68	11.62	5.80	32.45	23.80
Natural Gas	*0.07*	*0.15*	**805.41**	-4.63	60.09	159.46	-94.66	-21.79	53.54	-50.70	-54.01	-85.28	22.49	37.54
Gold	*-0.08*	*-0.16*	*-0.01*	**146.11**	164.45	6.66	-8.35	-4.19	-47.74	-39.13	11.51	-43.91	-21.55	-32.41
Silver	*-0.13*	*-0.00*	*0.12*	*0.77*	**316.22**	59.49	-21.80	9.80	-36.77	-51.15	16.38	-61.21	-46.76	-22.01
Copper	*0.30*	*0.24*	*0.34*	*0.03*	*0.20*	**269.51**	-70.43	13.29	89.35	13.11	24.30	-53.07	68.05	8.94
Live Cattle	*-0.19*	*-0.03*	*-0.25*	*-0.05*	*-0.09*	*-0.32*	**180.30**	18.30	-23.07	-16.29	-26.96	70.81	-6.47	80.24
Lean Hogs	*-0.00*	*-0.15*	*-0.03*	*-0.01*	*0.02*	*0.03*	*0.05*	**618.27**	-54.91	-7.75	9.10	-44.27	-88.09	-1.78
Coffee "C"	*-0.08*	*-0.01*	*0.09*	*-0.20*	*-0.10*	*0.27*	*-0.09*	*-0.11*	**407.69**	34.62	12.53	123.43	62.26	180.13
Sugar #1	*0.16*	*-0.07*	*-0.11*	*-0.20*	*-0.17*	*0.05*	*-0.07*	*-0.02*	*0.10*	**274.87**	65.41	-29.63	28.12	46.86
Cotton #2	*0.24*	*0.03*	*-0.09*	*0.04*	*0.04*	*0.07*	*-0.09*	*0.02*	*0.03*	*0.18*	**478.52**	87.19	102.05	53.99
Corn	*0.05*	*0.01*	*-0.12*	*-0.14*	*-0.14*	*-0.13*	*0.21*	*-0.07*	*0.24*	*-0.07*	*0.16*	**640.54**	288.11	412.80
Soybeans	*0.09*	*0.09*	*0.04*	*-0.09*	*-0.13*	*0.20*	*-0.02*	*-0.17*	*0.15*	*0.08*	*0.23*	*0.56*	**415.85**	232.42
Wheat	*-0.10*	*0.05*	*0.06*	*-0.11*	*-0.05*	*0.02*	*0.25*	*-0.00*	*0.38*	*0.12*	*0.10*	*0.69*	*0.48*	**553.00**

RV_{10min}	**SPY**	**CL**	**NG**	**GC**	**SV**	**HG**	**LC**	**LH**	**KC**	**SB**	**CT**	**CN**	**SY**	**WC**
S&P 500	**68.92**	90.57	-7.81	1.98	-20.04	26.06	3.25	24.47	-40.79	-5.38	35.40	-3.91	-24.18	-32.89
Light Crude	*0.59*	**342.05**	98.06	-69.66	-57.07	90.71	9.94	21.31	-11.96	-29.12	88.15	-8.19	20.25	18.49
Natural Gas	*-0.04*	*0.22*	**578.93**	45.71	101.82	160.88	-78.18	-104.34	104.28	-42.88	47.80	-36.26	66.09	51.45
Gold	*0.03*	*-0.43*	*0.21*	**78.43**	50.73	9.61	-17.97	-8.42	-35.05	-2.83	-28.70	-38.85	-18.87	-32.80
Silver	*-0.20*	*-0.26*	*0.36*	*0.48*	**139.93**	54.77	-23.71	-15.59	-17.48	-13.00	-33.72	-39.72	-6.86	-23.42
Copper	*0.18*	*0.28*	*0.38*	*0.06*	*0.26*	**317.27**	-102.54	114.21	123.27	35.28	-37.42	-78.76	48.24	49.66
Live Cattle	*0.03*	*0.04*	*-0.25*	*-0.16*	*-0.16*	*-0.45*	**164.68**	30.43	-25.38	2.54	88.65	151.23	65.74	148.58
Lean Hogs	*0.14*	*0.06*	*-0.21*	*-0.05*	*-0.06*	*0.31*	*0.12*	**418.34**	-94.24	41.35	31.23	-113.20	-66.97	4.08
Coffee "C"	*-0.23*	*-0.03*	*0.21*	*-0.19*	*-0.07*	*0.33*	*-0.09*	*-0.22*	**446.70**	60.91	-62.30	108.18	100.85	279.64
Sugar #1	*-0.04*	*-0.10*	*-0.11*	*-0.02*	*-0.07*	*0.13*	*0.01*	*0.13*	*0.19*	**242.43**	145.50	-26.27	-11.14	108.60
Cotton #2	*0.19*	*0.21*	*0.09*	*-0.14*	*-0.13*	*-0.09*	*0.31*	*0.07*	*-0.13*	*0.42*	**503.18**	54.13	73.48	147.24
Corn	*-0.02*	*-0.02*	*-0.06*	*-0.17*	*-0.13*	*-0.17*	*0.45*	*-0.21*	*0.19*	*-0.06*	*0.09*	**699.55**	403.33	542.96
Soybeans	*-0.14*	*0.05*	*0.13*	*-0.10*	*-0.03*	*0.13*	*0.24*	*-0.15*	*0.22*	*-0.03*	*0.15*	*0.71*	**463.78**	452.48
Wheat	*-0.14*	*0.03*	*0.07*	*-0.13*	*-0.07*	*0.10*	*0.39*	*0.01*	*0.45*	*0.24*	*0.22*	*0.70*	*0.72*	**860.64**

Table 4.5: The table presents two covariance matrices for the 14 commodities estimated using 5-minute (top panel) and 10-minute (bottom panel) RV on March 18th, 2013 (10am - 2pm). Variances and covariances are annualized and further scaled by 10^4. Correlations are reported below the main diagonal with an italic font.

MC_{1-comp}	**SPY**	**CL**	**NG**	**GC**	**SV**	**HG**	**LC**	**LH**	**KC**	**SB**	**CT**	**CN**	**SY**	**WC**
S&P 500	**116.84**	75.77	-34.48	-10.69	4.58	73.32	2.39	-5.64	21.01	-6.10	2.00	-4.96	2.38	1.81
Light Crude	*0.35*	**391.86**	-23.61	14.34	61.58	101.52	11.42	-31.64	41.60	5.07	11.34	14.14	5.57	3.12
Natural Gas	*-0.11*	*-0.04*	**921.74**	-1.88	-14.76	30.77	-26.48	8.46	29.60	-10.02	-61.60	52.74	-6.90	-1.04
Gold	*-0.09*	*0.07*	*-0.01*	**116.87**	151.25	9.18	-4.74	7.84	-16.83	10.41	-36.42	28.66	5.55	19.28
Silver	*0.02*	*0.16*	*-0.02*	*0.72*	**380.17**	31.04	30.64	30.20	-0.71	-3.73	85.57	-1.61	5.51	-17.91
Copper	*0.35*	*0.27*	*0.05*	*0.04*	*0.08*	**365.15**	-8.35	-21.68	-4.71	-36.98	13.93	-5.18	-23.51	-7.07
Live Cattle	*0.02*	*0.05*	*-0.07*	*-0.04*	*0.13*	*-0.04*	**142.36**	40.37	4.44	8.11	-1.72	13.95	-18.30	27.09
Lean Hogs	*-0.02*	*-0.06*	*0.01*	*0.03*	*0.06*	*-0.05*	*0.14*	**606.71**	4.55	-2.74	-22.64	-0.43	6.19	-19.86
Coffee "C"	*0.09*	*0.10*	*0.05*	*-0.08*	*-0.00*	*-0.01*	*0.02*	*0.01*	**421.94**	10.27	91.99	-1.41	39.64	-39.64
Sugar #1	*-0.03*	*0.01*	*-0.02*	*0.05*	*-0.01*	*-0.10*	*0.04*	*-0.01*	*0.03*	**345.21**	-37.34	3.76	-6.25	32.86
Cotton #2	*0.01*	*0.02*	*-0.09*	*-0.14*	*0.19*	*0.03*	*-0.01*	*-0.04*	*0.19*	*-0.09*	**544.65**	134.15	68.07	73.68
Corn	*-0.02*	*0.03*	*0.08*	*0.12*	*-0.00*	*-0.01*	*0.05*	*-0.00*	*-0.00*	*0.01*	*0.26*	**475.69**	204.56	276.23
Soybeans	*0.01*	*0.01*	*-0.01*	*0.03*	*0.01*	*-0.06*	*-0.08*	*0.01*	*0.10*	*-0.02*	*0.15*	*0.49*	**365.44**	132.19
Wheat	*0.01*	*0.01*	*-0.00*	*0.09*	*-0.05*	*-0.02*	*0.11*	*-0.04*	*-0.10*	*0.09*	*0.16*	*0.63*	*0.34*	**407.53**

MC_{2-comp}	**SPY**	**CL**	**NG**	**GC**	**SV**	**HG**	**LC**	**LH**	**KC**	**SB**	**CT**	**CN**	**SY**	**WC**
S&P 500	**116.84**	80.35	-12.55	-12.92	-7.44	76.88	8.40	-5.69	13.76	1.67	42.68	-9.93	10.89	0.47
Light Crude	*0.38*	**391.86**	7.67	-1.45	39.30	125.36	6.38	-11.56	29.37	-7.56	89.63	-25.61	22.06	-3.37
Natural Gas	*-0.04*	*0.01*	**921.74**	1.64	-44.28	84.25	-32.41	-10.85	74.22	-53.72	-56.59	34.81	-12.94	-22.93
Gold	*-0.11*	*-0.01*	*0.00*	**116.87**	149.72	15.23	-2.08	25.73	-27.00	-4.72	-37.49	0.15	5.05	-3.15
Silver	*-0.04*	*0.10*	*-0.07*	*0.71*	**380.17**	23.70	4.60	16.03	-59.90	23.58	-60.49	-36.39	-1.79	-38.08
Copper	*0.37*	*0.33*	*0.15*	*0.07*	*0.06*	**365.15**	-29.91	17.98	24.53	0.16	80.92	-7.86	-34.35	-16.89
Live Cattle	*0.07*	*0.03*	*-0.09*	*-0.02*	*0.02*	*-0.13*	**142.36**	18.71	-20.08	4.14	-15.27	9.25	-19.75	31.61
Lean Hogs	*-0.02*	*-0.02*	*-0.01*	*0.10*	*0.03*	*0.04*	*0.06*	**606.71**	69.63	7.01	426.46	-27.19	19.11	35.20
Coffee "C"	*0.06*	*0.07*	*0.12*	*-0.12*	*-0.15*	*0.06*	*-0.08*	*0.14*	**421.94**	23.83	5.64	-15.59	54.31	-39.85
Sugar #1	*0.01*	*-0.02*	*-0.10*	*-0.02*	*0.07*	*0.00*	*0.02*	*0.02*	*0.06*	**345.21**	95.41	11.69	-28.41	29.00
Cotton #2	*0.17*	*0.19*	*-0.08*	*-0.15*	*-0.13*	*0.18*	*-0.05*	*0.74*	*0.01*	*0.22*	**544.65**	83.64	19.19	68.24
Corn	*-0.04*	*-0.06*	*0.05*	*0.00*	*-0.09*	*-0.02*	*0.04*	*-0.05*	*-0.03*	*0.03*	*0.16*	**475.69**	236.91	320.47
Soybeans	*0.05*	*0.06*	*-0.02*	*0.02*	*-0.00*	*-0.09*	*-0.09*	*0.04*	*0.14*	*-0.08*	*0.04*	*0.57*	**365.44**	162.80
Wheat	*0.00*	*-0.01*	*-0.04*	*-0.01*	*-0.10*	*-0.04*	*0.13*	*0.07*	*-0.10*	*0.08*	*0.14*	*0.73*	*0.42*	**407.53**

Table 4.6: The table presents two covariance matrices for the 14 commodities estimated using 1-composite (top panel) and 2-composite (bottom panel) Markov chain estimators on March 18th, 2013 (10am - 2pm). Variances and covariances are annualized and further scaled by 10^4. Correlations are reported below the main diagonal with an italic font.

RV_{5min}	**SPY**	**CL**	**NG**	**GC**	**SV**	**HG**	**LC**	**LH**	**KC**	**SB**	**CT**	**CN**	**SY**	**WC**
S&P 500	**98.20**	72.55	7.53	-29.49	-40.80	59.85	11.82	8.20	14.35	-0.77	11.53	12.48	6.01	14.05
Light Crude	*0.42*	**304.64**	41.55	4.72	27.94	106.73	17.78	34.32	46.45	5.08	15.23	45.38	25.09	37.27
Natural Gas	*0.03*	*0.08*	**828.53**	-0.90	5.96	43.62	5.33	5.39	3.84	-21.71	4.43	18.29	28.41	-0.53
Gold	*-0.27*	*0.02*	*-0.00*	**123.93**	201.43	21.89	-5.48	-18.43	-10.53	-3.25	-3.96	8.03	14.74	6.89
Silver	*-0.18*	*0.07*	*0.01*	*0.80*	**512.90**	81.49	2.16	-45.85	15.06	13.05	-0.98	4.98	30.76	8.66
Copper	*0.38*	*0.38*	*0.09*	*0.12*	*0.22*	**257.37**	7.16	-10.27	49.36	11.44	-2.78	40.49	38.99	52.05
Live Cattle	*0.08*	*0.07*	*0.01*	*-0.03*	*0.01*	*0.03*	**199.14**	153.29	4.09	-6.71	8.24	11.19	10.78	15.68
Lean Hogs	*0.03*	*0.06*	*0.01*	*-0.05*	*-0.07*	*-0.02*	*0.36*	**935.02**	9.53	-60.11	20.26	7.31	35.86	-1.77
Coffee "C"	*0.06*	*0.10*	*0.01*	*-0.04*	*0.03*	*0.12*	*0.01*	*0.01*	**657.53**	74.46	16.09	63.27	22.37	61.59
Sugar #1	*-0.00*	*0.01*	*-0.03*	*-0.01*	*0.02*	*0.03*	*-0.02*	*-0.08*	*0.12*	**551.36**	2.53	79.07	42.36	72.83
Cotton #2	*0.05*	*0.04*	*0.01*	*-0.02*	*-0.00*	*-0.01*	*0.03*	*0.03*	*0.03*	*0.00*	**488.85**	-69.86	-17.14	-85.78
Corn	*0.04*	*0.08*	*0.02*	*0.02*	*0.01*	*0.08*	*0.03*	*0.01*	*0.08*	*0.11*	*-0.10*	**996.64**	335.43	711.51
Soybeans	*0.03*	*0.06*	*0.04*	*0.06*	*0.06*	*0.10*	*0.03*	*0.05*	*0.04*	*0.08*	*-0.03*	*0.45*	**552.12**	270.99
Wheat	*0.05*	*0.07*	*-0.00*	*0.02*	*0.01*	*0.11*	*0.04*	*-0.00*	*0.08*	*0.11*	*-0.14*	*0.79*	*0.40*	**818.54**

RV_{10min}	**SPY**	**CL**	**NG**	**GC**	**SV**	**HG**	**LC**	**LH**	**KC**	**SB**	**CT**	**CN**	**SY**	**WC**
S&P 500	**103.38**	83.04	1.75	-35.00	-43.92	67.72	18.68	17.29	13.72	1.82	14.03	11.27	13.61	20.55
Light Crude	*0.45*	**329.37**	59.02	3.89	37.04	127.49	17.70	46.91	68.14	12.95	7.74	50.17	29.61	36.87
Natural Gas	*0.01*	*0.12*	**771.89**	14.50	45.49	41.19	-2.69	8.37	0.61	-54.02	22.02	-18.33	4.83	-0.14
Gold	*-0.32*	*0.02*	*0.05*	**117.30**	193.10	14.79	-16.56	-28.31	-4.43	7.41	-6.04	2.63	11.57	2.25
Silver	*-0.19*	*0.09*	*0.07*	*0.79*	**510.20**	84.45	-18.71	-84.96	34.49	30.12	-0.59	1.41	21.38	14.52
Copper	*0.42*	*0.44*	*0.09*	*0.09*	*0.23*	**255.99**	16.10	-0.08	58.08	18.61	-18.35	39.15	40.07	63.66
Live Cattle	*0.13*	*0.07*	*-0.01*	*-0.11*	*-0.06*	*0.07*	**206.36**	182.41	13.66	-7.21	26.27	21.69	24.43	18.48
Lean Hogs	*0.06*	*0.09*	*0.01*	*-0.09*	*-0.13*	*-0.00*	*0.43*	**870.97**	-8.22	-37.37	1.44	38.99	50.73	23.20
Coffee "C"	*0.05*	*0.14*	*0.00*	*-0.02*	*0.06*	*0.14*	*0.04*	*-0.01*	**689.21**	75.36	9.40	96.86	77.15	83.21
Sugar #1	*0.01*	*0.03*	*-0.09*	*0.03*	*0.06*	*0.05*	*-0.02*	*-0.06*	*0.13*	**502.96**	17.81	92.54	76.18	91.50
Cotton #2	*0.07*	*0.02*	*0.04*	*-0.03*	*-0.00*	*-0.05*	*0.09*	*0.00*	*0.02*	*0.04*	**442.55**	-12.93	19.53	0.36
Corn	*0.04*	*0.09*	*-0.02*	*0.01*	*0.00*	*0.08*	*0.05*	*0.04*	*0.12*	*0.13*	*-0.02*	**999.24**	356.24	647.75
Soybeans	*0.05*	*0.06*	*0.01*	*0.04*	*0.04*	*0.10*	*0.07*	*0.07*	*0.11*	*0.13*	*0.04*	*0.44*	**653.10**	265.91
Wheat	*0.08*	*0.08*	*-0.00*	*0.01*	*0.02*	*0.15*	*0.05*	*0.03*	*0.12*	*0.16*	*0.00*	*0.78*	*0.40*	**683.20**

Table 4.7: The table presents two covariance matrices for the 14 commodities estimated using 5-minute (top panel) and 10-minute (bottom panel) RV for March, 2013 (pooled estimates). Variances and covariances are annualized and further scaled by 10^4. Correlations are reported below the main diagonal with an italic font.

MC_{1-comp}	**SPY**	**CL**	**NG**	**GC**	**SV**	**HG**	**LC**	**LH**	**KC**	**SB**	**CT**	**CN**	**SY**	**WC**
S&P 500	**106.38**	69.59	-1.52	-12.53	-9.17	49.64	5.09	-3.36	10.86	-1.08	5.40	11.11	12.02	8.20
Light Crude	*0.37*	**326.26**	14.89	9.34	28.47	82.09	6.53	2.76	16.30	-5.62	5.13	2.92	16.26	14.55
Natural Gas	*-0.00*	*0.03*	**932.63**	7.18	18.48	10.08	12.13	-5.79	22.71	-0.61	-9.77	-22.36	14.87	1.44
Gold	*-0.10*	*0.04*	*0.02*	**149.93**	219.11	26.36	-0.83	1.71	1.00	-2.37	-1.74	-0.26	2.41	-0.02
Silver	*-0.04*	*0.07*	*0.03*	*0.76*	**557.24**	73.28	0.12	-5.14	6.68	-13.18	1.91	4.63	5.79	-1.43
Copper	*0.30*	*0.28*	*0.02*	*0.13*	*0.19*	**259.93**	3.97	-6.91	18.82	-1.84	9.30	16.38	9.38	10.22
Live Cattle	*0.03*	*0.02*	*0.03*	*-0.00*	*0.00*	*0.02*	**213.33**	128.46	11.34	-1.90	16.85	-2.27	1.36	5.90
Lean Hogs	*-0.01*	*0.01*	*-0.01*	*0.00*	*-0.01*	*-0.01*	*0.31*	**830.83**	9.27	-5.11	9.64	3.63	11.97	6.54
Coffee "C"	*0.04*	*0.04*	*0.03*	*0.00*	*0.01*	*0.05*	*0.03*	*0.01*	**636.57**	36.08	7.91	2.97	7.70	8.80
Sugar #1	*-0.00*	*-0.01*	*-0.00*	*-0.01*	*-0.02*	*-0.00*	*-0.01*	*-0.01*	*0.06*	**528.18**	-1.88	33.55	15.23	28.24
Cotton #2	*0.02*	*0.01*	*-0.01*	*-0.01*	*0.00*	*0.03*	*0.05*	*0.02*	*0.01*	*-0.00*	**490.22**	12.67	11.50	13.48
Corn	*0.04*	*0.01*	*-0.03*	*-0.00*	*0.01*	*0.04*	*-0.01*	*0.00*	*0.00*	*0.05*	*0.02*	**835.96**	299.02	523.38
Soybeans	*0.05*	*0.04*	*0.02*	*0.01*	*0.01*	*0.02*	*0.00*	*0.02*	*0.01*	*0.03*	*0.02*	*0.43*	**578.95**	220.34
Wheat	*0.03*	*0.03*	*0.00*	*-0.00*	*-0.00*	*0.02*	*0.02*	*0.01*	*0.01*	*0.05*	*0.02*	*0.70*	*0.35*	**666.82**

MC_{2-comp}	**SPY**	**CL**	**NG**	**GC**	**SV**	**HG**	**LC**	**LH**	**KC**	**SB**	**CT**	**CN**	**SY**	**WC**
S&P 500	**106.74**	64.40	1.63	-6.86	-3.79	43.57	4.26	-0.65	12.19	0.56	5.81	7.09	9.53	8.31
Light Crude	*0.34*	**331.52**	13.26	13.49	37.73	72.97	4.36	1.84	13.61	1.68	2.90	7.64	12.45	7.86
Natural Gas	*0.01*	*0.02*	**961.44**	7.04	14.65	5.06	7.52	-6.65	12.05	2.72	-3.34	-13.22	-3.19	-7.69
Gold	*-0.05*	*0.06*	*0.02*	**157.12**	224.80	27.55	-1.73	0.16	0.72	-0.03	-0.43	1.52	1.57	1.51
Silver	*-0.02*	*0.09*	*0.02*	*0.75*	**572.28**	75.46	-0.12	-12.66	9.99	-8.22	0.32	4.02	5.11	-0.39
Copper	*0.26*	*0.24*	*0.01*	*0.13*	*0.19*	**267.97**	3.44	-3.69	17.74	4.27	2.94	10.85	12.63	5.37
Live Cattle	*0.03*	*0.02*	*0.02*	*-0.01*	*-0.00*	*0.01*	**203.06**	101.95	9.47	-3.76	5.64	1.08	1.29	1.33
Lean Hogs	*-0.00*	*0.00*	*-0.01*	*0.00*	*-0.02*	*-0.01*	*0.25*	**807.48**	9.01	3.25	11.19	4.66	1.52	2.10
Coffee "C"	*0.05*	*0.03*	*0.02*	*0.00*	*0.02*	*0.04*	*0.03*	*0.01*	**640.97**	23.88	5.15	9.98	1.77	7.60
Sugar #1	*0.00*	*0.00*	*0.00*	*-0.00*	*-0.02*	*0.01*	*-0.01*	*0.01*	*0.04*	**478.69**	5.00	16.03	5.51	24.56
Cotton #2	*0.03*	*0.01*	*-0.00*	*-0.00*	*0.00*	*0.01*	*0.02*	*0.02*	*0.01*	*0.01*	**484.68**	15.93	17.28	9.47
Corn	*0.02*	*0.02*	*-0.02*	*0.00*	*0.01*	*0.02*	*0.00*	*0.01*	*0.01*	*0.03*	*0.03*	**760.38**	232.67	605.64
Soybeans	*0.04*	*0.03*	*-0.00*	*0.01*	*0.01*	*0.03*	*0.00*	*0.00*	*0.00*	*0.01*	*0.03*	*0.35*	**572.85**	208.77
Wheat	*0.03*	*0.02*	*-0.01*	*0.00*	*-0.00*	*0.01*	*0.00*	*0.00*	*0.01*	*0.05*	*0.02*	*0.90*	*0.36*	**594.62**

Table 4.8: The table presents two covariance matrices for the 14 commodities estimated using 1-composite (top panel) and 2-composite (bottom panel) Markov chain estimators on March, 2013 (pooled estimates). Variances and covariances are annualized and further scaled by 10^4. Correlations are reported below the main diagonal with an italic font.

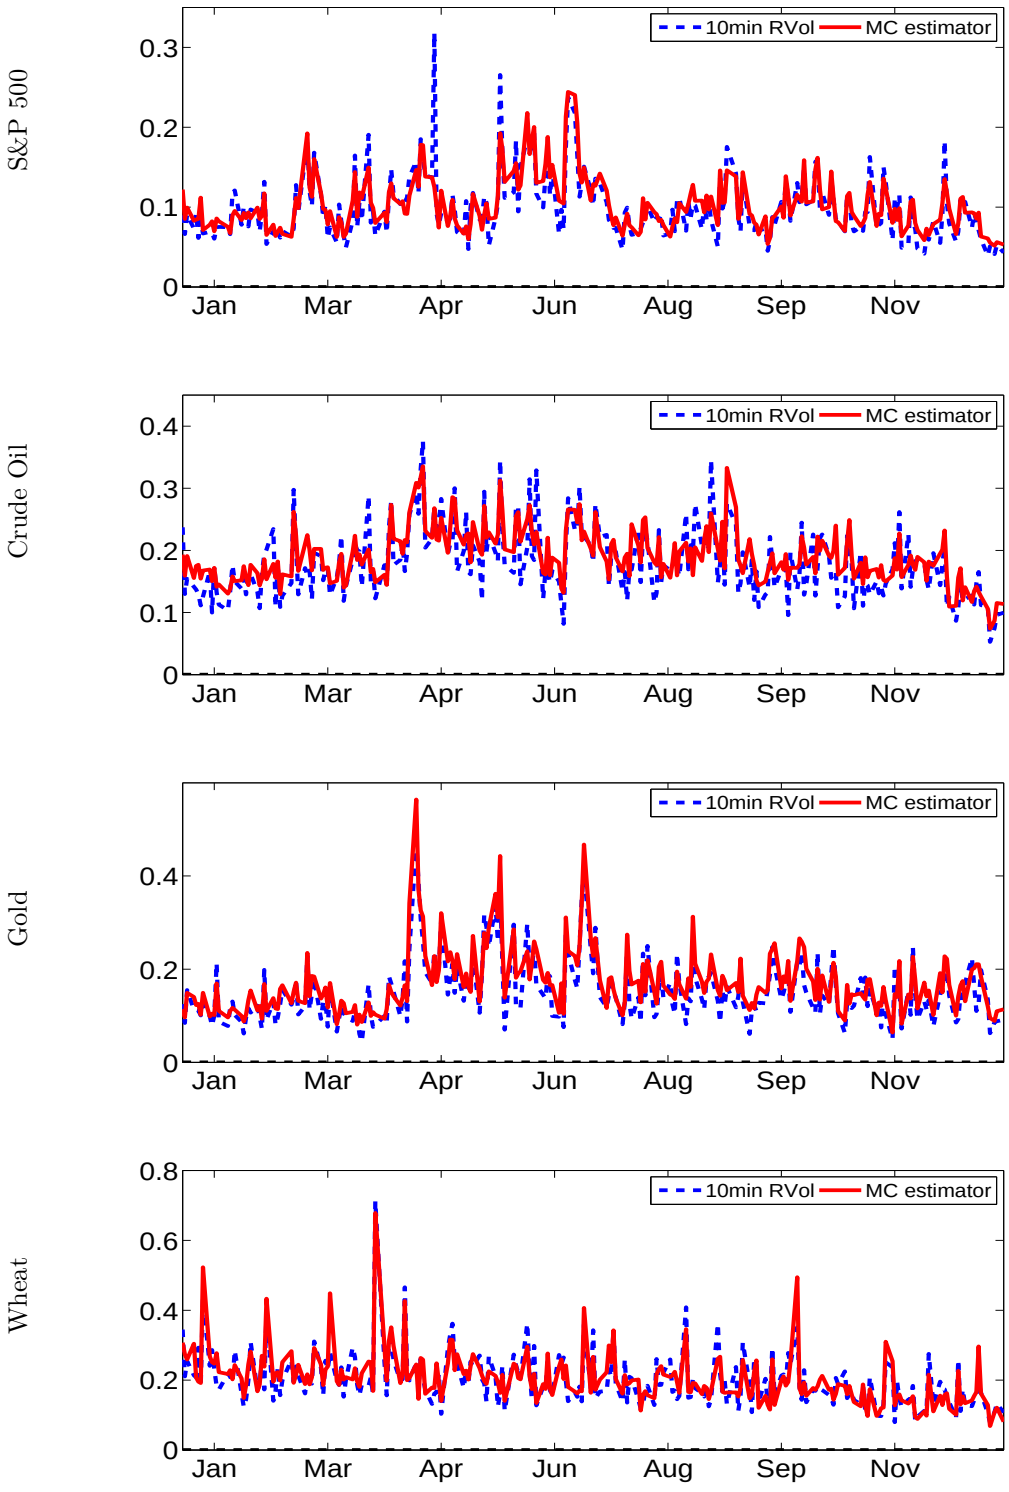

Figure 4.3: Realized volatility based on 10 min returns against volatility computed with MC estimator (order $k = 5$) in 2013. Estimated values are annualized.

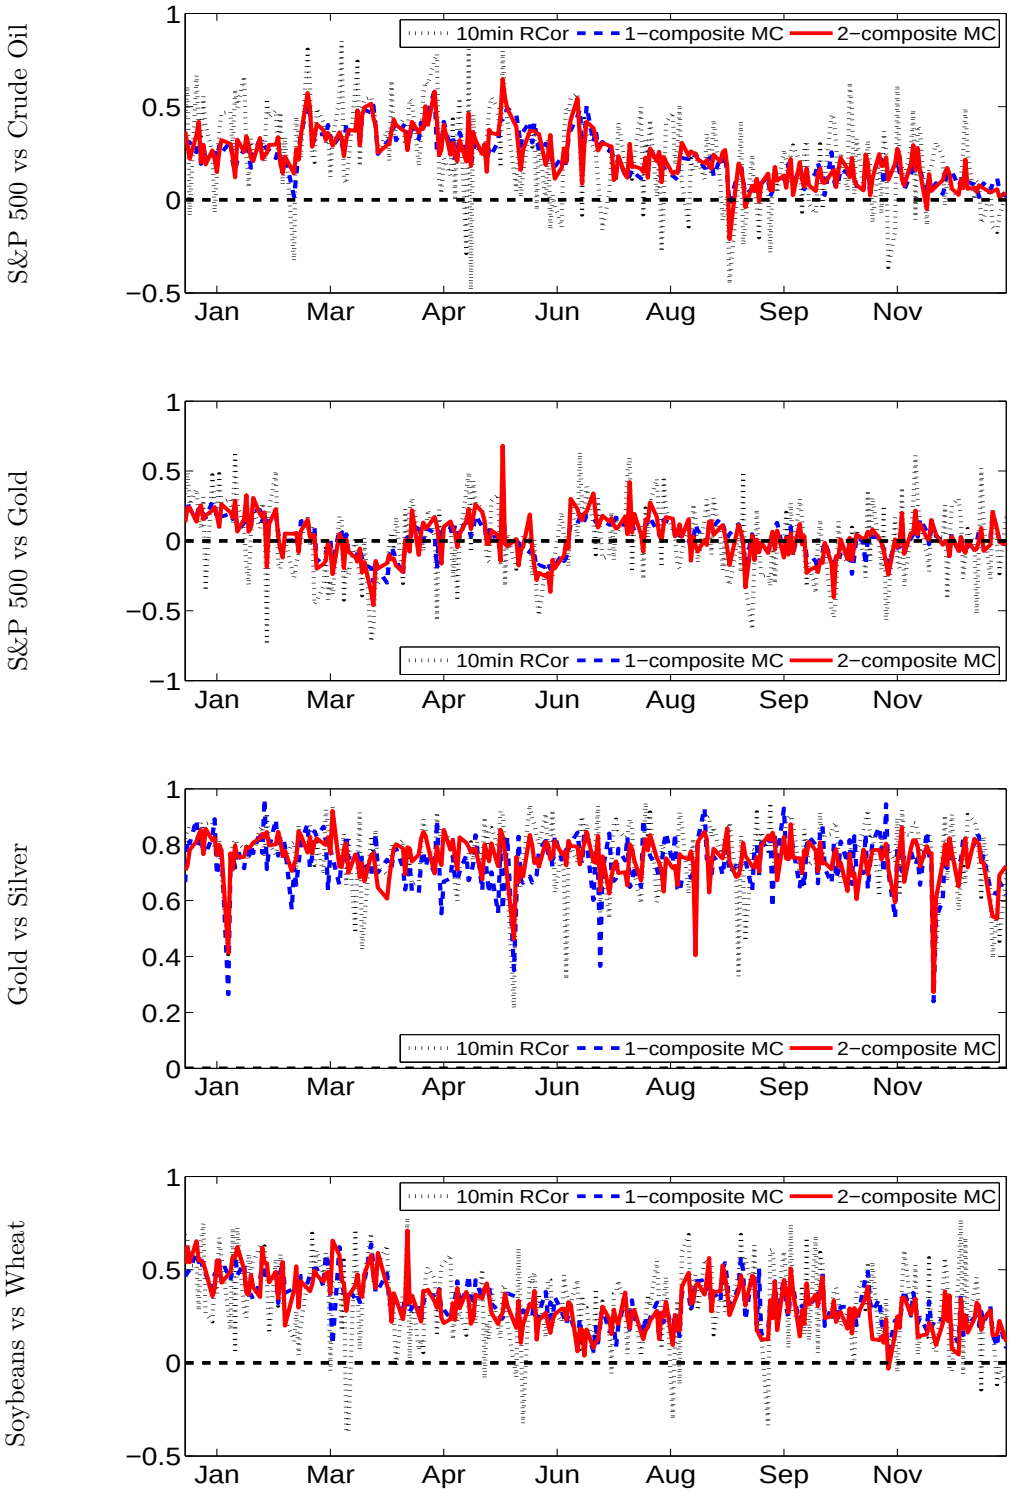

Figure 4.4: Realized correlation based on 10 min returns against correlation computed with MC estimator (order $k = 5$) in 2013.

Appendix D

Supplement to Chapter IV

D.1 Details on the Simulation Design with Stochastic Volatility

For comparison with the simulation design with constant volatility we seek to have the integrated variance be 1, in expectation. This is achieved as follows. Note that

$$\mathrm{E}((\mathrm{d}\log Y_{i,t})^2) = \mathrm{E}\left[\exp\left\{2\left(\beta_0 + \beta_1\tau_{i,t}\right)\right\}\Delta\right],$$

and since we (approximately) have $\tau_{i,t} \sim N(0, a^2)$ with

$$a^2 = \frac{\frac{1-\exp(2\alpha\Delta)}{-2\alpha}}{1-\exp(\alpha\Delta)^2} = \frac{1}{-2\alpha},$$

it follows that

$$2(\beta_0 + \beta_1\tau_{i,t}) \sim N(2\beta_0, 4\beta_1^2\frac{1}{-2\alpha}).$$

Hence

$$\mathrm{E}\left[\exp\left\{2\left(\beta_0 + \beta_1\tau_{i,t}\right)\right\}\right] = \exp\left(2\beta_0 + \beta_1^2\frac{1}{-\alpha}\right),$$

which will be equal to 1 if we set $\beta_0 = \beta_1^2/(2\alpha)$.

Bibliography

Admati, A. and Pfleiderer, P. (1988). A Theory of Intraday Patterns: Volume and Price Variability. *Review of Financial Studies*, 1(1):3–40.

Aielli, G. P. (2013). Dynamic conditional correlation: on properties and estimation. *Journal of Business & Economic Statistics*, 31:282–299.

Ait-Sahalia, Y., Fan, J., and Xiu, D. (2010). High-frequency covariance estimates with noisy and asynchronous financial data. *Journal of the American Statistical Association*, 105(492):1504–1517.

Ait-Sahalia, Y., Kalnina, I., and Xiu, D. (2014). The Idiosyncratic Volatility Puzzle: A Reassessment at High Frequency. Discussion paper, The University of Chicago.

Amisano, G. and Giacomini, R. (2007). Comparing Density Forecasts via Weighted Likelihood Ratio Tests. *Journal of Business & Economic Statistics*, 25:177–190.

Andersen, T., Dobrev, D., and Schaumburg, E. (2008). Duration-based volatility estimation. *working paper*.

Andersen, T. G. and Bollerslev, T. (1997). Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance*, 4(2-3):115–158.

Andersen, T. G. and Bollerslev, T. (1998a). Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts. *International Economic Review*, 39(4):885–905.

Andersen, T. G. and Bollerslev, T. (1998b). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. 39(4):885–905.

Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2007). Roughing It Up: Including Jump Components in the Measurement, Modeling, and Forecasting of Return Volatility. *The Review of Economics and Statistics*, 89(4):701–720.

Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2010). *Parametric and Nonparametric Volatility Measurement*. Handbook of Financial Econometrics. Elsevier.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Ebens, H. (2001a). The distribution of realized stock return volatility. *Journal of Financial Economics*, 61(1):43–76.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2001b). The Distribution of Realized Exchange Rate Volatility. *Journal of the American Statistical Association*, 96:42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003a). Modeling and Forecasting Realized Volatility. *Econometrica*, 71(2):579–625.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Vega, C. (2003b). Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange. *American Economic Review*, 93(1):38–62.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Wu, J. (2005). A Framework for Exploring the Macroeconomic Determinants of Systematic Risk. *American Economic Review*, 95(2):398–404.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Wu, J. (2006). *Realized Beta: Persistence and Predictability*, chapter 6, pages 1–39.

Andersen, T. G., Bollerslev, T., and Huang, X. (2011a). A reduced form framework for modeling volatility of speculative prices based on realized variation measures. *Journal of Econometrics*, 160(1):176–189.

Andersen, T. G., Bollerslev, T., and Meddahi, N. (2011b). Realized volatility forecasting and market microstructure noise. *Journal of Econometrics*, 160(1):220–234.

Ang, A., Liu, J., and Schwarz, K. (2010). Using Stocks or Portfolios in Tests of Factor Models. Working paper.

Aruoba, S. B., Diebold, F. X., and Scotti, C. (2009). Real-Time Measurement of Business Conditions. *Journal of Business & Economic Statistics*, 27(4):417–427.

Asai, M. and So, M. (2015). Long Memory and Asymmetry for Matrix-Exponential Dynamic Correlation Processes. *Journal of Time Series Econometrics*, 7(1):26.

Asness, C., Frazzini, A., and Pedersen, L. (2014). Quality minus junk. Working paper.

Bali, T. G., Engle, R. F., and Tang, Y. (2013). Dynamic Conditional Beta is Alive and Well in the Cross-Section of Daily Stock Returns. KoÃğ University-TUSIAD Economic Research Forum Working Papers 1305, Koc University-TUSIAD Economic Research Forum.

Balter, J., Dumitrescu, E., and Hansen, P. (2015). Forecasting Exchange Rate Volatility: Multivariate Realized GARCH Framework. Working paper.

Bandi, F. and Russell, J. (2005). Realized covariation, realized beta, and microstructure noise. Working paper, Graduate School of Business, University of Chicago.

Bandi, F. M. and Russell, J. R. (2008). Microstructure Noise, Realized Variance, and Optimal Sampling. *Review of Economic Studies*, 75:339–69.

Bannouh, K., Martens, M., Oomen, R., and van Dijk, D. (2012). Realized mixed-frequency factor models for vast dimensional covariance estimation. ERIM Report Series Research in Management ERS-2012-017-F&A, Erasmus Research Institute of Management (ERIM).

Banulescu, D., Hansen, P., Huang, Z., and Matei, M. (2014). Volatility During the Financial Crisis Through the Lens of High Frequency Data: A Realized GARCH Approach. Working paper.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2008). Designing Realized Kernels to Measure the ex post Variation of Equity Prices in the Presence of Noise. *Econometrica*, 76(6):1481–1536.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2009). Realized kernels in practice: trades and quotes. *Econometrics Journal*, 12(3):C1–C32.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2011a). Multivariate realised kernels: Consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading. *Journal of Econometrics*, 162(2):149–169.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A., and Shephard, N. (2011b). Subsampling realised kernels. *Journal of Econometrics*, 160(1):204–219.

Barndorff-Nielsen, O. E. and Shephard, N. (2002a). Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B*, pages 253–280.

Barndorff-Nielsen, O. E. and Shephard, N. (2002b). Estimating quadratic variation using realized variance. *Journal of Applied Econometrics*, 17(5):457–477.

Barndorff-Nielsen, O. E. and Shephard, N. (2002c). How accurate is the asymptotic approximation to the distribution of realised variance. Economics Series Working Papers 2001-W16, University of Oxford, Department of Economics.

Barndorff-Nielsen, O. E. and Shephard, N. (2004). Econometric Analysis of Realized Covariation: High Frequency Based Covariance, Regression, and Correlation in Financial Economics. *Econometrica*, 72(3):885–925.

Basak, S. and Pavlova, A. (2015). A Model of Financialization of Commodities. CEPR Discussion Papers 10651, C.E.P.R. Discussion Papers.

Bates, J. and Granger, C. (1969). The Combination of Forecasts. *Operations Research Quarterly*, 20:451–468.

Bauer, G. H. and Vorkink, K. (2011). Forecasting multivariate realized stock market volatility. *Journal of Econometrics*, 160(1):93–101.

Bauwens, L. and Storti, G. (2009). A Component GARCH Model with Time Varying Weights. *Studies in Nonlinear Dynamics & Econometrics*, 13(2):1–33.

Bauwens, L., Storti, G., and Violante, F. (2012). Dynamic conditional correlation models for realized covariance matrices. CORE Discussion Papers 2012060, Universite catholique de Louvain, Center for Operations Research and Econometrics (CORE).

Bekaert, G. and Wu, G. (2000). Asymmetric Volatility and Risk in Equity Markets. *Review of Financial Studies*, 13(1):1–42.

Bhardwaj, G., Gorton, G., and Rouwenhorst, G. (2015). Facts and Fantasies about Commodity Futures Ten Years Later. NBER Working Papers 21243, National Bureau of Economic Research, Inc.

Bianchi, F., Mumtaz, H., and Surico, P. (2009). The great moderation of the term structure of UK interest rates. *Journal of Monetary Economics*, 56(6):856–871.

Black, F. (1976). *Studies of Stock Price Volatility Changes*. Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section.

Black, F., Jensen, M., and Scholes, M. (1972). *The Capital Asset Pricing Model: Some Empirical Tests*. Jensen M. C. (ed.), Studies in the Theory of Capital Markets. Praeger, New York.

Blasques, F., Koopman, S. J., and Lucas, A. (2014). Optimal Formulations for Nonlinear Autoregressive Processes. Tinbergen Institute Discussion Papers 14-103/III, Tinbergen Institute.

Blume, M. E. (1970). Portfolio Theory: A Step Toward Its Practical Application. *The Journal of Business*, 43(2):152–73.

Bollerslev, T. (1990). Modelling the Coherence in Short-run Nominal Exchange Rates: A Multivariate Generalized ARCH Model. *The Review of Economics and Statistics*, 72(3):498–505.

Bollerslev, T., Cai, J., and Song, F. (2000). Intraday periodicity, long memory volatility, and macroeconomic announcement effects in the US treasury bond market. *Journal of Empirical Finance*, 7:37–55.

Bollerslev, T., Engle, R. F., and Wooldridge, J. M. (1988). A Capital Asset Pricing Model with Time-Varying Covariances. *Journal of Political Economy*, 96(1):116–31.

Bollerslev, T., Patton, A. J., and Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1):1–18.

Bollerslev, T. and Todorov, V. (2011). Tails, Fears, and Risk Premia. *Journal of Finance*, 66(6):2165–2211.

Bollerslev, T. and Zhang, B. Y. B. (2003). Measuring and modeling systematic risk in factor pricing models using high-frequency data. *Journal of Empirical Finance*, 10(5):533–558.

Bos, C. S., Janus, P., and Koopman, S. J. (2012). Spot Variance Path Estimation and Its Application to High-Frequency Jump Testing. *Journal of Financial Econometrics*, 10(2):354–389.

Boudt, K., Croux, C., and Laurent, S. (2011). Robust estimation of intraweek periodicity in volatility and jump detection. *Journal of Empirical Finance*, 18(2):353–367.

Braun, P. A., Nelson, D. B., and Sunier, A. M. (1995). Good News, Bad News, Volatility, and Betas. *Journal of Finance*, 50(5):1575–1603.

Brennan, M. J. and Schwartz, E. S. (1985). Evaluating Natural Resource Investments. *The Journal of Business*, 58(2):135–57.

Buyuksahin, B. and Robe, M. A. (2014). Speculators, commodities and cross-market linkages. *Journal of International Money and Finance*, 42(C):38–70.

Carhart, M. M. (1997). On Persistence in Mutual Fund Performance. *Journal of Finance*, 52(1):57–82.

Carriero, A., Clark, T., and Marcellino, M. (2014). No Arbitrage Priors, Drifting Volatilities, and the Term Structure of Interest Rates. CEPR Discussion Papers 9848, C.E.P.R. Discussion Papers.

Chiriac, R. and Voev, V. (2011). Modelling and forecasting multivariate realized volatility. *Journal of Applied Econometrics*, 26(6):922–947.

Chiu, T., Leonard, T., and Tsui, K.-W. (1996). The matrix-logarithmic covariance model. *Journal of the American Statistical Association*, 91(433):198–210.

Christensen, K., Oomen, R., and Renò, R. (2016). The Drift Burst Hypothesis. Working paper.

Christensen, K., Oomen, R. C., and Podolskij, M. (2008). Realised quantile-based estimation of the integrated variance. *working paper*.

Christensen, K., Oomen, R. C., and Podolskij, M. (2014). Fact or friction: Jumps at ultra high frequency. *Journal of Financial Economics*, 114(3):576–599.

Christensen, K. and Podolskij, M. (2007). Realized range-based estimation of integrated variance. *Journal of Econometrics*, 141(2):323–349.

Christie, A. A. (1982). The stochastic behavior of common stock variances : Value, leverage and interest rate effects. *Journal of Financial Economics*, 10(4):407–432.

Christoffersen, P., Lunde, A., and Olesen, K. V. (2014). Factor Structure in Commodity Futures Return and Volatility. CREATES Research Papers 2014-31, School of Economics and Management, University of Aarhus.

Clark, T. E. (2011). Real-Time Density Forecasts From Bayesian Vector Autoregressions With Stochastic Volatility. *Journal of Business & Economic Statistics*, 29(3):327–341.

Cochrane, J. (2001). *Asset pricing.* Princeton Univ. Press.

Cochrane, J. H. (1996). A Cross-Sectional Test of an Investment-Based Asset Pricing Model. *Journal of Political Economy*, 104(3):572–621.

Collier, P. and Goderis, B. (2012). Commodity prices and growth: An empirical investigation. *European Economic Review*, 56(6):1241–1260.

Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*, 7(2):174–196.

Corsi, F., Peluso, S., and Audrino, F. (2015). Missing in Asynchronicity: A Kalman-EM Approach for Multivariate Realized Covariance Estimation. *Journal of Applied Econometrics*, 30(3):377–97.

Corsi, F., Pirino, D., and Reno, R. (2010). Threshold bipower variation and the impact of jumps on volatility forecasting. *Journal of Econometrics*, 159(2):276–288.

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

Creal, D., Koopman, S. J., and Lucas, A. (2008). A General Framework for Observation Driven Time-Varying Parameter Models. Tinbergen Institute Discussion Papers 08-108/4, Tinbergen Institute.

Creal, D., Koopman, S. J., and Lucas, A. (2013). Generalized Autoregressive Score Models With Applications. *Journal of Applied Econometrics*, 28(5):777–795.

Creal, D., Koopman, S. J., Lucas, A., and Zamojski, M. (2015). Generalized Autoregressive Method of Moments. Tinbergen Institute Discussion Papers 15-138/III, Tinbergen Institute.

Creal, D. D. and Wu, J. C. (2014). Monetary Policy Uncertainty and Economic Fluctuations. NBER Working Papers 20594, National Bureau of Economic Research, Inc.

Delattre, S. and Jacod, J. (1997). A central limit theorem for normalized functions of the increments of a diffusion process, in the presence of round-off errors. *Bernoulli*, 3:1–28.

Diebold, F. X. and Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics*, 130(2):337–364.

Diebold, F. X. and Lopez, J. A. (1996). Forecast evaluation and combination. In Maddala, G. and Rao, C., editors, *Handbook of Statistics*, volume Volume 14: Statistical Methods in Finance, pages 241–268. Amsterdam: North-Holland.

Diebold, F. X., Rudebusch, G. D., and Boragan Aruoba, S. (2006). The macroeconomy and the yield curve: a dynamic latent factor approach. *Journal of Econometrics*, 131(1-2):309–338.

Dobrev, D. and Szerszen, P. J. (2010). The information content of high-frequency data for estimating equity return models and forecasting risk. International Finance Discussion Papers 1005, Board of Governors of the Federal Reserve System (U.S.).

Dumitrescu, E. and Hansen, P. R. (2013). Parameter Estimation with Out-of-Sample Objective. Working papers.

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods: Second Edition*. Number 9780199641178 in OUP Catalogue. Oxford University Press.

Engle, R. (2002). Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models. *Journal of Business & Economic Statistics*, 20(3):339–50.

Engle, R. and Kelly, B. (2011). Dynamic Equicorrelation. *Journal of Business & Economic Statistics*, 30(2):212–228.

Engle, R. F. (2012). Dynamic Conditional Beta. NYU Working Paper 2451/31582.

Engle, R. F. and Gallo, G. M. (2006). A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics*, 131(1-2):3–27.

Engle, R. F., Ledoit, O., and Wolf, M. (2016). Large dynamic covariance matrices. ECON - Working Papers 231, Department of Economics - University of Zurich.

Engle, R. F. and Ng, V. K. (1993). Measuring and Testing the Impact of News on Volatility. *Journal of Finance*, 48(5):1749–78.

Engle, R. F., Ng, V. K., and Rothschild, M. (1990). Asset pricing with a factor-arch covariance structure : Empirical estimates for treasury bills. *Journal of Econometrics*, 45(1-2):213–237.

Engle, R. F. and Sheppard, K. (2001). Theoretical and Empirical properties of Dynamic Conditional Correlation Multivariate GARCH. NBER Working Papers 8554, National Bureau of Economic Research, Inc.

Engle, R. F. and Sokalska, M. E. (2012). Forecasting intraday volatility in the US equity market. Multiplicative component GARCH. *Journal of Financial Econometrics*, 10(1):54–83.

Epps, T. W. (1979). Comovements in stock prices in the very short run. *Journal of the American Statistical Association*, 74(366):291–298.

Fama, E. F. and French, K. R. (1992). The Cross-Section of Expected Stock Returns. *Journal of Finance*, 47(2):427–65.

Fama, E. F. and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56.

Fama, E. F. and French, K. R. (1997). Industry costs of equity. *Journal of Financial Economics*, 43(2):153–193.

Fama, E. F. and French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22.

Fama, E. F. and MacBeth, J. D. (1973). Risk, Return, and Equilibrium: Empirical Tests. *Journal of Political Economy*, 81(3):607–36.

Fan, J., Furger, A., and Xiu, D. (2015). Incorporating Global Industrial Classification Standard into Portfolio Allocation: A Simple Factor-Based Large Covariance Matrix Estimator with High Frequency Data. Working paper no. 15-01, Fama-Miller Center for Research in Finance, The University of Chicago, Booth School of Business.

Ferson, W. E. and Harvey, C. R. (1991). The Variation of Economic Risk Premiums. *Journal of Political Economy*, 99(2):385–415.

Frijns, B. and Margaritis, D. (2008). Forecasting daily volatility with intraday data. *The European Journal of Finance*, 14(6):523–540.

Gallo, G. M. and Otranto, E. (2015). Forecasting realized volatility with changing average levels. *International Journal of Forecasting*, 31(3):620–634.

Gallo, G. M. and Otranto, E. (2016). Combining Markov Switching and Smooth Transition in Modeling Volatility: A Fuzzy Regime MEM. Econometrics working papers archive.

Geman, H. (2005). Commodities and commodity derivatives : modeling and pricing for agriculturals, metals and energy. Economics Papers from University Paris Dauphine 123456789/607, Paris Dauphine University.

Genizi, A. (1993). Decomposition of R2 in multiple regression with correlated regressors. *Statistica Sinica*, 3:407–420.

Geweke, J. and Amisano, G. (2010). Comparing and evaluating Bayesian predictive distributions of asset returns. *International Journal of Forecasting*, 26(2):216–230.

Geweke, J. and Amisano, G. (2011). Optimal prediction pools. *Journal of Econometrics*, 164(1):130–141.

Golosnoy, V., Gribisch, B., and Liesenfeld, R. (2012). The conditional autoregressive Wishart model for multivariate stock market volatility. *Journal of Econometrics*, 167(1):211–223.

Gorton, G. and Rouwenhorst, K. G. (2006). Facts and Fantasies about Commodity Futures. *Financial Analysts Journal*, 62(2):47–68.

Gospodinov, N., Kan, R., and Robotti, C. (2014). Misspecification-Robust Inference in Linear Asset-Pricing Models with Irrelevant Risk Factors. *Review of Financial Studies*, 27(7):2139–2170.

Grant, M. C., Boyd, S. P., and Ye, Y. (2014). *CVX: Matlab Software for Disciplined Convex Programming*.

Gromping, U. (2007). Estimators of Relative Importance in Linear Regression Based on Variance Decomposition. *The American Statistician*, 61:139–147.

Hall, S. G. and Mitchell, J. (2007). Combining density forecasts. *International Journal of Forecasting*, 23(1):1–13.

Hansen, L. P. (1982). Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica*, 50(4):1029–54.

Hansen, L. P. and Jagannathan, R. (1997). Assessing Specification Errors in Stochastic Discount Factor Models. *Journal of Finance*, 52(2):557–90.

Hansen, N., Lunde, A., Olesen, K., and Elst, H. (2014a). Realizing Commodity Correlations and the Market Beta. Creates research papers, School of Economics and Management, University of Aarhus.

Hansen, N. S. and Lunde, A. (2013). Analyzing Oil Futures with a Dynamic Nelson-Siegel Model. CREATES Research Papers 2013-36, School of Economics and Management, University of Aarhus.

Hansen, P. R. (2015). A martingale decomposition of discrete Markov chains. *Economics Letters*, 133(C):14–18.

Hansen, P. R. and Horel, G. (2009). Quadratic Variation by Markov Chains. CREATES Research Papers 2009-13, Department of Economics and Business Economics, Aarhus University.

Hansen, P. R. and Horel, G. (2014). Limit theory for the long-run variance of finite Markov chains. *Working paper*.

Hansen, P. R. and Huang, Z. (2016). Exponential garch modeling with realized measures of volatility. *Journal of Business & Economic Statistics*, 34(2):269–287.

Hansen, P. R., Huang, Z., and Shek, H. H. (2012). Realized GARCH: a joint model for returns and realized measures of volatility. *Journal of Applied Econometrics*, 27(6):877–906.

Hansen, P. R., Janus, P., and Koopman, S. J. (2016). Realized Wishart-GARCH: A Score-driven Multi-Asset Volatility Model. Tinbergen Institute Discussion Papers 16-061/III, Tinbergen Institute.

Hansen, P. R., Large, J., and Lunde, A. (2008). Moving average-based estimators of integrated variance. *Econometric Reviews*, 27:79–111.

Hansen, P. R. and Lunde, A. (2005). A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data. *Journal of Financial Econometrics*, 3(4):525–554.

Hansen, P. R. and Lunde, A. (2006a). Consistent ranking of volatility models. *Journal of Econometrics*, 131(1-2):97–121.

Hansen, P. R. and Lunde, A. (2006b). Realized Variance and Market Microstructure Noise. *Journal of Business & Economic Statistics*, 24:127–161.

Hansen, P. R. and Lunde, A. (2011). *Forecasting Volatility Using High-Frequency Data*. The Oxford Handbook of Economic Forecasting.

Hansen, P. R. and Lunde, A. (2014). Estimating The Persistence And The Autocorrelation Function Of A Time Series That Is Measured With Error. *Econometric Theory*, 30(01):60–93.

Hansen, P. R., Lunde, A., and Nason, J. M. (2011). The Model Confidence Set. *Econometrica*, 79(2):453–497.

Hansen, P. R., Lunde, A., and Voev, V. (2014b). Realized Beta Garch: A Multivariate Garch Model With Realized Measures Of Volatility. *Journal of Applied Econometrics*, 29(5):774–799.

Harvey, A. (2013). *Dynamic Models for Volatility and Heavy Tails: with Applications to Financial and Economic Time Series*. Econometric society monographs. Cambridge University Press.

Harvey, A., Ruiz, E., and Sentana, E. (1992). Unobserved component time series models with Arch disturbances. *Journal of Econometrics*, 52(1-2):129–157.

Harvey, A. C. (1991). *Forecasting. Structural Time Series Models and the Kalman Filter*. Cambridge University Press.

Hassine, M. and Roncalli, T. (2013). Measuring Performance of Exchange Traded Funds. MPRA Paper 44298, University Library of Munich, Germany.

Hautsch, N., Kyj, L. M., and Oomen, R. C. A. (2012). A blocking and regularization approach to high-dimensional realized covariance estimation. *Journal of Applied Econometrics*, 27(4):625–645.

Hautsch, N. and Ou, Y. (2012). Analyzing interest rate risk: Stochastic volatility in the term structure of government bond yields. *Journal of Banking & Finance*, 36(11):2988–3007.

Hodrick, R. J. and Zhang, X. (2001). Evaluating the specification errors of asset pricing models. *Journal of Financial Economics*, 62(2):327–376.

Horel, G. (2007). *Estimating Integrated Volatility with Markov Chains*. PhD thesis, Stanford University.

Ishihara, T., Omori, Y., and Asai, M. (2014). Matrix Exponential Stochastic Volatility with Cross Leverage. CIRJE F-Series CIRJE-F-932, CIRJE, Faculty of Economics, University of Tokyo.

Jacod, J., Li, Y., Mykland, P. A., Podolskij, M., and Vetter, M. (2009). Microstructure noise in the continuous case: The pre-averaging approach. *Stochastic Processes and their Applications*, 119(7):2249–2276.

Jagannathan, R., Skoulakis, G., and Wang, Z. (2002). Generalized Method of Moments: Applications in Finance. *Journal of Business & Economic Statistics*, 20(4):470–81.

Jagannathan, R. and Wang, Z. (1996). The Conditional CAPM and the Cross-Section of Expected Returns. *Journal of Finance*, 51(1):3–53.

Kalnina, I. and Tewou, K. (2015). Cross-sectional Dependence in Idiosyncratic Volatility. Cahiers de recherche 08-2015, Centre interuniversitaire de recherche en economie quantitative, CIREQ.

Kan, R. and Zhang, C. (1999). GMM tests of stochastic discount factor models with useless factors. *Journal of Financial Economics*, 54(1):103–127.

Karstanje, D., van der Wel, M., and van Dijk, D. (2015). Common Factors in Commodity Futures Curves. Technical report.

Kawakatsu, H. (2006). Matrix exponential GARCH. *Journal of Econometrics*, 134(1):95–128.

Kemeny, J. and Snell, J. (1976). *Finite Markov Chains*. Springer.

Kilian, L. (2008). The Economic Effects of Energy Price Shocks. *Journal of Economic Literature*, 46(4):871–909.

Kilian, L. (2009). Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market. *American Economic Review*, 99(3):1053–69.

Kleibergen, F. (2009). Tests of risk premia in linear factor models. *Journal of Econometrics*, 149(2):149–173.

Koopman, S. J., Jungbacker, B., and Hol, E. (2005). Forecasting daily variability of the S&P 100 stock index using historical, realised and implied volatility measurements. *Journal of Empirical Finance*, 12(3):445–475.

Koopman, S. J., Lucas, A., and Scharth, M. (2015). Predicting time-varying parameters with parameter-driven and observation-driven models. *Review of Economics and Statistics*, Forthcoming.

Koopman, S. J., Mallee, M. I. P., and Van der Wel, M. (2010). Analyzing the Term Structure of Interest Rates Using the Dynamic Nelson-Siegel Model With Time-Varying Parameters. *Journal of Business & Economic Statistics*, 28(3):329–343.

Koopman, S. J. and Scharth, M. (2012). The Analysis of Stochastic Volatility in the Presence of Daily Realized Measures. *Journal of Financial Econometrics*, 11(1):76–115.

Laurent, S., Bauwens, L., and Rombouts, J. V. K. (2006). Multivariate GARCH models: a survey. *Journal of Applied Econometrics*, 21(1):79–109.

Ledoit, O., Santa-Clara, P., and Wolf, M. (2003). Flexible Multivariate GARCH Modeling with an Application to International Stock Markets. *Review of Economics and Statistics*, 85:735–747.

Lewellen, J. and Nagel, S. (2006). The conditional CAPM does not explain asset-pricing anomalies. *Journal of Financial Economics*, 82(2):289–314.

Lewellen, J., Nagel, S., and Shanken, J. (2010). A skeptical appraisal of asset pricing tests. *Journal of Financial Economics*, 96(2):175–194.

Li, Y. and Mykland, P. A. (2006). Determining the volatility of a price process in the presence of rounding errors. Working paper.

Liu, L. Y., Patton, A. J., and Sheppard, K. (2015). Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes. *Journal of Econometrics*, 187(1):293–311.

Liu, Q. (2009). On portfolio optimization: How and when do we benefit from high-frequency data? *Journal of Applied Econometrics*, 24(4):560–582.

Lopez, J. A. (2001). Evaluating the Predictive Accuracy of Volatility Models. *Journal of Forecasting*, 20(2):87–109.

Lucas, A. and Zhang, X. (2016). Score-driven exponentially weighted moving averages and Value-at-Risk forecasting. *International Journal of Forecasting*, 32(2):293–302.

Lunde, A. and Olesen, K. V. (2014). Modeling and Forecasting the Distribution of Energy Forward Returns - Evidence from the Nordic Power Exchange. CREATES Research Papers 2013-19, School of Economics and Management, University of Aarhus.

Lunde, A., Shephard, N., and Sheppard, K. (2014). Econometric analysis of vast covariance matrices using composite realized kernels. *working paper*.

Maheu, J. M. and McCurdy, T. H. (2002). Nonlinear features of realized FX volatility. 84:668–681.

Mancini, C. (2001). Disentangling the jumps of the diffusion in a geometric jumping Brownian motion. *Giornale dell'Istituto Italiano degli Attuari*, 64:19–47.

Mancini, C. (2009). Non-parametric Threshold Estimation for Models with Stochastic Diffusion Coefficient and Jumps. *Scandinavian Journal of Statistics*, 36(2):270–296.

Martens, M. and van Dijk, D. (2007). Measuring volatility with the realized range. *Journal of Econometrics*, 138(1):181–207.

McAleer, M. and Medeiros, M. (2008). Realized Volatility: A Review. *Econometric Reviews*, 27(1-3):10–45.

Meddahi, N. (2002). A theoretical comparison between integrated and realized volatility. *Journal of Applied Econometrics*, 17(5):479–508.

Mincer, J. A. and Zarnowitz, V. (1969). The Evaluation of Economic Forecasts. In *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, NBER Chapters, pages 3–46. National Bureau of Economic Research, Inc.

Muth, J. F. (1960). Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association*, 22(4):299–306.

Nelson, C. R. and Siegel, A. F. (1987). Parsimonious Modeling of Yield Curves. *The Journal of Business*, 60(4):473–89.

Newey, W. K. and West, K. D. (1987). A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica*, 55(3):703–08.

Noureldin, D., Shephard, N., and Sheppard, K. (2012). Multivariate high-frequency-based volatility (HEAVY) models. *Journal of Applied Econometrics*, 27(6):907–933.

Pagan, A. R. and Schwert, G. W. (1990). Alternative models for conditional stock volatility. *Journal of Econometrics*, 45(1-2):267–290.

Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1):246–256.

Patton, A. J. and Sheppard, K. (2015). Good Volatility, Bad Volatility: Signed Jumps and The Persistence of Volatility. *The Review of Economics and Statistics*, 97(3):683–697.

Petrella, I. and Delle Monache, D. (2016). Adaptive models and heavy tails. Bank of England working papers 577, Bank of England.

Protter, P. (2004). *Stochastic Integration and Differential Equations.* New York: Springer-Verlag.

Ravazzolo, F. and Lombardi, M. J. (2012). Oil price density forecasts: Exploring the linkages with stock markets. Working Papers 0008, Centre for Applied Macro- and Petroleum economics (CAMP), BI Norwegian Business School.

Renò, R. (2003). A closer look at the epps effect. *International Journal of Theoretical and Applied Finance*, 6:87–102.

Roll, R. (1984). A Simple Implicit Measure of the Effective Bid-Ask Spread in an Efficient Market. *Journal of Finance*, 39(4):1127–39.

Samuelson, P. (1965). Proof that properly anticipated prices fluctates randomly. *Industrial management review*, 6(Spring):13–31.

Schwartz, E. and Smith, J. E. (2000). Short-Term Variations and Long-Term Dynamics in Commodity Prices. *Management Science*, 46(7):893–911.

Schwartz, E. S. (1997). The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging. *Journal of Finance*, 52(3):923–73.

Shephard, N. and Sheppard, K. (2010). Realising the future: forecasting with high-frequency-based volatility (HEAVY) models. *Journal of Applied Econometrics*, 25(2):197–231.

Shephard, N., Sheppard, K., and Engle, R. F. (2008). Fitting vast dimensional time-varying covariance models. Economics Series Working Papers 403, University of Oxford, Department of Economics.

Shephard, N. and Xiu, D. (2012). Econometric analysis of multivariate realised QML: efficient positive semi-definite estimators of the covariation of equity prices. Economics Series Working Papers 604, University of Oxford, Department of Economics.

Shin, M. and Zhong, M. (2015). Does Realized Volatility Help Bond Yield Density Prediction? Finance and economics discussion series 2015-115, Washington: Board of Governors of the Federal Reserve System.

Silvennoinen, A. and Terasvirta, T. (2008). Multivariate GARCH models. CREATES Research Papers 2008-06, Department of Economics and Business Economics, Aarhus University.

Stock, J. and Watson, M. (2011). Dynamic factor models. *The Oxford Handbook of Economic Forecasting*.

Takahashi, M., Omori, Y., and Watanabe, T. (2009). Estimating stochastic volatility models using daily returns and realized volatility simultaneously. *Computational Statistics & Data Analysis*, 53(6):2404–2426.

Timmermann, A. (2006). *Forecast Combinations*, volume 1 of *Handbook of Economic Forecasting*, chapter 4, pages 135–196. Elsevier.

Tse, Y. K. and Tsui, A. K. C. (2002). A Multivariate Generalized Autoregressive Conditional Heteroscedasticity Model with Time-Varying Correlations. *Journal of Business & Economic Statistics*, 20(3):351–62.

Wasserfallen, W. and Zimmermann, H. (1985). The behavior of intra-daily exchange rates. *Journal of Banking & Finance*, 9(1):55–72.

Weigand, R. (2014). Matrix Box-Cox Models for Multivariate Realized Volatility. University of Regensburg Working Papers in Business, Economics and Management Information Systems 478, University of Regensburg, Department of Economics.

White, H. (1982). Maximum Likelihood Estimation of Misspecified Models. *Econometrica*, 50(1):1–25.

Williams, P. M. (1999). Matrix logarithm parameterizations for regularized neural network covariance models. *Neural Networks*, 12:299–308.

Zebedee, A., Bentzen, E., Hansen, P., and Lunde, A. (2008). The Greenspan years: an analysis of the magnitude and speed of the equity market response to FOMC announcements. *Financial Markets and Portfolio Management*, 22(1):3–20.

Zhang, L. (2006). Efficient estimation of stochastic volatility using noisy observations: a multi-scale approach. *Bernoulli*, 12:1019–1043.

Zhang, L., Mykland, P. A., and Ait-Sahalia, Y. (2005). A Tale of Two Time Scales: Determining Integrated Volatility With Noisy High-Frequency Data. *Journal of the American Statistical Association*, 100:1394–1411.

Zhou, B. (1996). High-frequency data and volatility in foreign exchange rates. 14(1):45–52.

Zhou, B. (1998). Parametric and nonparametric volatility measurement. In Dunis, C. L. and Zhou, B., editors, *Nonlinear Modelling of High Frequency Financial Time Series*, chapter 6, pages 109–123. John Wiley Sons Ltd.