

Estimating the Conditional CAPM with Overlapping Data Inference

Esben Hedegaard*

Robert J. Hodrick†

W.P. Carey School of Business

Graduate School of Business

Arizona State University

Columbia University

and NBER

February 25, 2013

We thank Eric Ghysels, Christian Lundblad, Alberto Plazzi, Rossen Valkanov, Jules van Binsbergen and the seminar participants at the Stanford Graduate School of Business for helpful comments; and we thank Norman White for research computing assistance. This research was supported by a grant from the Network for Study on Pensions, Aging, and Retirement to the Columbia Business School.

*esben.hedegaard@asu.edu.

†rh169@columbia.edu.

Abstract

Asset pricing models such as the conditional CAPM are typically estimated with MLE using a monthly or quarterly horizon with data sampled to match the horizon even though daily data are available. We develop an overlapping data inference methodology (ODIN) that uses all of the data while maintaining the monthly or quarterly forecasting period, and we apply it to the conditional CAPM. Our approach recognizes that the first order conditions of MLE can be used as orthogonality conditions of GMM. We simulate from GARCH and MIDAS models and examine the substantial reductions in standard errors and increases in power that arise from our methodology. Using historical data, we find considerable differences in the estimates from the non-overlapping samples that begin on different days. Using our overlapping data inference, we find a significant risk-return trade-off in the monthly data from 1955 to 2011 with a symmetric GARCH model.

1 Introduction

This paper demonstrates how to use all the available data in estimating models such as the conditional capital asset pricing model (CAPM) and explores the resulting improvements in power. We thus extend the analysis of Hansen and Hodrick (1980), who introduced the idea of what we call overlapping data inference (ODIN), to situations in which researchers typically sample the data to use maximum likelihood estimation (MLE). We find substantial improvements in power that translate into large increases in effective sample sizes of non-overlapping data. Whenever the data are sampled, alternative non-overlapping estimates can be generated by shifting the starting date, and we compare the one ODIN estimation to the various estimates generated in this way.

We use the conditional CAPM as our laboratory. The conditional CAPM is the simplest version of Merton's (1973) intertemporal capital asset pricing model (ICAPM) which demonstrates that expected returns on assets respond dynamically to changes in the investment opportunity set. The conditional CAPM implies that the conditional expected returns on assets depend on their conditional covariances with the market return. Because the market return is also priced by the conditional CAPM, the model implies a linear trade-off between the conditional expected return on the market and the conditional variance of the market, as in Merton's (1980) Model 1. The first econometric analyses of the ICAPM thus focused on estimation of this risk-return trade-off. Ghysels, Santa Clara, and Valkanov (2005), Lundblad (2007), Lettau and Ludvigson (2010), and Nyberg (2012) provide extensive references to some of the vast empirical literature that investigates this risk-return trade-off.

Although the ICAPM was developed as a continuous time model, no one seriously thinks that agents act continuously, and empirical investigations of the ICAPM are typically done using monthly or quarterly returns. Market microstructure frictions, non-synchronous portfolio investment decisions, and individual stock illiquidity are features of the data that induce

difficult short-run serial correlation issues that are outside the realm of the theoretical model. Indeed, Merton (1980, p. 336) proposed a one-month time interval as “not an unreasonable choice” for the horizon to examine the predictions of the model. While we agree that a monthly or quarterly horizon is a reasonable frequency for testing the conditional CAPM or the ICAPM, daily data are available, and estimation with monthly or quarterly calendar returns is clearly arbitrary. An interval labeled a ‘month’ need not start on the first day of a calendar month.

Many researchers estimate the conditional CAPM using variants of Bollerslev’s (1985) generalized autoregressive conditional heteroskedastic (GARCH) model appropriately modified to incorporate a variance-in-mean term (GARCH-M), as in Engle, Lilien, and Robbins (1987) and French, Schwert, and Stambaugh (1987). More recently, Ghysels, Santa Clara, and Valkanov (2005) introduce the mixed data sampling (MIDAS) approach as an alternative to GARCH.¹ The empirical evidence in support of the conditional CAPM can only be described as decidedly mixed, with some studies finding positive and significant estimates of the risk-return trade-off and others finding insignificant or negative coefficients. Perhaps this should be unsurprising in light of Lundblad’s (2007) simulations that show the difficulty of finding a significant risk-return trade-off in the available sample sizes when such a trade-off actually exists.

In these GARCH and MIDAS models, researchers typically sample returns as calendar months or quarters. We observe daily data on the market return from 1926 to the present, and if we think of the typical month (quarter) as having 22 (66) working days, there are 22 (66) different starting days for estimation that uses a one-month (one-quarter) sampling interval. While it would be quite difficult to formally model the conditional distributions of the errors

¹Ghysels, Santa Clara, and Valkanov (2005) found a positive, highly statistically significant estimate of the risk-return trade-off. When we were unable to reproduce their results, we contacted them and found an error in their MatLab code. Ghysels, Plazzi, and Valkanov (2013) revise the estimation and extend it to include a discussion of the effects of financial crises on the risk-return trade-off.

within MLE that uses all of the daily data and the one-month or one-quarter forecasting interval, our ODIN approach allows estimation with all of the data constraining the various non-overlapping samples to have the same coefficients. We do this by viewing the first order conditions of the MLE as orthogonality conditions in Hansen's (1982) Generalized Method of Moments (GMM) that should be satisfied simultaneously by the same set of parameters for each of the different starting days.

The plan of the paper is as follows. Section 2 motivates the study of the conditional CAPM, discusses the choices of horizon and assets, notes that including a constant in the conditional mean is necessary to test the prediction of the model, and introduces the original MIDAS model as well as a modified version that allows simulations. Section 3 presents our ODIN methodology. Section 4 examines power issues in testing the conditional CAPM, first within the basic models and then in the ODIN versions. Section 5 examines estimation of the GARCH and MIDAS models with calendar-sampled data. Section 6 presents the estimation of the ODIN models, documents the differences between the results from the one ODIN model and those of the various non-overlapping monthly models generated by shifting the sampling start date, and examines the estimator that is the average of the non-overlapping estimates. Section 7 provides conclusions, and an Appendix provides some technical details. An Online Appendix contains additional figures and technical details.

2 The Conditional CAPM

We first introduce the conditional CAPM and discuss the importance of examining explicitly conditional econometric models.² Fundamentally, theoretical models imply a link between

²The classical reference on the unconditional implications of conditional asset pricing models is Hansen and Richard (1987), who note that the market portfolio could be on the conditional mean-variance frontier, but not on the unconditional frontier. Those who are familiar with the ideas in Lewellen and Nagel (2006) could skip this motivational section as the general idea is that examining the unconditional implications of conditional models by assessing their average pricing errors is problematic when the betas of assets covary with the prices of risks.

first and second moments that must be estimated simultaneously. We present the argument in terms of the conditional CAPM, but it is more general and applies to multifactor models as well. Then, we discuss the choices of horizon and assets for testing the implications of the model. We then discuss the importance of including a constant in the conditional mean when testing the conditional CAPM. The section concludes with a description of the MIDAS model of the conditional variance.

2.1 The Importance of Conditional Models

Let $R_{i,t}$ denote the excess return on asset i , and let $R_{M,t}$ denote the excess return on the market. Then, the conditional CAPM postulates that the expected excess return on an asset depends on the conditional covariance of the return with the market return:

$$E_t(R_{i,t+1}) = \gamma C_t(R_{i,t+1}, R_{M,t+1}) \quad (1)$$

where γ parameterizes the risk-return trade-off. Since equation (1) also applies to the market return, we have

$$E_t(R_{M,t+1}) = \gamma \sigma_{M,t}^2 \quad (2)$$

which expresses the linear risk-return trade-off between the conditional expected excess return on the market and its conditional variance. Since the estimation of the conditional covariances and variances requires estimating at least one conditional mean, a straightforward implication of equations (1) and (2) is that a test of the theory should involve developing models that simultaneously estimate conditional means and variances.

Cochrane (2005) notes that equation (1) can be written in a conditional beta representation:

$$E_t(R_{i,t+1}) = \lambda_t \beta_{i,t} \quad (3)$$

where the conditional betas are given by

$$\beta_{i,t} \equiv C_t(R_{i,t+1}, R_{M,t+1})/\sigma_{M,t}^2$$

and the time-varying price of risk is given by

$$\lambda_t \equiv \gamma\sigma_{M,t}^2.$$

Taking the unconditional expectation of equation (3) gives

$$E(R_{i,t+1}) = E(\lambda_t\beta_{i,t}) = E(\lambda_t)E(\beta_{i,t}) + Cov(\lambda_t, \beta_{i,t}). \quad (4)$$

Equation (4) clearly demonstrates that if the conditional betas covary with the conditional price of risk, a cross-sectional regression of the average excess returns on estimates of the unconditional expectations of the $\beta_{i,t}$ s does not result in zero pricing errors.³ It is for this reason that we desire to directly test the conditional implications of equation (1).

2.2 The Choice of Horizon

Since we want to test explicitly conditional theoretical models, we have to choose the horizon for which to say the theory holds. While daily returns are available, most of the existing literature prefers to test the conditional CAPM using longer horizon monthly or quarterly returns. Researchers then sample the data so that the frequency of the data is the same as the forecast interval of the model. Two considerations motivate this choice of horizon.

First, aspects of the trading process induced by market microstructure frictions, non-

³Lewellen and Nagel (2006) also note that $E(\beta_{i,t})$ is not the same as the unconditional β_i that would be recovered in a time series multiple regression of $R_{i,t+1}$ on $R_{M,t+1}$. Thus, the usual Fama-MacBeth (1973) tests in which a sequence of cross-sectional regressions is run on the unconditional betas estimated over the full sample is not an appropriate test of the unconditional implications of a conditional model. Ang and Kristensen (2012) propose non-parametric methods to test the implications of conditional models.

synchronous portfolio investment decisions, and individual stock illiquidity that are outside the theory dominate the autocorrelations in short-horizon returns. More importantly, when more volatile trading environments arise, theory predicts that stock returns are expected to be contemporaneously negatively correlated with the increase in volatility because prices must fall to provide an increase in expected returns, as in Campbell and Hentschel (1992). If the adjustment of expected returns to news that increases the conditional variance is not precisely contemporaneously correlated with the increase in the conditional variance because of market illiquidity or the non-synchronous trading of investors, using a short horizon for testing the conditional risk-return trade-off may problematically find a negative relation as volatility increases and prices fall slightly later. Thus, researchers use a longer horizon to balance the theoretical idea that there is a risk-return trade-off over a particular horizon against the loss of power that arises from sampling the data. Then, the data are sampled to allow the use of MLE. The ODIN model attempts to improve this situation. The econometrician can use any forecast horizon while maximizing power from using all available data.

In Section 4 we show in simulations that if the data were actually generated by a GARCH process for which the risk-return relationship is present at the short-term (continuous) horizon, using daily data with a daily horizon would indeed provide the highest power. But, if it is desirable to set the forecast horizon to be a longer interval, which is preferred in practice, we demonstrate that the power of the ODIN model is superior to that of the basic model in which the frequency of observations is the forecast interval. We also find that the improvement in power increases with the length of the sampling interval. Here, though, we rely on large samples in which the overlap is a small fraction of the total sample size, so this statement should not be extrapolated literally. In what follows we use the term ‘basic’ to refer to models estimated with non-overlapping observations to distinguish them from ODIN models.

2.3 The Choice of Assets

Researchers must also choose which assets to use in measuring the risk-return trade-off of the conditional CAPM. Any set of assets could potentially be used. Although simulations with multiple assets calibrated to realistic data indicate that the ability to reject $\gamma = 0$ is enhanced by using multiple assets, we focus only on the market portfolio primarily because much of the literature uses the market as the only test asset even though the theory applies generally. Thus, we present the nature of the improvement that using overlapping data offers within a familiar context. We leave discussion of the improvements in power from using multiple assets to a later paper that will explore additional risk factors. Given that the only asset analyzed is the market return, we drop the M subscript on the market return and its conditional variance to simplify the notation.

2.4 On Including a Constant in the Risk-Return Trade-off

Lanne and Saikkonen (2006) explicitly advocate estimating the conditional CAPM without a constant, as in equation (2), and they find strong support for the conditional CAPM. Scruggs (1998) and Nyberg (2012) estimate both with and without a constant finding much higher significance of the risk-return trade-off without a constant. We now demonstrate that including a constant in the conditional mean is necessary to test the prediction of the conditional CAPM that the conditional mean of the market return is dynamically linked to its conditional variance, even though under the null hypothesis that the model is true, the constant is zero. Estimating without a constant simply relates the average future return to the average conditional variance, whereas if a constant is included, the estimate of γ will only be significantly different from zero if the covariance of the future return with its conditional variance is positive and significant.

To make the point most easily, we treat the conditional variance, σ_t^2 , as known, and we as-

sume that the market return is conditionally normally distributed, $R_{t+1} \sim N(E_t(R_{t+1}), \sigma_t^2)$, where $E_t(R_{t+1})$ is either specified by equation (2) or

$$E_t(R_{t+1}) = \mu + \gamma\sigma_t^2. \quad (5)$$

Without a constant, the relevant part of the log-likelihood function is

$$\sum_{t=1}^T -\frac{1}{2} \frac{(R_{t+1} - \gamma\sigma_t^2)^2}{\sigma_t^2}.$$

The first order condition for γ gives

$$\hat{\gamma} = \frac{\sum_{t=1}^T R_{t+1}}{\sum_{t=1}^T \sigma_t^2} = \frac{(\frac{1}{T}) \sum_{t=1}^T R_{t+1}}{(\frac{1}{T}) \sum_{t=1}^T \sigma_t^2} = \frac{\bar{R}_{t+1}}{\bar{\sigma}_t^2}$$

where \bar{R}_{t+1} and $\bar{\sigma}_t^2$ denote the sample means of the return and the conditional variance, respectively. Because the numerator and the denominator are both positive in the historical data, $\hat{\gamma}$ is positive and significant, but in no sense does finding a significant $\hat{\gamma}$ indicate support for the dynamic prediction of the conditional CAPM that the conditional mean increases when the conditional variance increases.

If the conditional mean contains a constant, the relevant part of the log-likelihood function is

$$\sum_{t=1}^T -\frac{1}{2} \frac{(R_{t+1} - \mu - \gamma\sigma_t^2)^2}{\sigma_t^2}.$$

The first order condition with respect to μ and γ are the following:

$$\sum_{t=1}^T \frac{(R_{t+1} - \hat{\mu} - \hat{\gamma}\sigma_t^2)}{\sigma_t^2} = 0$$

$$\sum_{t=1}^T (R_{t+1} - \hat{\mu} - \hat{\gamma}\sigma_t^2) = 0.$$

Solving for $\hat{\gamma}$ gives

$$\hat{\gamma} = \frac{\overline{\left(\frac{R_{t+1}}{\sigma_t^2}\right)} - \overline{R_{t+1}} \overline{\frac{1}{\sigma_t^2}}}{1 - \overline{\sigma_t^2} \overline{\frac{1}{\sigma_t^2}}}.$$

where bars above variables indicate sample averages. The solution for $\hat{\gamma}$ is the sample covariance of R_{t+1} with $1/\sigma_t^2$ divided by the sample covariance of σ_t^2 with $1/\sigma_t^2$. By Jensen's inequality, the denominator of $\hat{\gamma}$ is negative. Therefore, $\hat{\gamma}$ is positive if the covariance of future returns with the reciprocal of the conditional variance is negative. Thus, including the constant is necessary to test the prediction that the conditional mean responds positively to its conditional variance. Variability of the conditional variance is clearly necessary for identification of $\hat{\gamma}$. For this example, the MLE standard error of $\hat{\gamma}$ would be

$$s.e.(\hat{\gamma}) = \frac{\sqrt{\overline{\frac{1}{\sigma_t^2}}}}{\sqrt{T} \left[\overline{\sigma_t^2 \frac{1}{\sigma_t^2}} - 1 \right]^{1/2}}. \quad (6)$$

From equation (6), the more variable is the conditional variance, the smaller is the standard error, which is also what Lundblad (2007) finds in his simulations. Thus, rejecting the null hypothesis should be easier, the more variable is σ_t^2 . While this example is highly stylized because we assume that the conditional variance is known, the intuition is certainly correct. Including a constant in the conditional mean of the risk-return trade-off also captures possible effects of any omitted state-variables that might appear in alternative specifications of the ICAPM. Consequently, we include a constant in all of our estimations.

2.5 The MIDAS Framework

One of the challenges in estimating the risk-return trade-off is that the conditional variance is not observable, and the GARCH model is but one model for the conditional variance. This section first explains the MIDAS model of the conditional variance introduced by Ghysels,

Santa Clara, and Valkanov (2005). Because simulating from that model introduces unrealistic sample paths, we develop an alternative formulation of the MIDAS model that avoids this problem.

2.5.1 The Original MIDAS Model

Ghysels, Santa Clara, and Valkanov (2005) model the expected monthly excess return on the market as a linear function of the conditional variance of the monthly market return as in equation (5). The monthly conditional variance is modeled as a weighted average of past squared daily excess returns on the market, where lower case r 's refer to daily excess returns:⁴

$$\sigma_t^2 \equiv V_t^{MIDAS} = 22 \sum_{d=0}^D w_d r_{t-d}^2. \quad (7)$$

The weights on the past squared excess returns sum to one and are initially modeled as

$$w_d(\kappa_1, \kappa_2) = \frac{\exp(\kappa_1 d + \kappa_2 d^2)}{\sum_{j=0}^{250} \exp(\kappa_1 j + \kappa_2 j^2)}, d = 0, \dots, D. \quad (8)$$

The number of days into the past, D , is set at 250, and multiplying by 22 in equation (7) scales the variance to a monthly value. The numerators of the weights and thus the shape of the distributed lag function are exponential functions of two parameters, κ_1 and κ_2 , that must be estimated. While this specification allows for considerable flexibility in the shape of the distributed lag, as long as $\kappa_2 < 0$, the weights die out eventually as the days get further into the past. Ghysels, Santa Clara, and Valkanov (2005) use MLE to estimate the four parameters of the model, μ , γ , κ_1 and κ_2 , assuming that $R_{t+1} \sim N(E_t(R_{t+1}), V_t^{MIDAS})$, where the conditional mean and variance are given in equation (5) and (7). The approach is thus similar to a GARCH-M model but with a different specification of the conditional

⁴The mixture of monthly and daily data is the hallmark of the MIDAS approach. See Ghysels, Sinko, and Valkanov (2007) for an introduction to the MIDAS literature and examples of its use in other applications.

variance.

We note in the introduction that the results in Ghysels, Santa Clara, and Valkanov (2005) are incorrect, and we refer the reader to Ghysels, Plazzi, and Valkanov (2013) for corrected results and an update of their model. Their original sample periods use data from 1928 to 2000 for 850 observations and 1964 to 2000, for 450 observations.

To gain additional understanding of their modeling approach, we attempted to simulate from this MIDAS model. However, even though the unconditional variance of r_t is the same for all t , the conditional distribution of r_t becomes increasingly peaked as t grows large.⁵ This leads to unrealistic data because most simulated paths essentially ‘die out’ as the distribution of returns becomes more peaked around zero until a tail-event occurs, and the variance increases dramatically.

2.5.2 An Alternative MIDAS Specification

It is easy to modify the original MIDAS specification of the conditional variance to eliminate this problem. Inspired by the usual GARCH model, let the conditional variance for a horizon of K days be

$$V_t^{MIDAS} = K \left(\omega + \phi \sum_{d=0}^D w_d(\theta) r_{t-d}^2 \right) \quad (9)$$

where $\omega > 0, 0 < \phi < 1$, and as above, the weights, w_d , sum to 1. The unconditional variance of the daily return process is $E(\sigma^2) = \frac{\omega}{1-\phi}$. Using this specification results in realistic simulated sample paths.

⁵To see this, consider the simplest example of a MIDAS model: The variance is estimated based on 1 past return and is used to describe only the next day. For simplicity, consider the case where $\mu = \gamma = 0$. With r_t denoting daily returns, the model is

$$r_{t+1} = z_{t+1} \sqrt{V_t^{MIDAS}}, \quad z_{t+1} \sim N(0, 1), \quad V_t^{MIDAS} = r_t^2.$$

Let $V_0^{MIDAS} = \sigma_0^2$ be given. We can then explicitly write out the sequence of returns and show that $r_t = z_t \prod_{s=1}^{t-1} |z_s| \sigma_0$. As the z_t 's are independent, $E(r_t) = 0$ for all t , and $V(r_t) = E(r_t^2) = \sigma_0^2$. This shows that all returns have the same unconditional mean and variance. However, the distribution of r_t clearly changes with t , and a small shock will cause the conditional distribution to be very tight around zero.

While the two-parameter exponential weight function in equation (8) provides flexibility in the pattern of the weights on past squared returns, our experience with MIDAS indicated that estimates of κ_2 are often insignificantly different from 0 and are sometimes disturbingly positive, in which case the past weights do not die out.⁶ An alternative weight function can be parameterized with the one-parameter beta-polynomial:

$$w_d(\kappa) = \frac{f(u_d, \kappa)}{\sum_{j=0}^D f(u_j, \kappa)}, \quad f(u_d, \kappa) = u_d^{(\kappa-1)}, \quad d = 0, \dots, D \quad (10)$$

where $u_d = (1 - d/D)$ are the points where f is evaluated.⁷ We also set $D = 500$, as we find that the choice of lag length influences the estimate of γ , but that the variation is very small once D is larger than 300. In this MIDAS specification, the variance process is parameterized with three parameters as in the GARCH(1,1) model.

3 Overlapping Data Inference with GMM

This section derives our ODIN estimation strategy for the conditional CAPM. Throughout, we maintain the usual assumption in the literature that monthly or quarterly returns satisfy the conditional CAPM. The Appendix describes construction of the overlapping monthly or quarterly excess returns.

Consider the ODIN-GARCH estimation first. Instead of presenting a general case, for ease of exposition and consistency with the past literature, we assume that the model holds at the monthly frequency. Thus, the GARCH-M model is

$$R_{t_m+1} = \mu + \gamma \sigma_{t_m}^2 + \varepsilon_{t_m+1}, \quad \varepsilon_{t_m+1} \sim N(0, \sigma_{t_m}^2) \quad (11)$$

⁶One sees this problem in the new results of Ghysels, Plazzi, and Valkanov (2013).

⁷Ghysels, Sinko, and Valkanov (2007) discuss various weighting schemes including exponential Almon lags and beta functions. Our specification corresponds to a restricted version of their equation (4).

where

$$\sigma_{t_m+1}^2 = \omega + \alpha\varepsilon_{t_m+1}^2 + \beta\sigma_{t_m}^2. \quad (12)$$

The time index $t_m = 0, \dots, T_m - 1$ counts 22 day periods or approximately one month of trading days, and there are T_m of these monthly periods in the sample. Our notation distinguishes between monthly data indexed with t_m and daily data indexed with t . Capital R 's represent monthly excess returns, and lower case r 's represent daily returns. Hence, the excess monthly return, R_{t_m+1} , can also be represented as an excess return over 22 trading days, $R_{t+22,t}$.⁸ Although the model in equations (11) and (12) is specified at the monthly frequency, we assume that only the number of days in the forecast matters, in which case the starting date for the month does not matter. Thus, we can write

$$R_{t+22,t} = \mu + \gamma\sigma_{t,t+22}^2 + \varepsilon_{t+22,t}, \quad \varepsilon_{t+22,t} \sim N(0, \sigma_{t,t+22}^2) \quad (13)$$

where the notation indicates that the monthly model can also be written with daily subscripts, and $\varepsilon_{t+22,t}$ denotes the innovation in the monthly return realized on day $t + 22$. The conditional variance evolves as

$$\sigma_{t+22,t+44}^2 = \omega + \alpha\varepsilon_{t+22,t}^2 + \beta\sigma_{t,t+22}^2 \quad (14)$$

We assume that equations (13) and (14) hold *for all* $t = 0, 1, \dots, T - 1$.

For the MIDAS version of the model, monthly returns indexed by t_m are generated by

$$R_{t_m+1} = \mu + \gamma V_{t_m}^{MIDAS} + \varepsilon_{t_m+1}, \quad \varepsilon_{t_m+1} \sim N(0, V_{t_m}^{MIDAS}) \quad (15)$$

⁸The first subscript of a variable denotes the day that the variable enters the information set, and the second subscript is a second day that indicates the period of time, either from the past for a return or into the future for a conditional variance, that is necessary to describe the variable.

where

$$V_{t_m}^{MIDAS} = 22 \left(\omega + \phi \sum_{d=0}^D w_{t_m-d}(\kappa) r_{t_m,-d}^2 \right) \quad (16)$$

and $r_{t_m,-d}$ is the daily return d days before the start of month t_m . Again, we assume that equations (15) and (16) are true independently of the day we use as the starting date for the month. Thus, we have

$$R_{t+22,t} = \mu + \gamma V_{t,t+22}^{MIDAS} + \varepsilon_{t+22,t}, \quad \varepsilon_{t+22,t} \sim N(0, V_{t,t+22}^{MIDAS}) \quad (17)$$

where

$$V_{t,t+22}^{MIDAS} = 22 \left(\omega + \phi \sum_{d=0}^D w_{t-d}(\kappa) r_{t-d}^2 \right) \quad (18)$$

and equations (17) and (18) hold *for all* $t = 0, 1, \dots, T - 1$.

As a caveat to our analysis, it is not at all clear whether there exists a data generating process for daily returns that has the postulated return properties over the ‘monthly’ intervals. Our point is that if the model is viewed as an abstraction that holds at the monthly horizon better than it does at any other horizon, and if there is nothing particular about calendar months, the starting date becomes irrelevant. We thus have the opportunity to increase the power of the tests by using overlapping data. We should not be wedded to calendar months given that we observe daily data.

3.1 ODIN-GARCH

This section derives the first order conditions from the MLEs of the monthly models that must hold for each starting date. We use Hansen’s (1982) GMM to derive the asymptotic distribution of the parameter estimates that must satisfy the average of the monthly first order conditions. Because it is difficult to know how the serial correlation induced by the creation of overlapping observations affects small sample inference, we examine extensive

simulations in the sections that follow.

The log likelihood function for the GARCH-M model with a monthly sampling interval in equations (11) and (12) is

$$\log(L) = \sum_{t_m=0}^{T_m-1} \left(-\frac{1}{2} \log(2\pi) - \log(\sigma_{t_m}) - \frac{1}{2} \frac{\varepsilon_{t_m+1}^2}{\sigma_{t_m}^2} \right).$$

Rather than estimating ω as a free parameter, we estimate ω by variance targeting as in Engle and Kroner (1995).⁹ Consequently, the first order conditions of the MLE are the following:

$$\begin{aligned} \frac{\partial \log(L)}{\partial \mu} &= \sum_{t_m=0}^{T_m-1} \left(\frac{\varepsilon_{t_m+1}}{\sigma_{t_m}^2} + \frac{1}{\sigma_{t_m}} \frac{\partial \sigma_{t_m}}{\partial \mu} \left(\frac{\varepsilon_{t_m+1}^2}{\sigma_{t_m}^2} + 2\gamma \varepsilon_{t_m+1} - 1 \right) \right) = 0 \\ \frac{\partial \log(L)}{\partial \gamma} &= \sum_{t_m=0}^{T_m-1} \left(\varepsilon_{t_m+1} + \frac{1}{\sigma_{t_m}} \frac{\partial \sigma_{t_m}}{\partial \gamma} \left(\frac{\varepsilon_{t_m+1}^2}{\sigma_{t_m}^2} + 2\gamma \varepsilon_{t_m+1} - 1 \right) \right) = 0 \\ \frac{\partial \log(L)}{\partial \alpha} &= \sum_{t_m=0}^{T_m-1} \frac{1}{\sigma_{t_m}} \frac{\partial \sigma_{t_m}}{\partial \alpha} \left(\frac{\varepsilon_{t_m+1}^2}{\sigma_{t_m}^2} + 2\gamma \varepsilon_{t_m+1} - 1 \right) = 0 \\ \frac{\partial \log(L)}{\partial \beta} &= \sum_{t_m=0}^{T_m-1} \frac{1}{\sigma_{t_m}} \frac{\partial \sigma_{t_m}}{\partial \beta} \left(\frac{\varepsilon_{t_m+1}^2}{\sigma_{t_m}^2} + 2\gamma \varepsilon_{t_m+1} - 1 \right) = 0 \end{aligned}$$

Cochrane (2005) notes that the first order conditions of MLE are the sample counterparts to unconditional moment conditions in GMM estimation. We think of these orthogonality conditions as holding for the 22 possible daily starting dates that index the months associated with t_m . By utilizing GMM, we can use all of the daily data with these orthogonality conditions and appropriately account for the serial correlation induced by the overlapping data in constructing the GMM standard errors.

Let the parameter vector be $\theta = (\mu, \gamma, \alpha, \beta)$. Then, the vector of sample orthogonality

⁹Variance targeting guarantees that the unconditional estimate of the variance of a GARCH model equals the sample variance. Francq, Horvath, and Zakoian (2011) examine the econometric properties of this popular estimation strategy.

conditions with t indexing daily data is

$$G_T(R; \theta) = \frac{1}{T} \sum_{t=0}^{T-1} g_t(R_{t+1}; \theta) = \begin{pmatrix} \frac{1}{T} \sum_{t=0}^{T-1} \left(\frac{\varepsilon_{t+22,t}}{\sigma_{t,t+22}^2} + \frac{1}{\sigma_{t,t+22}} \frac{\partial \sigma_{t,t+22}}{\partial \mu} \left(\frac{\varepsilon_{t+22,t}^2}{\sigma_{t,t+22}^2} + 2\gamma \varepsilon_{t+22,t} - 1 \right) \right) \\ \frac{1}{T} \sum_{t=0}^{T-1} \left(\varepsilon_{t+22,t} + \frac{1}{\sigma_{t,t+22}} \frac{\partial \sigma_{t,t+22}}{\partial \gamma} \left(\frac{\varepsilon_{t+22,t}^2}{\sigma_{t,t+22}^2} + 2\gamma \varepsilon_{t+22,t} - 1 \right) \right) \\ \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{\sigma_{t,t+22}} \frac{\partial \sigma_{t,t+22}}{\partial \alpha} \left(\frac{\varepsilon_{t+22,t}^2}{\sigma_{t,t+22}^2} + 2\gamma \varepsilon_{t+22,t} - 1 \right) \\ \frac{1}{T} \sum_{t=0}^{T-1} \frac{1}{\sigma_{t,t+22}} \frac{\partial \sigma_{t,t+22}}{\partial \beta} \left(\frac{\varepsilon_{t+22,t}^2}{\sigma_{t,t+22}^2} + 2\gamma \varepsilon_{t+22,t} - 1 \right) \end{pmatrix} \quad (19)$$

where $g_t(R_{t+1}; \theta)$ denotes the vector of right-hand-side functions. Because the system of equations (19) is just identified, GMM chooses the parameter estimates, $\hat{\theta}$, to set $G_T(R, \hat{\theta}) = 0$. Intuitively, the parameters may be estimated by maximizing the average of the 22 monthly log-likelihood functions.

Let the gradient of the sample orthogonality conditions with respect to the parameters be

$$D_T(\hat{\theta}) = \nabla_{\theta} G_T(R; \hat{\theta}).$$

Then, the asymptotic distribution theory of Hansen's (1982) GMM implies that

$$\sqrt{T} \left(\hat{\theta} - \theta_0 \right) \rightarrow N \left(0, D_T(\hat{\theta})^{-1} S(\hat{\theta}) D_T^{\top}(\hat{\theta})^{-1} \right)$$

where

$$S(\hat{\theta}) = \sum_{j=-21}^{21} C_T \left(g_t(R_t; \hat{\theta}), g_{t-j}(R_{t-j}; \hat{\theta})^{\top} \right) \quad (20)$$

and the $C_T \left(g_t(R_t; \hat{\theta}), g_{t-j}(R_{t-j}; \hat{\theta})^{\top} \right)$ are the sample autocovariances of $g_t(R_t, \hat{\theta})$. Under the null hypothesis, these autocovariances will be non-zero until $j = 22$. Note that we equally weight the sample covariances, as in Hansen and Hodrick (1980), to estimate the spectral

density matrix, $S(\hat{\theta})$.

3.2 ODIN-MIDAS

The log likelihood function for the MIDAS model specified at the monthly horizon is

$$\log(L) = \sum_{t_m=0}^{T_m-1} \left(-\frac{1}{2} \log(2\pi) - \log \left(\sqrt{V_{t_m}^{MIDAS}} \right) - \frac{1}{2} \frac{\varepsilon_{t_m+1}^2}{V_{t_m}^{MIDAS}} \right).$$

We again estimate ω by variance targeting, such that the first order conditions are the following:

$$\begin{aligned} \frac{\partial \log(L)}{\partial \mu} &= \sum_{t_m=0}^{T_m-1} -\frac{\varepsilon_{t_m+1}}{V_{t_m}^{MIDAS}} \frac{\partial \varepsilon_{t_m+1}}{\partial \mu} = \sum_{t_m=0}^{T_m-1} \frac{\varepsilon_{t_m+1}}{V_{t_m}^{MIDAS}} = 0 \\ \frac{\partial \log(L)}{\partial \gamma} &= \sum_{t_m=0}^{T_m-1} -\frac{\varepsilon_{t_m+1}}{V_{t_m}^{MIDAS}} \frac{\partial \varepsilon_{t_m+1}}{\partial \gamma} = \sum_{t_m=0}^{T_m-1} \varepsilon_{t_m+1} = 0 \\ \frac{\partial \log(L)}{\partial \phi} &= \sum_{t_m=0}^{T_m-1} \left(\frac{-1}{2V_{t_m}^{MIDAS}} + \gamma \frac{\varepsilon_{t_m+1}}{V_{t_m}^{MIDAS}} + \frac{1}{2} \frac{\varepsilon_{t_m+1}^2}{(V_{t_m}^{MIDAS})^2} \right) \frac{\partial V_{t_m}^{MIDAS}}{\partial \phi} = 0 \\ \frac{\partial \log(L)}{\partial \kappa} &= \sum_{t_m=0}^{T_m-1} \left(\frac{-1}{2V_{t_m}^{MIDAS}} + \gamma \frac{\varepsilon_{t_m+1}}{V_{t_m}^{MIDAS}} + \frac{1}{2} \frac{\varepsilon_{t_m+1}^2}{(V_{t_m}^{MIDAS})^2} \right) \frac{\partial V_{t_m}^{MIDAS}}{\partial \kappa} = 0 \end{aligned}$$

Notice that these first order conditions for μ and γ , in contrast to the analogous equations for the GARCH model, do not involve derivatives of the conditional variances. This arises because the conditional variances depend on daily data, and there is no link in the MIDAS model between the conditional mean of the low frequency returns and the conditional mean of the higher frequency returns. Consequently, there is no indirect effect of conditional mean errors on the conditional mean parameters through their effect on the conditional variances.

Once we have the first order conditions of the MLE for the MIDAS model, we proceed with the GMM estimation of the ODIN-MIDAS model as in the ODIN-GARCH model, above. We treat the first order conditions of the MLE as GMM orthogonality conditions, estimate

the parameters to satisfy the average of the 22 possible starting values, and construct the standard errors to allow for the serial correlation induced by the use of overlapping data.

4 Power Analysis of the Conditional CAPM

The section analyzes the power of the basic models before examining the power of the ODIN models. We first consider the improvement in power from using t -statistics instead of the distribution of the coefficient estimates in the GARCH model. Then, we examine power properties of the MIDAS model, and then the improvements in power from using the ODIN models.

4.1 Power in the Basic Monthly GARCH Model

We note above that the evidence for a significantly positive risk-return trade-off, γ , is decidedly mixed. Lundblad (2007) demonstrates that the negative estimates of γ often found in the literature are not inconsistent with the conditional CAPM being the true data-generating process. He demonstrates this intuitively by simulating 500 months of returns from a GARCH-M, in which $\gamma = 2$ and the volatility process is calibrated to historical data. He finds that 19% of the estimates of γ fall below zero.

We now ask a different but related question. If $\gamma \neq 0$, how likely are we to reject the null of $\gamma = 0$ at a 5% marginal level of significance in particular sample sizes? This is a question about the power of the test of $\gamma = 0$ for different sample sizes against the alternative hypothesis that $\gamma = 1, 2$, or 3 , for example.¹⁰ To address these power issues, we simulate from the GARCH-M model using the same parameters as Lundblad (2007). Specifically, the

¹⁰This analysis uses a non-overlapping data framework to maintain consistency with the analysis in Lundblad (2007). In Section 7, we show that using the ODIN framework increases the power of the t -test even more.

data-generating process imposes the restrictions of the conditional CAPM such that

$$R_{t_{m+1}} = \gamma\sigma_{t_m}^2 + \varepsilon_{t_{m+1}}$$

where $\varepsilon_{t_{m+1}} \sim N(0, \sigma_{t_m}^2)$, and the monthly conditional variance evolves according to a GARCH model

$$\sigma_{t_{m+1}}^2 = \omega + \alpha\varepsilon_{t_{m+1}}^2 + \beta\sigma_{t_m}^2$$

with $\omega = 0.0002$, $\alpha = 0.1$, $\beta = 0.85$. We simulate under the null hypothesis, $\gamma = 0$, and under the alternative hypotheses, $\gamma = 1, 2$ and 3 , and we vary the sample length from $T = 500$ to $5,000$. For each sample we jointly estimate the parameters of the model, including a constant, μ , with MLE, and we conduct 5,000 simulations of each sample size. Rather than focusing strictly on the distributions of $\hat{\gamma}$, as in Lundblad (2007), we also examine the distributions of the test statistics under both the null and the alternative hypotheses to examine the powers of the tests.

Our simulations indicate that with at least 500 observations, the MLE of γ is only slightly biased, but it is fairly spread out, as in Lundblad (2007). In particular, with $T = 500$, we confirm that 19% of the estimates are less than zero. We also find that the 95th-percentile of $\hat{\gamma}$ is 4.48 and that the 97.5-percentile is 6.03. Hence, with 500 observations, if one were to base a significance test strictly on the distribution of the coefficient estimate, one could argue that to reject the null hypothesis of $\gamma = 0$ at the 5% marginal level of significance one would need to observe $\hat{\gamma} > 6.03$ when using a two-sided test or $\hat{\gamma} > 4.48$ for a one-sided test, which is more relevant because we think that the true γ is positive. Increasing the sample size to 1,000 months, the 95th and 97.5-percentiles of $\hat{\gamma}$ under the null hypothesis decrease to 2.56 and 3.27, respectively. When simulating under the alternative hypothesis that $\gamma = 2$ with 500 months of data, only 9.2% of the samples have a $\hat{\gamma} > 6.03$, while with 1,000 months of data, 21.1% of the samples have a $\hat{\gamma} > 3.27$. Viewed in this way, the power of MLE in

this environment seems quite low.

Appropriately assessing the power of an estimation strategy, though, requires examination of a test statistic. Under the null hypothesis of $\gamma = 0$, the t -statistic is a pivot whose distribution does not depend on the actual underlying parameters and which should be asymptotically $N(0, 1)$. The Online Appendix shows QQ-plots of the t -statistics against the quantiles of a standard normal distribution, and with at least 500 observations we find that convergence of the test statistic is excellent as the QQ-plot virtually overlays the diagonal line.¹¹ Thus, the 95% confidence interval of the t -test is approximately $(-1.96, 1.96)$ for these sample sizes. Under the alternative hypotheses that $\gamma = 1, 2$, or 3 , the t -statistic follows a non-central t -distribution where the non-centrality parameter depends on the value of γ and the sample size. As the sample size increases, the distribution moves farther to the right, and the power of the test increases.

To illustrate the differences between using the distribution of the coefficient estimates and the t -test, we examine the power of both. We first calculate 95% confidence intervals for the point estimates and the t -statistics using the distributions simulated under the null of $\gamma = 0$. Next, we calculate the percentages of the observations that fall outside these confidence intervals when simulating under the different alternative hypotheses. These percentages represent the powers of the two tests.

Figure 1 presents the power functions for three ‘pairs’ of lines associated with the alternative hypotheses that $\gamma = 1, 2$, or 3 . For each pair, the solid line graphs the power based on the t -test, and the dashed line graphs the power based on coefficient estimates as a function of the sample size.

¹¹Rather than present these figures, we have placed them in an Online Appendix. The standard errors are calculated using the quasi-maximum likelihood (QMLE) approach of Bollerslev and Wooldridge (1992). We take care in evaluating the Hessian of the likelihood function. In particular, the Hessian returned from the MatLab optimization routine is not reliable, as it is a so-called ‘pseudo-Hessian’ constructed with the purpose of choosing sensible step-sizes, not to be a high-precision estimate of the second derivatives. Instead, we use the DERIVEST suite by D’Errico (2011), an adaptive numerical differentiation toolbox that provides high-precision first-order and second-order derivatives.

Note that, as expected, the power of the t -test is always higher than the power of the test based on coefficient estimates, and for very large sample sizes, the powers of both tests converge to one. Recall that 500 months is about 42 years, and 1,000 months is about 83 years. For a 500 month sample, using the t -test roughly doubles the power compared to using the point estimates. Comparing the two power curves for $\gamma = 1, 2$, and 3 to see the number of observations for which the lower power test has the same power as the t -test with 500 observations indicates that we would need about 1,200 observations when $\gamma = 1$, 950 observations when $\gamma = 2$, and about 800 observations when $\gamma = 3$. For 1,000 months, the increase in power is a little less, but it is still substantial. Because Lundblad (2007) focuses directly on the distributions of $\hat{\gamma}$, he somewhat overstates the difficulty of rejecting the null hypothesis of $\gamma = 0$. We nevertheless agree with his central point: If the true risk-return trade-off is $\gamma = 2$, with 1,000 months of data, we only have a 30% chance of rejecting the false null (up from 21% if one uses coefficient estimates instead of the t -statistic). If $\gamma = 1$, power drops to 11%, while if $\gamma = 3$, power increases to 57%. Section 4.3 demonstrates that using the ODIN framework increases the power of the tests, and hence the effective sample size, even more.

4.2 Power Analysis in the Basic Monthly MIDAS Model

As for the GARCH model, all MIDAS simulations have $\mu = 0$ in the data generating process, but μ is estimated with the other parameters using MLE. Figure 2 analyzes the power of the test of the null hypothesis, $\gamma = 0$, against different alternatives in the MIDAS model for two sets of parameters.¹² The first set of parameters, associated with the left part of Figure 2, use $\phi = 0.87$ and $\kappa = 21$, which are chosen to match the volatility-of-volatility in the GARCH

¹²The Online Appendix contains QQ-plots of the simulated t -statistics against the quantiles of a standard normal distribution. We first calculate 95% confidence intervals for the t -statistics using the distribution simulated under the null of $\gamma = 0$. Next, we calculate the percentages of the observations that fall outside these confidence intervals when simulating under the alternatives.

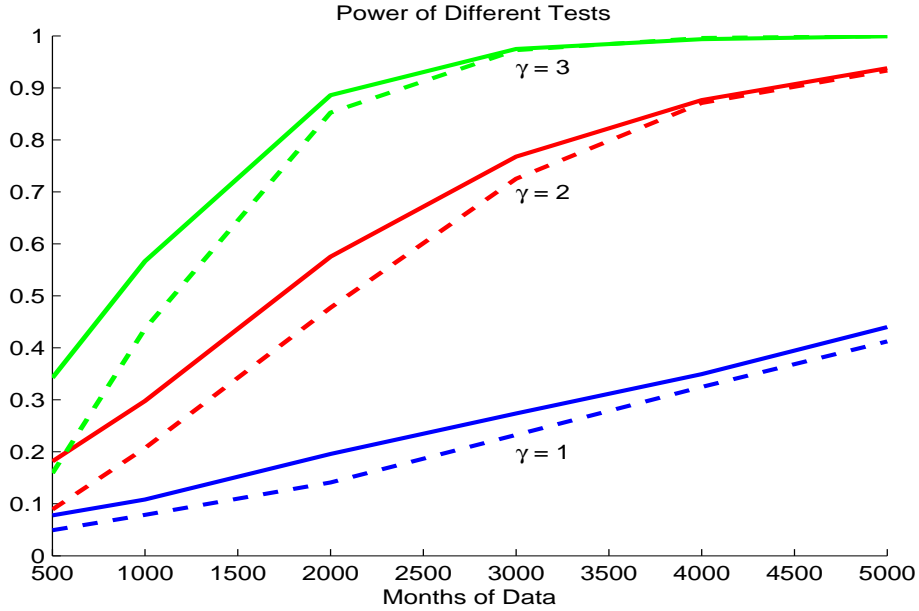


Figure 1: Power of the test of $\gamma = 0$ against different alternatives in the GARCH-M model for a 5% marginal level of significance. The solid lines represent the power of the t -test, and the dotted lines represent the power of the test using the distribution of the coefficient estimates.

model considered in Figure 1. We do this because, as shown above, the volatility-of-volatility is important for the identification of γ . As might be anticipated in this case, the power of the MIDAS model is not much different from that of the GARCH model. The right plot shows the power of the MIDAS model calibrated to the value-weighted market return over 1927-2011. In this case, the volatility-of-volatility is much lower as $\phi = 0.69$ and $\kappa = 5.55$. The smaller values of both ϕ and κ decrease the implied volatility-of-volatility, directly with a lower value of ϕ but also because a lower value of κ makes the beta-polynomial weight function flatter which ‘averages’ the monthly volatility over a longer history of past daily squared returns.

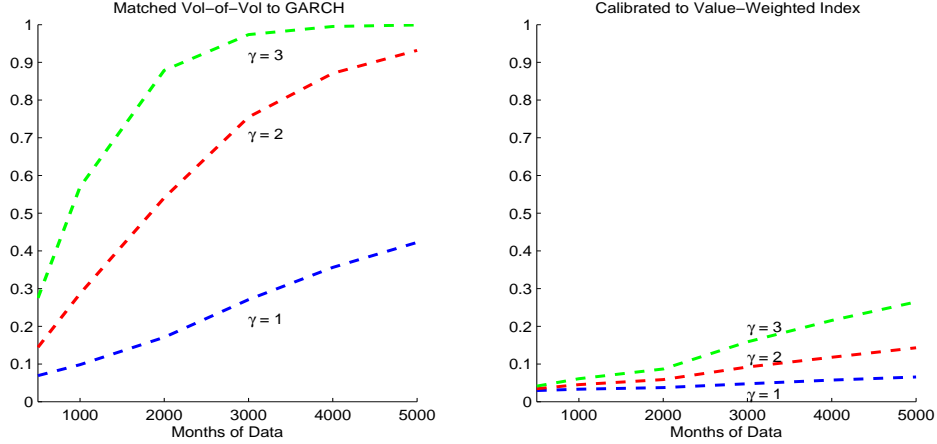


Figure 2: Power of the test $\gamma = 0$ against different alternatives in the MIDAS model for a 5% marginal level of significance. In the left plot the parameters are chosen to match the volatility-of-volatility for the GARCH model in Figure 1. In the right plot, the model is calibrated to the value-weighted index over 1929-2011 (see Table 5.)

4.3 Power Analysis of ODIN Models

The ODIN estimation method shares a lot with the basic estimation method. We think of the underlying model as an abstraction that holds at a given frequency, but we assume that the starting date is irrelevant. As noted above, the ODIN estimator specified at the monthly horizon maximizes the average of the 22 likelihood functions based on the different starting dates. Hence, under the assumption that the starting date is irrelevant, the probability limit of the ODIN estimator is the same as the probability limit of the basic estimator. The asymptotic variance of the ODIN estimates is however always smaller than the asymptotic variance of the basic estimator. As shown above in equation (20), the spectral density matrix that enters the asymptotic variance of the monthly ODIN estimator is $S = \sum_{j=-21}^{21} C_T(j)$, where $C_T(j)$ is the sample autocovariances of g_t with g_{t-j} . In the Online Appendix we show that the corresponding spectral density matrix for the basic monthly model is $S = 22C_T(0)$. Hansen and Hodrick (1980) demonstrate that $S = \sum_{j=-21}^{21} C_T(j)$ is less than $22C_T(0)$ by a positive definite matrix, and it follows that the asymptotic variance of the ODIN estimator

is smaller than the asymptotic variance of the basic estimator.

While the asymptotic results are interesting, to assess the performance of ODIN estimation on historical sample sizes we simulate 22,000 days or 1,000 months of data from a continuous time GARCH model. Following Nelson (1990), Drost and Nijman (1993), Andersen and Bollerslev (1998), and Lundblad (2007) the continuous-time limit for a GARCH(1,1)-M is specified as

$$\begin{aligned}\frac{dP_t}{P_t} &= \gamma\sigma_t^2 dt + \sigma_t dW_{P,t} \\ d\sigma_t^2 &= \theta(\omega - \sigma_t^2)dt + \sqrt{2\lambda\theta} \sigma_t^2 dW_{\sigma,t}\end{aligned}$$

where P_t is the market price level, σ_t^2 is the stochastic instantaneous variance process, and $W_{P,t}$ and $W_{\sigma,t}$ are independent Brownian motions. Andersen and Bollerslev (1998) derive the mapping between discrete time GARCH parameters estimated on monthly data ($\omega = 0.0002, \alpha = 0.10, \beta = 0.85$), and the continuous-time parameters ($\theta = 0.0023, \omega = 1.8182 \cdot 10^{-4}, \lambda = 0.459$, assuming 22 trading days per month), as in Lundblad (2007). We then simulate from the continuous time model in 5-minute increments using a standard Euler scheme and different values of γ . Finally, we sample the process to get daily log prices and compute daily log returns. Summing these gives log returns for any given start date and any forecasting horizon, and these returns satisfy a weak GARCH model. The Online Appendix contains further details on the simulations. We then estimate basic GARCH models and ODIN GARCH models with forecasting horizons of one, five, ten, 22, 33, 44, 55, and 66 days. In the basic model, the sampling frequency is the same as the forecast horizon, whereas the ODIN-GARCH model always uses all available daily data. Although no analytical results are available for QMLE or GMM applied to weak GARCH models, Drost and Nijman (1993) show that the asymptotic bias of the QMLE estimates is small, which we confirm is also the case for the ODIN estimator.

Table 1 presents the means of $\hat{\gamma}$ from the 10,000 simulations, both for the null hypothesis, $\gamma = 0$, and for three alternative hypotheses, $\gamma = 1, 2$, and 3 . We find that $\hat{\gamma}$ is slightly biased, which confirms the results in Drost and Nijman (1993) that QMLE is not consistent but that the asymptotic bias is small. Under the null hypothesis, there is only a slight bias in $\hat{\gamma}$ for the ODIN estimate, as the largest mean is 0.043. The bias for the ODIN estimator is comparable to the bias of the basic estimator for small forecast horizons, but the bias fluctuates more for the basic model for long forecasting horizons.

As we increase the value of γ , the bias with daily sampling grows from -2% at $\gamma = 1$, to -3.5% at $\gamma = 2$, and -4.3% at $\gamma = 3$. For the basic GARCH, this negative bias is offset by a positive bias induced by the shrinking of the sample size as the forecast horizon increases. In some cases, these biases offset such that there appears to be little overall bias for sampling frequency 22 and $\gamma = 1$, for sampling frequency 33 and $\gamma = 2$, and for a sampling frequency between 33 and 44 for $\gamma = 3$. For the ODIN GARCH specification, as the overlap in forecasts increases, the negative bias is still present for $\gamma = 2$ and 3 . As the forecasting interval is increased, the bias becomes less negative for $\gamma = 1$, but more negative when $\gamma = 2$ or 3 . In general, the bias is smaller for the ODIN model than for the basic model for long forecasting horizons.

The Online Appendix shows QQ-plots of $\hat{\gamma}$ against the quantiles of a normal distribution. For all sampling frequencies and all values of γ , the distributions of $\hat{\gamma}$ for the ODIN model are closer to a normal distribution than is the distributions of $\hat{\gamma}$ for the basic model. The Online Appendix also shows the empirical means of $\hat{\gamma}$ when the models are estimated on 5,000 months of simulated data. In this case, the biases of the basic model and the ODIN model are very similar. Further, the empirical means of $\hat{\gamma}$ based on 1,000 months of data are much closer to the large-sample means for the ODIN model, suggesting that the small-sample bias is larger for the basic model than for the ODIN model.

While these simulations from a weak GARCH model indicate small biases in the pa-

parameter estimates, we are interested in the increased precision that ODIN offers. Indeed, correctly assessing any bias that basic and ODIN estimations may introduce when calibrated on real data requires taking explicit stand on the true data generating process, and the weak GARCH model is but one example. To illustrate the increased precision that ODIN offers, Figure 3 shows box-plots of $\hat{\gamma}$ for $\gamma = 0$.¹³ The Online Appendix show similar box-plots for $\gamma = 1, 2$, and 3. For each sampling frequency, the left blue box-plot shows the results for the basic GARCH, and the right green box-plot shows the results for the ODIN GARCH. In the right figure, showing sampling frequencies 33-66, the values have been truncated at 50. Beneath each box-plot is the number of estimations out of 10,000 that did not converge (meaning that the estimates of the GARCH process had $\hat{\alpha} = 0, \hat{\beta} = 0$, or $\hat{\alpha} + \hat{\beta} = 1$). For instance, for $\gamma = 0$ and a 33 day sampling frequency, 88 of the basic estimations did not converge, and 5 of the ODIN estimations did not converge. Note that the two models are estimated on exactly the same simulated data. For all sampling frequencies, the ODIN model converges more often than the basic model. It is clear from the figure that the precision of the ODIN GARCH estimates is much higher than the precision of the basic GARCH estimates for all sampling frequencies. Further, the increase in precision increases with the sampling frequency.

Having illustrated the increased precision for the ODIN estimates, we now turn to hypothesis testing. Hypothesis testing is based on the estimated standard errors, and we first compare the estimated standard errors to the empirical standard deviations of $\hat{\gamma}$ in Table 9. The first two lines of the table show the empirical standard deviations of the 10,000 $\hat{\gamma}$'s for the basic and ODIN models, respectively. The next two lines show the averages of the 10,000 estimated standard errors. The last two rows show the bias in the estimated standard errors as the percentage deviation of the average estimated standard error from the actual

¹³The central mark is the median, the edges of the box are the 25th and 75th percentiles, the whisker length is 1.5 of the interquartile range which for a normal distribution corresponds to 99.3% coverage. Outliers beyond the whiskers are marked individually with a '+'.

Table 1: Empirical Means.

Forecast Horizon	1	5	10	22	33	44	55	66
$\gamma = 0$								
Basic GARCH	0.0273	0.0353	0.0371	0.0215	0.0042	0.0739	0.1118	-0.0222
ODIN GARCH	0.0273	0.0273	0.0258	0.0228	0.0229	0.0360	0.0393	0.0434
$\gamma = 1$								
Basic GARCH	0.9808	0.9777	0.9791	0.9889	1.0621	1.1657	1.2817	1.5190
ODIN GARCH	0.9808	0.9752	0.9686	0.9702	0.9795	1.0003	1.0149	1.0031
$\gamma = 2$								
Basic GARCH	1.9305	1.9139	1.9100	1.9302	2.0185	2.1966	2.4449	2.8619
ODIN GARCH	1.9305	1.9189	1.9023	1.8943	1.8997	1.9182	1.9195	1.8949
$\gamma = 3$								
Basic GARCH	2.8722	2.8379	2.8194	2.8235	2.9354	3.1225	3.4697	3.7735
ODIN GARCH	2.8722	2.8542	2.8185	2.7750	2.7535	2.7412	2.7100	2.6651

Note: The Table shows the empirical means of $\hat{\gamma}$ for different values of γ and for different forecast horizons, which correspond to different sampling frequencies for the Basic GARCH model.




Figure 3: Box-plots of $\hat{\gamma}$ when $\gamma = 0$, for different sampling frequencies. For each sampling frequency, the left box-plots shows the results for the basic GARCH, and the right box-plot shows the results for the ODIN GARCH. In the right, the values have been truncated at 50. Beneath each box-plot is the number of estimations that did not converge (out of 10,000).

Table 2: Precision of Basic and ODIN Model Estimates.

Sampling Frequency	1	5	10	22	33	44	55	66
Empirical Std of $\hat{\gamma}$, Basic	1.05	1.19	1.33	1.71	2.98	4.34	9.78	10.66
Empirical Std of $\hat{\gamma}$, ODIN	1.05	1.15	1.26	1.49	1.72	1.94	2.13	2.25
Average of $\hat{\sigma}(\hat{\gamma})$, Basic	1.02	1.15	1.28	1.61	2.28	3.86	7.70	9.40
Average of $\hat{\sigma}(\hat{\gamma})$, ODIN	1.02	1.11	1.20	1.41	1.60	1.78	1.93	2.06
Basic Bias	-0.03	-0.04	-0.04	-0.05	-0.24	-0.11	-0.21	-0.12
ODIN Bias	-0.03	-0.04	-0.04	-0.06	-0.07	-0.08	-0.09	-0.09

Note: Table for $\gamma = 0$. The first two lines of the table show the empirical standard deviation of $\hat{\gamma}_i$ for the basic and ODIN models, based on all 10,000 estimates. The next two lines show the average of the 10,000 estimated standard errors. The last two rows show the bias in the estimated standard errors as the percentage deviation of the average estimated standard error from the actual standard deviation.

standard deviation. As we saw above, the standard deviation of the 10,000 estimates of γ is much smaller for the ODIN model which we quantify in rows 1 and 2. Rows 3 and 4 also demonstrate that the average estimated standard errors, $\hat{\sigma}(\hat{\gamma})$, are smaller for the ODIN model. Finally, rows 5 and 6 show that the asymptotic standard error estimates are generally too small, but the bias in the basic model is far larger. This analysis shows that the standard error estimates for the ODIN model are generally more appropriate than the standard errors for the basic model for the available historical sample sizes.

Finally, we assess the increase in the power of the tests from using ODIN versus sampling the data. The QQ-plots of the empirical distributions of the t -statistics versus a standard normal distribution in the Online Appendix indicate that with these sample sizes, the asymptotic distributions of the t -statistics are very close to a standard normal, but that the convergence is better for the ODIN model. We nevertheless use the empirical distribution under the null hypothesis for inference.

Figure 4 displays the power of the test of the null hypothesis $\gamma = 0$ against the alternative hypotheses $\gamma = 1, 2$, and 3 for different sampling horizons, using a two-sided test. The actual values are calculated for forecast horizons of 1, 5, 10, 22, 33, 44, 55, and 66 days.¹⁴ The

¹⁴The Online Appendix contains QQ-plots of the simulated t -statistics against the quantiles of a standard

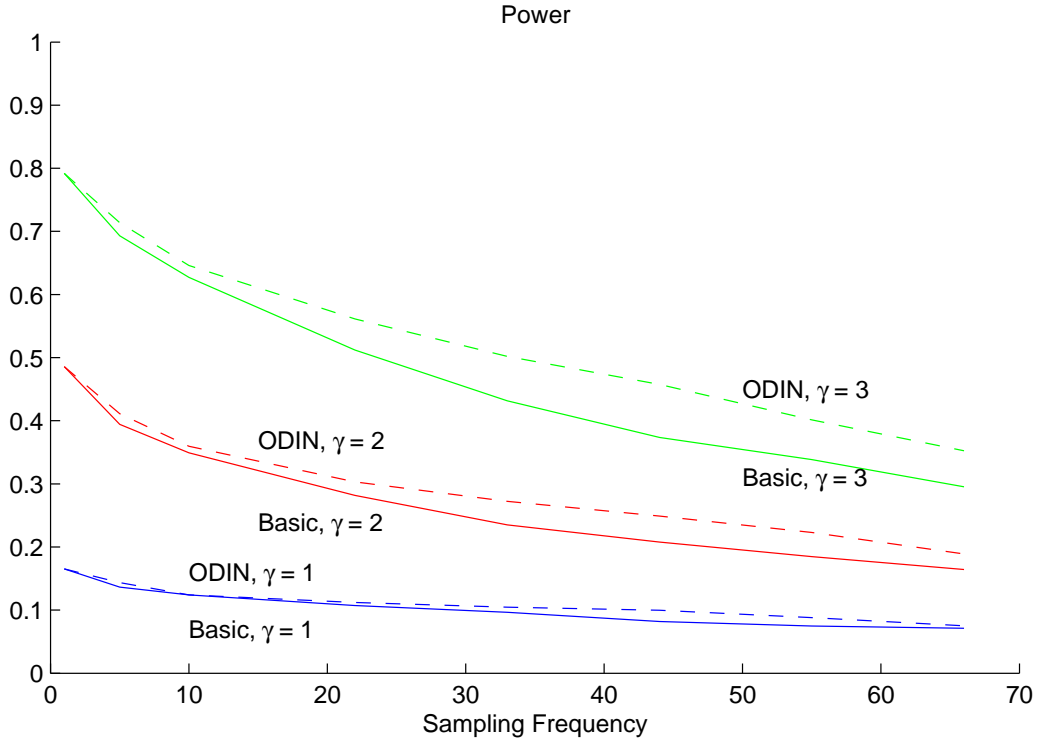


Figure 4: Power of ODIN GARCH vs. Basic GARCH for various values of the risk-return trade-off and a sample size of 1,000 months of daily data.

power of the ODIN model is always higher. As seen in Table 3, for monthly sampling, the power of the ODIN model is between 5% and 10% higher than the power of the basic model, and for quarterly sampling the power of the ODIN model is 15% higher for $\gamma = 2$ (19% higher for $\gamma = 3$, and 5% higher for $\gamma = 1$). These results appear less dramatic than the box-plots in Figure 3, which is due to the increased bias of both $\hat{\gamma}$ and $\hat{\sigma}(\hat{\gamma})$ for the basic model, when $\gamma = 1, 2$ and 3.

Also note that because there is a true risk-return relationship at short horizons in the simulations, daily sampling does provide a substantive increase in power. This result is

normal distribution. We first calculate 95% confidence intervals for the t -statistics using the distribution simulated under the null of $\gamma = 0$. Next, we calculate the percentages of the observations that fall outside these confidence intervals when simulating under the alternatives. The t -statistics for both the basic GARCH model and the ODIN GARCH model are slightly skewed, but in opposite directions. Using the traditional $(-1.96, 1.96)$ interval increases the improvement of the ODIN model over the basic model.

Table 3: Power Improvement for ODIN GARCH Models over Standard GARCH Models

Forecast Horizon	1	5	10	22	33	44	55	66
$\gamma = 1$	1.000	1.054	1.001	1.048	1.085	1.221	1.178	1.053
$\gamma = 2$	1.000	1.044	1.030	1.076	1.158	1.199	1.207	1.149
$\gamma = 3$	1.000	1.030	1.030	1.097	1.162	1.225	1.186	1.194

Note: The table presents the improvement in power for the ODIN over the basic model.

contrary to Lundblad’s (2007) conclusion that the sampling frequency does not matter much. The different conclusions are due to Lundblad focusing on point estimates only, whereas we focus on t -statistics.

4.4 A Caveat on the Sample Size

The previous discussion could leave the reader with the impression that ODIN is a free lunch. One can increase power and not suffer any ill consequences. But, we use only relatively long samples in which the overlap remains a small fraction of the sample size. We know from Richardson and Stock (1989) and Valkanov (2003) that building up highly serially correlated error processes can cause the standard asymptotic distribution theory underlying test statistics to provide poor approximations if the sample size is not sufficiently large. We have not worked out the asymptotic distribution theory associated with the functional central limit theorem discussed in these papers, but this is a worthwhile idea. With this caveat in mind, we turn to the evidence from actual data.

5 Estimating the Basic Models

This section presents estimates for the basic GARCH and MIDAS models for three samples of monthly and quarterly non-overlapping calendar data. Table 4 presents the GARCH estimation, and Table 5 presents the modified MIDAS estimation as described in equations

(9) and (10). As above, we estimate ω by variance targeting.

We use the same observations on the dependent variable, R_{t+1} , for the GARCH and MIDAS models. Although CRSP data begin in 1926, the full sample for monthly data is 1927:10 to 2011:12 because of the lags introduced in the MIDAS model. The corresponding full sample for quarterly data is 1927:4 to 2011:4. We split the sample after 1952 to recognize that the Great Depression, World War II, and the lack of Federal Reserve independence prior to the Treasury-Fed Accord of 1952 may have produced data that require more complex modeling than the conditional CAPM. The second monthly sample is consequently 1927:10 to 1952:12, and its quarterly counterpart is 1927:4 to 1952:4. To avoid use of data from before 1952, the third monthly sample is 1955:1 to 2011:12, and its quarterly counterpart is 1955:1 to 2011:4.

The tables report the estimates of γ , with the QMLE asymptotic standard errors in parenthesis, and the associated p -values of the t -tests of $\gamma = 0$. We also report bootstrap p -values under the null hypothesis that the true coefficient is 0 (see Appendix B for details).

In Table 4, for the monthly GARCH model, the $\hat{\gamma}$'s range from 0.44 to 3.01, but only the latter estimate from the post 1952 sample has a bootstrap p -value less than .10. With quarterly data, the GARCH model $\hat{\gamma}$'s range from 1.75 to 6.34, and each of the values is significantly different from zero at the 10% marginal level of significance from the bootstrapped p -values. Note, though, that for the early quarterly sample from 1927-1952, $\hat{\alpha}$ is unusually high, and $\hat{\beta}$ is unusually low.

For the MIDAS results in Table 5, the $\hat{\gamma}$'s do not approach statistical significance at traditional levels. The $\hat{\gamma}$'s for the monthly full sample and the first sub-sample are negative, and the $\hat{\gamma}$ for the second sub-sample has a bootstrap p -value that is .904. With quarterly data, only the $\hat{\gamma}$ of 3.56 in the second sub-sample approaches traditional statistical significance with a bootstrap p -value of .17.

To understand how the MIDAS model weights past squared daily returns in the con-

Table 4: GARCH Estimation Results and Bootstrapped p -values

Panel A: Monthly GARCH(1,1)-M model							
Period	μ	γ	$\omega \times 10,000$	α	β	Obs	LLF
1927:10–2011:12	0.005	1.331	0.764	0.129	0.846	1011	1653.61
Standard error	(0.002)	(0.917)	(0.260)	(0.021)	(0.026)		
p -value	.029	.146	.003	.000	.000		
Bootstrap p		.176					
1927:10–1952:12	0.010	0.443	0.707	0.141	0.846	303	419.91
Standard error	(0.004)	(1.017)	(0.468)	(0.034)	(0.039)		
p -value	.012	.663	.131	.000	.000		
Bootstrap p		.725					
1955:1–2011:12	0.001	3.011	1.028	0.105	0.843	684	1188.81
Standard error	(0.004)	(1.903)	(0.404)	(0.028)	(0.037)		
p -value	.801	.114	.011	.000	.000		
Bootstrap p		.083					
Panel B: Quarterly GARCH(1,1)-M model							
Period	μ	γ	$\omega \times 10,000$	α	β	Obs	LLF
1927:12–2011:12	0.006	1.747	16.930	0.334	0.536	337	321.24
Standard error	(0.011)	(1.008)	(8.678)	(0.085)	(0.137)		
p -value	.570	.083	.051	.000	.000		
Bootstrap p		.090					
1927:12–1952:12	−0.012	1.884	58.054	0.690	0.094	101	69.98
Standard error	(0.020)	(0.401)	(19.197)	(0.072)	(0.072)		
p -value	.554	.000	.002	.000	.189		
Bootstrap p		.012					
1955:3–2011:12	−0.025	6.337	17.208	0.178	0.588	228	246.31
Standard error	(0.043)	(5.320)	(7.183)	(0.104)	(0.064)		
p -value	.565	.234	.017	.086	.000		
Bootstrap p		.098					

Note: The table shows estimation results for the Basic GARCH model, using non-overlapping monthly and quarterly returns. The bootstrapped p -values are based on 5,000 simulations with $\gamma = 0$, keeping the remaining parameters at their estimated values.

Table 5: MIDAS Estimation Results and Bootstrapped p -values

Panel A: Monthly Alternative MIDAS model with β -weights							
Period	μ	γ	$\omega \times 10,000$	ϕ	κ	Obs	LLF
1927:10–2011:12	0.006	−0.149	0.351	0.693	5.550	1011	1655.76
Standard error	(0.003)	(1.318)	(0.057)	(0.050)	(3.433)		
p -value	.015	.910	.000	.000	.106		
Bootstrap p		.789					
1927:10–1952:12	0.013	−1.544	0.304	0.811	2.929	303	420.14
Standard error	(0.004)	(1.693)	(0.090)	(0.056)	(1.137)		
p -value	.004	.362	.001	.000	.010		
Bootstrap p		.285					
1955:1–2011:12	0.005	0.209	0.396	0.557	43.312	684	1204.40
Standard error	(0.003)	(1.694)	(0.062)	(0.070)	(17.418)		
p -value	.102	.902	.000	.000	.013		
Bootstrap p		.904					
Panel B: Quarterly Alternative MIDAS model with β -weights							
Period	μ	γ	$\omega \times 10,000$	ϕ	κ	Obs	LLF
1927:12–2011:12	0.015	0.620	0.401	0.650	6.427	337	320.18
Standard error	(0.012)	(2.044)	(0.115)	(0.100)	(7.859)		
p -value	.222	.762	.000	.000	.414		
Bootstrap p		.774					
1927:12–1952:12	0.036	−1.111	0.165	0.897	3.365	101	70.41
Standard error	(0.013)	(2.017)	(0.174)	(0.108)	(1.312)		
p -value	.007	.582	.344	.000	.010		
Bootstrap p		.390					
1955:3–2011:12	−0.004	3.564	0.488	0.454	9.431	228	244.96
Standard error	(0.017)	(3.016)	(0.134)	(0.150)	(4.750)		
p -value	.816	.237	.000	.003	.047		
Bootstrap p		.172					

Note: The table presents estimation results for the alternative MIDAS specification of the conditional variance in equation (9) with beta-polynomial weights as in equation (10). The estimates of μ have been multiplied by 22 (66) for the monthly (quarterly) results to make them comparable to the GARCH estimates. The bootstrapped p -values are based on 5,000 simulations with $\gamma = 0$, keeping the remaining parameters as their estimated values.

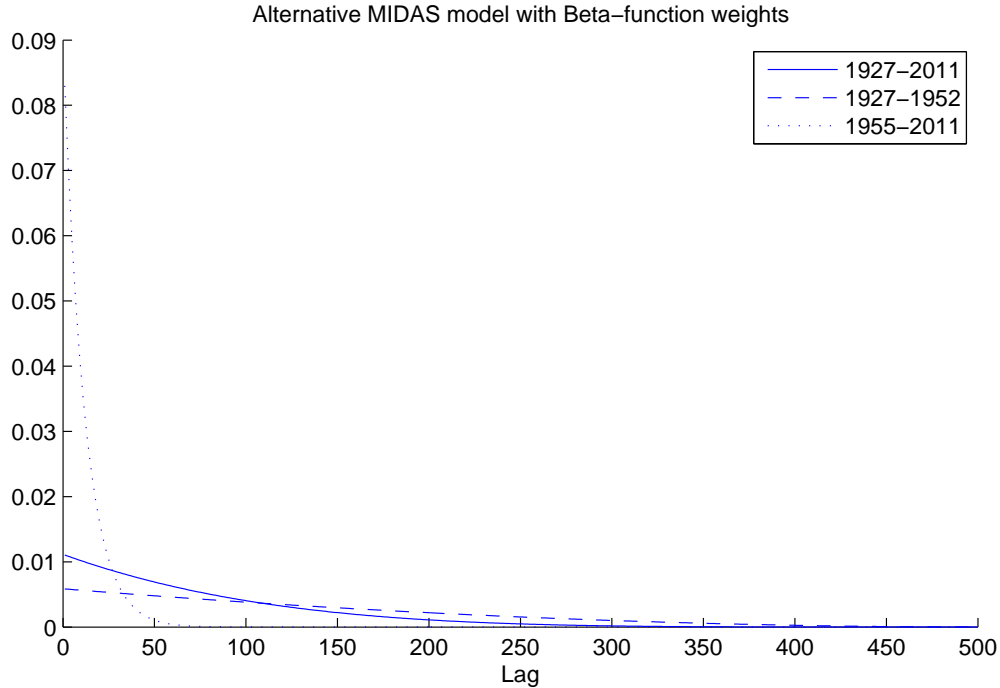


Figure 5: The weight functions in the MIDAS models.

ditional variance, Figure 5 shows the estimated weight functions. The importance of the recent past is much higher in the second sub-sample than in the whole sample or the first sub-sample. While the weight functions are all declining, this is simply due to the fact that we use the beta-polynomial with one parameter, which ensures monotonicity. Nothing in the data inherently makes the weight functions decline. We impose it as a modeling assumption. Using the exponential weight function in equation (8) produces estimates of κ_2 that are positive and induces U-shaped weight functions.

Figure 6 presents the conditional standard deviations for the GARCH and MIDAS models, in annualized values. For the full sample period, 1927-2011, the two estimates are very similar although the volatility-of-volatility is smaller for the MIDAS model. For the first sub-period 1927-1952, the conditional standard deviation in the MIDAS model is smoother and reacts less quickly to new innovations, due to the flatter weight function as seen in Figure 5. On the other hand, for the second sub-period 1955-2011, the conditional standard

deviation from the MIDAS model is much more ‘choppy’ than that of the GARCH model, which is seen most clearly for the 2004-2010 period. This is due to the steep weight function shown in Figure 5.

6 Estimating the ODIN Models

The results of the ODIN-GARCH estimation are presented in Table 6 with the results for monthly and quarterly data in Panels A and B, respectively. Table 7 presents the comparable ODIN-MIDAS results. Each panel contains results for the three sample periods examined above: the full sample, 1927:10-2011:12; the first sub-sample, 1927:10-1952:12; and the second sub-sample, 1955:1-2011:12. Standard errors are presented in parenthesis below the point estimates with p -values below the standard errors.

In Table 6, the $\hat{\gamma}$'s for the monthly samples are similar to the basic GARCH model. The $\hat{\gamma}$ for the full sample ODIN model is 1.678 with a p -value of .210 compared to a $\hat{\gamma}$ of 1.331 with a p -value of .146 for the basic model. The $\hat{\gamma}$ for the first sub-sample ODIN model is 0.654 with a p -value of .684 compared to a $\hat{\gamma}$ of 0.443 with a p -value of .663 for the basic model. The $\hat{\gamma}$ for the second sub-sample ODIN model is now 3.354 with a p -value of .022 compared to 3.011 and .114 for the basic model.

For the quarterly results, the $\hat{\gamma}$ in the full sample decreases from 1.747 to 1.572, and the p -value increases from .083 to .097. The p -value for $\hat{\gamma}$ in the first sub-sample increases substantially from .012 to .666, while the estimates of α and β now have the more traditional values expected in a GARCH (1,1) model. This suggests that the calendar results from the first sub-sample are spuriously significant. The $\hat{\gamma}$ for the second sub-sample ODIN model is now 4.894 with a p -value of .213 compared to 6.337 and .234 for the basic model.

The results for the ODIN-MIDAS model show some improvements in the standard errors for the quarterly data relative to the results in Table 5, but the standard errors of the




Figure 6: Conditional standard deviations in the MIDAS and GARCH models, for the three different sample periods.

Table 6: Estimation Results for ODIN GARCH Model.

Panel A: Monthly ODIN-GARCH						
Period	μ	γ	α	β	Obs	LLF
1927:9:30–2011:12:8	0.004	1.678	0.122	0.841	22220	35509.63
Standard error	(0.003)	(1.340)	(0.024)	(0.033)		
<i>p</i> -value	.159	.210	.000	.000		
1927:9:30–1952:12:12	0.008	0.654	0.109	0.869	7370	10374.01
Standard error	(0.005)	(1.608)	(0.034)	(0.044)		
<i>p</i> -value	.136	.684	.001	.000		
1955:1:26–2011:12:13	0.001	3.354	0.133	0.783	14300	24114.36
Standard error	(0.004)	(1.464)	(0.026)	(0.037)		
<i>p</i> -value	.853	.022	.000	.000		
Panel B: Quarterly ODIN-GARCH						
Period	μ	γ	α	β	Obs	LLF
1927:9:30–2011:12:8	0.011	1.572	0.220	0.655	22176	21532.11
Standard error	(0.009)	(0.947)	(0.062)	(0.099)		
<i>p</i> -value	.244	.097	.000	.000		
1927:9:30–1952:12:12	0.022	0.439	0.133	0.856	7326	5540.65
Standard error	(0.014)	(1.017)	(0.066)	(0.092)		
<i>p</i> -value	.111	.666	.043	.000		
1955:1:26–2011:12:13	−0.013	4.894	0.170	0.500	14256	15508.53
Standard error	(0.031)	(3.931)	(0.101)	(0.098)		
<i>p</i> -value	.687	.213	.092	.000		

Note: The table presents results from estimation of the GMM estimation of the ODIN-GARCH model specified in the orthogonality conditions of equation (19). Standard errors are in parenthesis with *p*-values below.

Table 7: Estimation Results for ODIN MIDAS Model.

Panel A: Monthly ODIN-MIDAS						
Period	μ	γ	ϕ	κ	Obs	LLF
1927:9:30–2011:12:8	0.005	0.483	0.656	21.520	22220	35721.81
Standard error	(0.002)	(1.230)	(0.058)	(13.577)		
<i>p</i> -value	.048	.695	.000	.113		
1927:9:30–1952:12:12	0.012	−1.359	0.804	3.319	7370	10562.47
Standard error	(0.005)	(1.741)	(0.065)	(1.388)		
<i>p</i> -value	.010	.435	.000	.017		
1955:1:26–2011:12:13	0.004	0.868	0.534	64.657	14300	24311.54
Standard error	(0.003)	(1.708)	(0.064)	(27.948)		
<i>p</i> -value	.219	.611	.000	.021		
Panel A: Quarterly ODIN-MIDAS						
Period	μ	γ	ϕ	κ	Obs	LLF
1927:9:30–2011:12:8	0.019	−0.011	0.620	5.035	22176	21917.19
Standard error	(0.009)	(1.502)	(0.083)	(1.383)		
<i>p</i> -value	.046	.994	.000	.000		
1927:9:30–1952:12:12	0.036	−1.331	0.776	3.571	7326	5811.72
Standard error	(0.015)	(1.812)	(0.115)	(1.264)		
<i>p</i> -value	0.016	0.463	0.000	0.005		
1955:1:26–2011:12:13	0.001	2.619	0.454	7.762	14256	15587.65
Standard error	(0.015)	(2.710)	(0.152)	(2.411)		
<i>p</i> -value	.954	.334	.003	.001		

Note: The table presents results from the GMM estimation of the ODIN-MIDAS model. The values of μ in the MIDAS model have been multiplied by 22 (66) to measure them in monthly (quarterly) values, making them comparable to the estimates in the GARCH model. Standard errors are in parenthesis with *p*-values below.

monthly models are actually larger for the two sub-samples. None of the $\hat{\gamma}$'s has a p -value smaller than .334. The first sub-sample ODIN model $\hat{\gamma}$'s for both the monthly and quarterly samples are negative. For monthly data, the full sample $\hat{\gamma}$ is now positive, but for quarterly data, the full sample $\hat{\gamma}$ is now negative. The estimate of γ that is the most comparable to the ODIN-GARCH model is the second sub-sample in which $\hat{\gamma} = 2.619$ with a p -value of .334.

We next examine the differences in the individual estimates of the various non-overlapping samples to determine whether the reason we do not see uniform improvement in the standard errors of the ODIN models relative to their basic counterparts can be traced to instability in the parameters.¹⁵

6.1 Individual Estimations vs. ODIN

Realizing that the starting date of the sample does not matter raises the question of how different can the various non-overlapping estimates be. The surprising answer is, quite different. Figures 7, 8, and 9 illustrate the differences between the 22 possible basic, non-overlapping, monthly estimations and the ODIN estimations for the three sample periods: the full sample, labeled A; the first sub-sample, labeled B; and the second sub-sample, labeled C.

Figure 7 shows the results for the full sample, 1927–2011. The top four plots present the estimates of the four parameters in the GARCH model, and the bottom four plots present the results for the MIDAS model. The solid black line shows the 22 point estimates from estimating the basic GARCH model on 22 day returns, shifting the sampling start date one day at a time. The shaded gray area shows the 95% confidence intervals for the estimates.

¹⁵It is also possible to do ODIN estimation with alternative conditional density functions, such as Student's t -distribution. We did not find that allowing for this alternative specification improved the performance of the conditional CAPM. These results are available in the Online Appendix, which also presents results for models in which the conditional variance responds asymmetrically to the innovation in returns. When we allow for this asymmetry, we find that the risk-return trade-off is imprecisely estimated.

These 22 estimates are obviously not independent as they are based on the same data, which are just sampled differently, and as expected the estimate moves slowly with the sampling start date. Next, the ‘Basic’ estimates represent the parameter estimates from the basic models estimated from monthly calendar returns, with horizontal lines showing the 95% confidence intervals. Finally, the ODIN estimate is shown together with its 95% confidence interval.

The top right plot presents $\hat{\gamma}$ in the GARCH model, and the (3,2) plot presents $\hat{\gamma}$ in the MIDAS model. One clearly sees the negative correlation between $\hat{\mu}$ and $\hat{\gamma}$ for both the GARCH and MIDAS models, as the starting date varies. The individual $\hat{\gamma}$ ’s vary from 1 to 3 for the GARCH model, and between -0.1 and 0.8 for the MIDAS model. For the GARCH model, one sees a negative correlation between $\hat{\alpha}$ and $\hat{\beta}$, and for the MIDAS model, one sees a negative correlation between $\hat{\phi}$ and $\hat{\kappa}$ because increasing either of these parameters gives rise to higher volatility-of-volatility. None of the 22 individual $\hat{\gamma}$ ’s are significantly different from zero, for either the GARCH or the MIDAS model. For the GARCH model, the ODIN standard error of $\hat{\gamma}$ is actually larger than that of the basic monthly GARCH model. This situation arises because of the variation in $\hat{\gamma}$ that comes from varying the starting date. For some starting dates, the basic non-overlapping standard error is much larger than for other starting dates. While the basic monthly model does not ‘see’ this variation, the ODIN model recognizes the variation that comes from changing the starting date resulting in a larger standard error. Also, note that the point estimate from the ODIN model is closer to the average value of the individual estimates than is the point estimate for the basic monthly model.

Figure 8 presents the results for the first sub-sample, 1927–1952. Again, none of the individual $\hat{\gamma}$ ’s are significantly different from zero, for either of the models. The $\hat{\gamma}$ ’s vary between -0.1 and 2.4 for the GARCH model, and between -1.5 and -0.8 for the MIDAS model. As for the full sample, the standard error of $\hat{\gamma}$ is actually larger in the ODIN

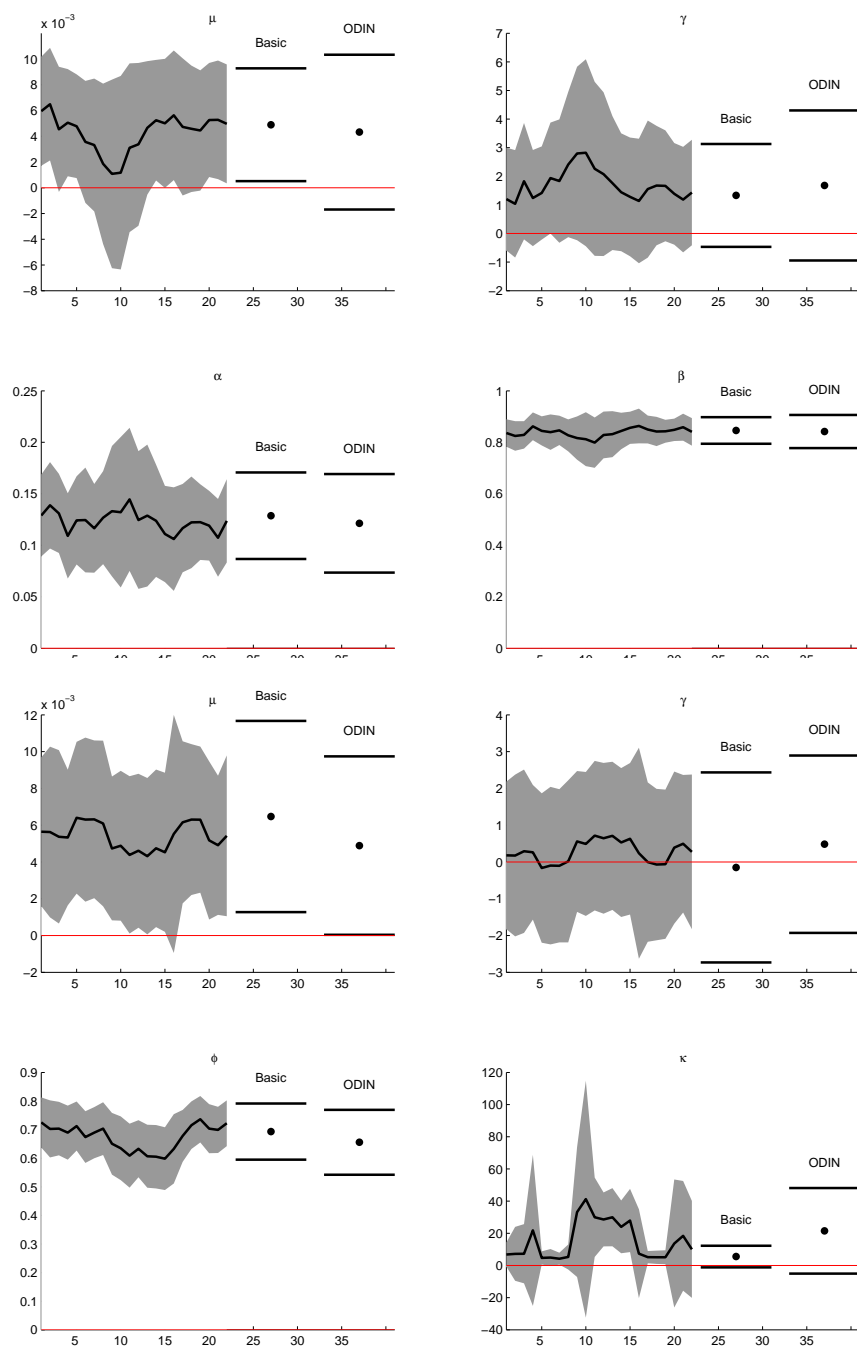


Figure 7: Monthly Estimates, Sample A: 1927–2011. The plots show the 22 individual estimates obtained by shifting the start-date and their 95% confidence interval in shaded grey. The top four plots show the GARCH estimates, the bottom four plots show the MIDAS estimates. The ‘basic monthly’ estimate is based on calendar months and the 95% confidence interval is indicated with horizontal lines. Finally, the ODIN estimate is shown along with its 95% confidence interval.

GARCH model than for the basic GARCH model, again due to the variation in the $\hat{\gamma}$'s and their standard errors as the starting date changes.

Figure 9 presents results for the second sub-sample, 1955–2011. For the GARCH model, most of the individual $\hat{\gamma}$'s are significantly different from zero, even though the estimate from the basic model estimated on calendar month returns happens not to be significantly different from zero. Nevertheless, the ODIN estimation recognizes that it is highly unlikely that the true γ is zero because so many of the individual estimates are significantly different from zero. For the MIDAS model, none of the individual $\hat{\gamma}$'s are significantly different from zero.

Note that there are some similarities between the GARCH and MIDAS estimates. The $\hat{\gamma}$'s are lowest for sample B, slightly higher for sample A, and highest for sample C. The Online Appendix shows similar plots for quarterly estimations. As might be anticipated from the variation in the monthly samples, the variation in the quarterly estimates as the sample starting date changes is much larger for the quarterly estimations. For instance, for the second sub-sample, the GARCH $\hat{\gamma}$'s vary between 2 and 20, and the MIDAS $\hat{\gamma}$'s vary between 0 and 14. These plots also indicate that the basic models fail to converge for some of the starting dates.

6.2 The Average of the Individual Estimates

ODIN constrains the various non-overlapping estimates to have the same value. One could also consider the estimator that is the average of the individual estimates from the correlated non-overlapping samples. In the Online Appendix, we demonstrate that the standard error of this estimator is the same asymptotically as the standard error for the ODIN estimator. For monthly data, the averages of the $\hat{\gamma}$'s for samples A, B, and C, with the ODIN standard errors in parenthesis for the basic GARCH model are 1.699 (1.340), 0.498 (1.608), and 3.302 (1.464), respectively. For the basic MIDAS model, the averages of the $\hat{\gamma}$'s for samples A, B,

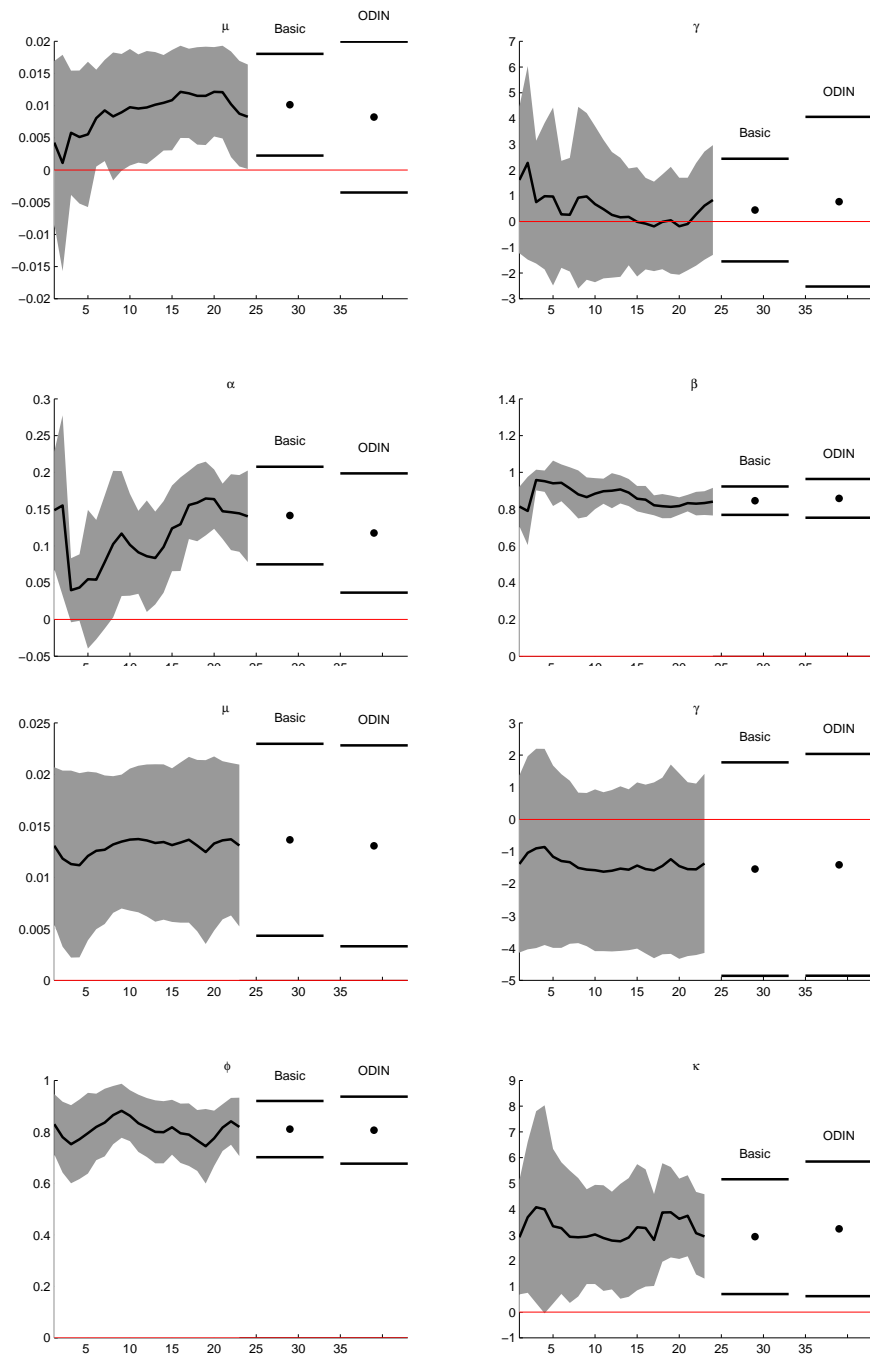


Figure 8: Monthly Estimates, Sample B: 1927–1952. The plots show the 22 individual estimates obtained by shifting the start-date and their 95% confidence interval in shaded grey. The top four plots show the GARCH estimates, the bottom four plots show the MIDAS estimates. The ‘basic monthly’ estimate is based on calendar months and the 95% confidence interval is indicated with horizontal lines. Finally, the ODIN estimate is shown along with its 95% confidence interval.




Figure 9: Monthly Estimates, Sample C: 1955–2011. The plots show the 22 individual estimates obtained by shifting the start-date and their 95% confidence interval in shaded grey. The ‘basic monthly’ estimate is based on calendar months and the 95% confidence interval is indicated with horizontal lines. Finally, the ODIN estimate is shown along with its 95% confidence interval.

and C, with the ODIN standard errors in parenthesis, are 0.277 (1.230), -1.396 (1.741), and 0.768 (1.708), respectively.

7 Conclusions

When financial economists empirically investigate the predictions of their models, they must choose the horizon over which the agents in the model act. For example, Merton's (1973) ICAPM is a theoretical continuous time model, but empirical researchers usually choose a one-month or one-quarter horizon as the most appropriate test environment even though daily data are available. The most popular methods for modeling the conditional variances and covariances that are the sources of risk in these models are GARCH and MIDAS, which are usually implemented with MLE by sampling the data at the same frequency as the horizon chosen for the model. Here we demonstrate that when the data are sampled more finely than the horizon of the model, we can use all of the available data to lower the standard errors of the estimates and improve the power of the tests of the theories by using overlapping data inference (ODIN). Our insight is to use the first order conditions of MLE as orthogonality conditions of GMM. We estimate the parameters of the model from the average of the overlapping MLE samples and construct appropriate standard errors that account for the serial correlation induced by the overlapping data.

We apply this ODIN methodology to investigate the risk-return trade-off implied by the conditional CAPM using GARCH and MIDAS modeling of the conditional variance of the market return. Simulations of the ODIN methodology indicate that decreases in standard errors and increases in power can be substantial and correspond to large increases in the sample size. For example, with a 1,000 month sample of daily data, a horizon of 22 days, and $\gamma = 3$; the average standard error for the ODIN-GARCH model is 84.8% of its basic counterpart. Because standard errors decrease linearly in the square root of the sample

length, a 15.2% reduction in the standard error corresponds to a 38.9% increase in the sample length, which is effectively an additional 389 months of non-overlapping data or more than 32 years. For quarterly horizons and quarterly sampling of the data in the basic model, we find that ODIN cuts standard errors by approximately 30%, which corresponds to more than doubling the non-overlapping sample length.

The simulations indicate that if the true model were the conditional CAPM with a one-month horizon, the ODIN approach would be substantially more powerful than the basic approach. When we examine actual data, the basic GARCH approach produces a positive conditional risk-return trade-off for the sample period 1955:1 to 2011:12 that has a p -value of .08. When we use the ODIN methodology, the p -value falls to .022. As with much of the literature, though, we find insignificant or even negative trade-offs in other samples and with asymmetric responses to shocks.

Of course, the conditional CAPM is the simplest specification of the ICAPM. Many authors, including Campbell (1996), Scruggs (1998), Guo and Whitelaw (2006), Bali and Engle (2010), and Campbell, Giglio, Polk, and Turley (2012) estimate ICAPMs that include additional state variables.

Some of these papers could be done with ODIN. For example, Scruggs (1998) uses monthly data on the excess market return, the excess return on a long-term bond index, and the risk free rate with QMLE. Monthly measurements of these variables are all available at a daily frequency. Campbell, Giglio, Polk, and Turley (2012) use quarterly data and a six variable vector autoregression. The variables are the quarterly real stock return, the within-quarter realized return volatility from daily data, the price-earnings ratio measured as the price of the S&P 500 index divided by a ten-year trailing moving average of aggregate earnings of companies in the S&P 500 index, the term spread, the small-stock value spread, and the default spread. Only the aggregate earnings variable is truly only measured at the quarterly frequency, and the use of the ten-year moving average of earnings implies that the

earnings part changes very slowly. Thus, one could change the stock price across days within a quarter while keeping the earnings constant throughout the quarter without much loss of content or induced measurement error. This model could therefore be estimated with ODIN, either at the quarterly frequency or the monthly frequency.

We certainly agree that additional state variables, such as the change in the interest rate, are no doubt necessary to adequately capture the changing investment environment faced by investors. We plan to include the conditional covariances of returns with such state variables in future research that investigates the conditional expected returns on assets. More generally, any study that estimates betas with financial data that are available at the daily frequency is a candidate for the ODIN modeling strategy.

A Data and Returns

We start with daily rates of returns, r_{t+1}^d , on the value-weighted stock index as well as monthly returns, $R_{t_m}^f$, on one-month T-bills from CRSP. As the MIDAS model uses daily returns, we proceed as follows: 1) For each month, we construct daily risk-free returns as $r_t^f = (R_{t_m}^f)^{(1/N_m)} - 1$, where N_m is the number of trading days in the month. Hence, we get N_m daily risk-free returns, which are all the same within the month. 2) For the conditional variances of the MIDAS model, we construct daily excess returns as $r_{t+1} = r_{t+1}^d - r_t^f$. 3) For the monthly excess returns that are the dependent variable in the MIDAS and GARCH models, we first compute monthly stock returns and monthly risk-free returns as $R_{t_m}^m = (1 + r_{t+1}^d)(1 + r_{t+2}^d) \cdots (1 + r_{t+N}^d)$ and $R_{t_m}^f = (1 + r_t^f)(1 + r_{t+1}^f) \cdots (1 + r_{t+N-1}^f)$ and then take the difference, $R_{t_m} = R_{t_m}^m - R_{t_m}^f$.

When we estimate the basic monthly or quarterly GARCH-M and MIDAS models, we use actual calendar periods as this has been the standard in the literature. For the ODIN models, we construct returns over 22-day periods, or 66-day periods, for any given starting date and

always estimate the GARCH-M and MIDAS models on the same dependent variable excess returns.

B Simulation and Bootstrapping

B.1 Simulation and Bootstrapping from the GARCH-M Model

Simulating from the GARCH-M model is straightforward. To bootstrap the model, we first construct standardized residuals as

$$\hat{\varepsilon}_{t_m+1} = \frac{R_{t_m+1} - \hat{\mu} - \hat{\gamma}\hat{\sigma}_{t_m}^2}{\hat{\sigma}_{t_m}}$$

where $\hat{\mu}$ and $\hat{\gamma}$ are the estimated parameters, and $\hat{\sigma}_{t_m}^2$ is the estimated conditional variance of R_{t_m+1} . Because the process of standardized residuals does not necessarily have a sample mean of zero and variance of one, we ensure the standardized residuals have mean zero and variance one by calculating $u_{t_m} = (\hat{\varepsilon}_{t_m} - \mu_{\hat{\varepsilon}_{t_m}})/\sigma_{\hat{\varepsilon}_{t_m}}$, where $\mu_{\hat{\varepsilon}_{t_m}}$ and $\sigma_{\hat{\varepsilon}_{t_m}}$ are the sample mean and standard deviation of the $\hat{\varepsilon}_{t_m}$. We then simulate from the GARCH-M model using innovations drawn with replacement from u_1, u_2, \dots, u_T . Estimating the GARCH-M model based on real or bootstrapped data with non-normal innovations can be viewed as quasi-maximum-likelihood (QMLE) and is thus consistent. Note that simulating from the model using innovations that do not have mean zero and variance one would not provide consistent parameter estimates.

B.2 Simulation and Bootstrapping from the MIDAS Model

The MIDAS model is harder to simulate than the GARCH-M model because the MIDAS model is based on daily returns but is estimated on monthly or quarterly returns. When simulating using normal innovations, each daily return is drawn from a standard normal

distribution and is scaled by the conditional standard deviation for that month. We proceed as follows: 1) The conditional variance is based on the previous 500 daily returns. Hence, we first draw 500 standard normal variables and scale them by the unconditional standard deviation of the model, $\sqrt{\omega/(1-\phi)}$. 2) We next calculate the conditional variance for the daily returns over the next month, $\sigma_{t_m}^2 = V_{t_m}^{MIDAS}$ as in equation (7). 3) If there are N_m days in the next month, we draw N_m standard normal variables $u_d, d = 1, \dots, N_m$ and calculate daily returns as $r_{m,d} = \mu + \gamma\sigma_{t_m}^2 + \sqrt{\sigma_{t_m}^2}u_d, d = 1, \dots, N_m$, where $r_{m,d}$ is the return on day d in month m . 4) We repeat steps 2-3 for the following months. The result is a series of daily returns, and the MIDAS model is then calibrated to these daily returns.

To bootstrap the MIDAS model we need daily innovations for the simulation, even though the model is estimated based on monthly data. To obtain this, we ‘interpolate’ daily means and variances in the following way: 1) For each day, we use the estimated parameters to construct a daily forecast of the conditional variance of the daily return, $\hat{\sigma}_t^2 = \hat{\omega} + \hat{\phi} \sum_{d=1}^{500} w(\hat{\kappa})r_{t-d}^2$. 2) For each day, we use the estimated parameters to construct a daily forecast of the conditional mean of the daily return: $\hat{\mu}_t = \hat{\mu} + \hat{\gamma}\hat{\sigma}_t^2$. 3) For each day, we obtain the residual for the daily return as $\hat{\varepsilon}_{t+1} = \frac{r_{t+1} - \hat{\mu}_t}{\sqrt{\hat{\sigma}_t^2}}$.

As above for the GARCH-M model, this process of standardized residuals does not necessarily have a sample mean of zero and a variance of one. Hence, we ensure the standardized residuals have mean zero and variance one by calculating $u_t = (\hat{\varepsilon}_t - \mu_{\hat{\varepsilon}_t})/\sigma_{\hat{\varepsilon}_t}$, and we simulate from the MIDAS model using innovations drawn with replacement from u_1, u_2, \dots, u_T .

References

Andersen, T. G. and T. Bollerslev, 1998. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885-905.

Ang, A. and D. Kristensen, 2012. Testing conditional factor models, *Journal of Financial*

Economics 106, 132-156.

Bali, T. and R. F. Engle, 2010. The intertemporal capital asset pricing model with dynamic conditional correlations, *Journal of Monetary Economics* 57, 377-390.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307-327.

Bollerslev, T. and J. M. Wooldridge, 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time varying covariances, *Econometric Reviews* 11, 143-172.

Campbell, J. Y., 1996. Understanding risk and return, *Journal of Political Economy* 104, 298-345.

Campbell, J. Y. and L. Hentschel, 1992. No news is good news: An asymmetric model of changing volatility. *Journal of Financial Economics* 31, 281-318.

Campbell, J. Y., S. Giglio, C. Polk, and R. Turley, 2012. An Intertemporal CAPM with Stochastic Volatility, manuscript, Harvard University.

Cochrane, J. H. 2005. *Asset Pricing: Revised Edition*. Princeton, NJ: Princeton University Press.

D'Ericco, J., 2011. DERIVEST, <http://www.mathworks.com/matlabcentral/fileexchange/13490-adaptive-robust-numerical-differentiation>.

Drost, F.C. and T.J. Nijman, 1993, Temporal aggregation of GARCH processes, *Econometrica* 61, 909-927.

Engle, R. F. and K. F. Kroner, 1995. Multivariate simultaneous generalized ARCH, *Econometric Theory* 11, 122-150.

Engle, R. F., D. M. Lilien, and R. P. Robins, 1987. Estimating time varying risk premia in the term structure: The ARCH-M model, *Econometrica* 55, 391-407.

Engle, R. F. and V. Ng, 1993. Measuring and testing the impact of news on volatility. *Journal of Finance* 48, 1749-1778.

Fama, E. F. and J. MacBeth, 1973. Risk, return, and equilibrium: Empirical tests, *Journal of Political Economy* 71, 607-636.

Francq, C., L. Horvath, and J. Zakoian, 2011. Merits and drawbacks of variance targeting in GARCH models, *Journal of Financial Econometrics* 9, 619-656.

French, K. R., G. W. Schwert, and R. Stambaugh, 1987. Expected stock returns and volatility, *Journal of Financial Economics* 19, 3-29.

Ghysels, E., P. Santa-Clara, and R. Valkanov, 2005. There is a risk-return trade-off after all, *Journal of Financial Economics* 76, 509-548.

Ghysels, E., A. Sinko, and R. Valkanov, 2007. MIDAS regressions: Further results and new directions, *Econometric Reviews*, 26, 53-90.

Ghysels, E., A. Plazzi, and R. Valkanov, 2013. The risk-return relationship and financial crises, manuscript, University of North Carolina.

Glosten, L. R., R. Jagannathan, and D. E. Runkle, 1993. On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* 48, 1779-1801.

Guo, H., and R. F. Whitelaw, 2006. Uncovering the risk-return relation in the stock market. *Journal of Finance* 61, 1433-1463.

Hansen, L. P., 1982. Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029-1054.

Hansen, L. P., and R. J. Hodrick, 1980. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy* 88, 829-853.

Hansen, L. P. and S. F. Richard, 1987. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models, *Econometrica* 55, 587-613.

Lanne, M. and P. Saikkonen, 2006. Why is it so difficult to uncover the risk-return trade-off in stock returns? *Economic Letters* 92, 118-125.

Lettau, M. and S. C. Ludvigson, 2010. Measuring and modeling variation in the risk-

return trade-off, Chapter 11 in Y. Ait-Sahalia and L. P. Hansen, eds., *Handbook of Financial Econometrics: Volume 1 - Tools and Techniques*, Elsevier: Amsterdam, The Netherlands.

Lewellen, J. and S. Nagel, 2006. The conditional CAPM does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289–314.

Lundblad, C., 2007. The risk return trade off in the long run: 1836–2003, *Journal of Financial Economics* 85, 123–150.

Merton, R. C., 1973. An intertemporal asset pricing model, *Econometrica* 41, 867–887.

Merton, R. C., 1980. On estimating the expected return on the market: an exploratory investigation, *Journal of Financial Economics* 8, 323–361.

Nelson, D. B., 1991. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* 59, 347-370.

Nyberg, H., 2012. Risk-return trade-off in U.S. stock returns over the business cycle, *Journal of Financial and Quantitative Analysis* 47, 137-158.

Richardson, M. and J. Stock, 1989. Drawing inferences from statistics based on multiyear asset returns, *Journal of Financial Economics* 25, 323-348.

Scruggs, J. T., 1998. Resolving the puzzling intertemporal relation between the market risks premium and conditional market variance: A two-factor approach, *Journal of Finance* 53, 575-603.

Valkanov, R., 2003, Long-horizon regressions: Theoretical results and applications, *Journal of Financial Economics* 68, 201-232.

C Online Appendix

C.1 Basic Monthly GARCH model

For the basic monthly GARCH model analyzed in Section 4, this section presents additional figures showing the convergence of the parameter estimates and t -statistics. Figure 10 shows density plots of $\hat{\gamma}$ for the sample lengths $T = 500, 2000,$ and 5000 months under the null of $\gamma = 0$ in the left plots and under the alternative $\gamma = 2$ in the right plot. Similarly, Figure 11 shows density plots of the t -statistics under the null $\gamma = 0$ in the left plot, and under the alternative $\gamma = 2$ in the right plot.

To further assess the convergence of the t -statistics, Figure 12 shows QQ-plots of simulated t -statistics under the null $\gamma = 0$ for different sample lengths. The convergence is excellent for all sample sizes. Further, Figure 13 to 15 show QQ-plots of the t -statistics for the test $\gamma = \gamma_0$, where γ_0 is the true parameter, under the alternatives $\gamma_0 = 1, 2, 3$. Here, we focus on the ‘centralized’ t -statistics

$$t_c = \frac{\hat{\gamma} - \gamma_0}{\hat{\sigma}(\hat{\gamma})}$$

instead of the actual t -statistics

$$t = \frac{\hat{\gamma}}{\hat{\sigma}(\hat{\gamma})}$$

since we know the former should asymptotically follow a standard normal distribution. The convergence is excellent for sample sizes larger than $T = 500$.

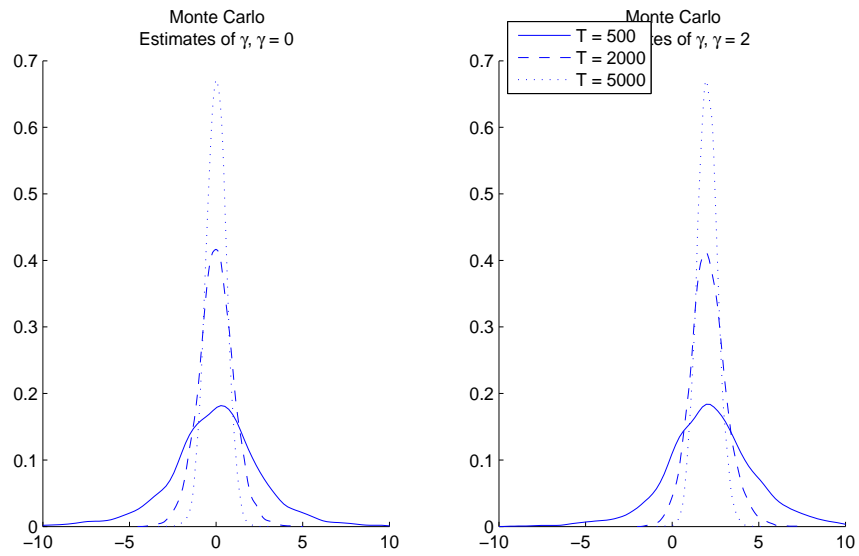


Figure 10: Density Plots of Simulated $\hat{\gamma}$ in the Basic Monthly GARCH Model.

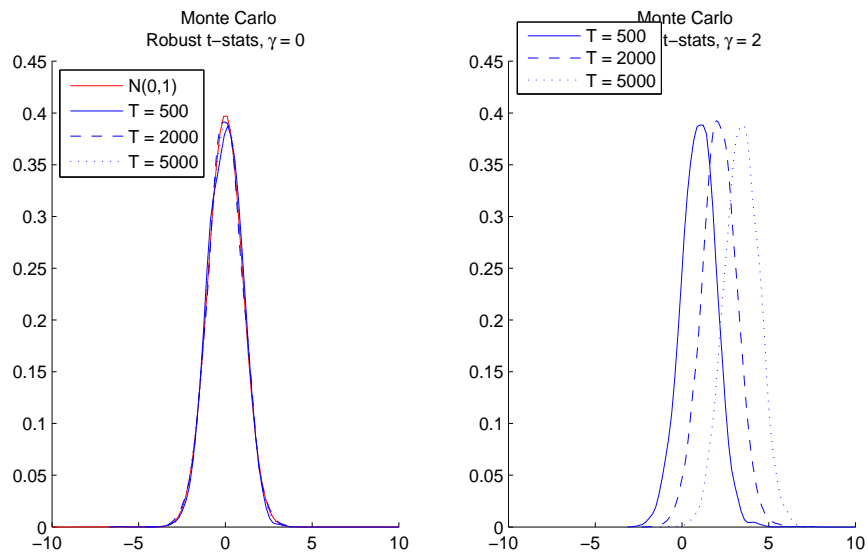


Figure 11: Density Plots of Simulated t -Statistics in the Basic Monthly GARCH Model.

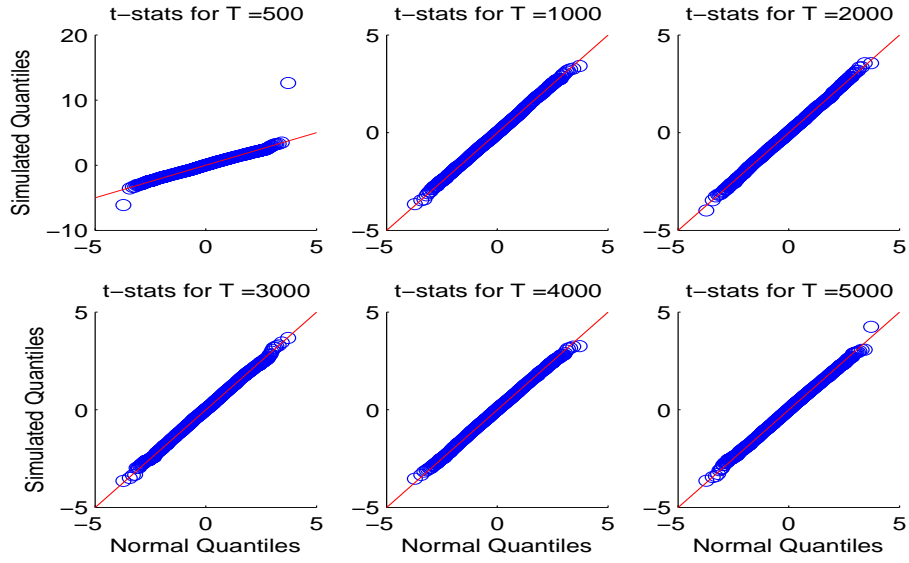


Figure 12: QQ-Plots of Simulated t -Statistics in the Basic Monthly GARCH Model for the test $\gamma = \gamma_0$, $\gamma_0 = 0$.

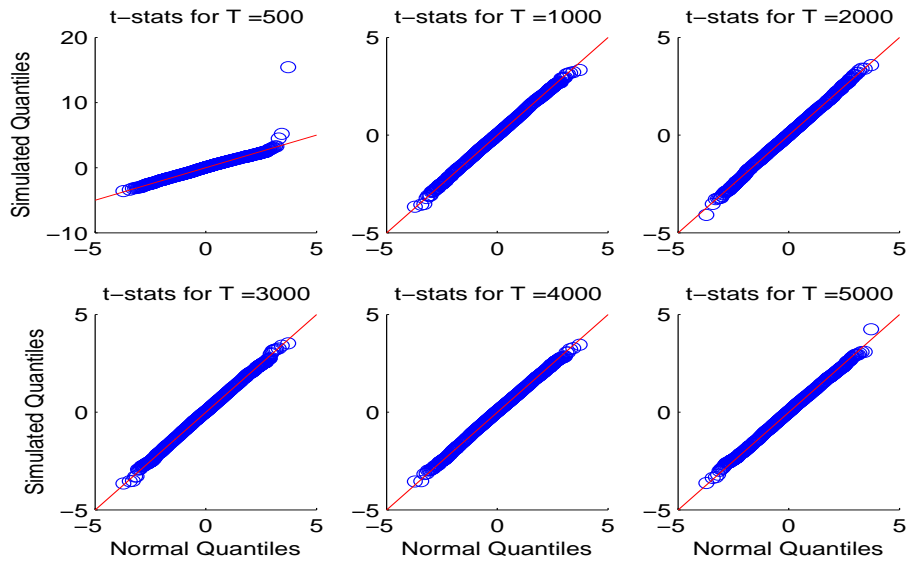


Figure 13: QQ-Plots of Simulated t -Statistics in the Basic Monthly GARCH Model for the test $\gamma = \gamma_0$, $\gamma_0 = 1$.

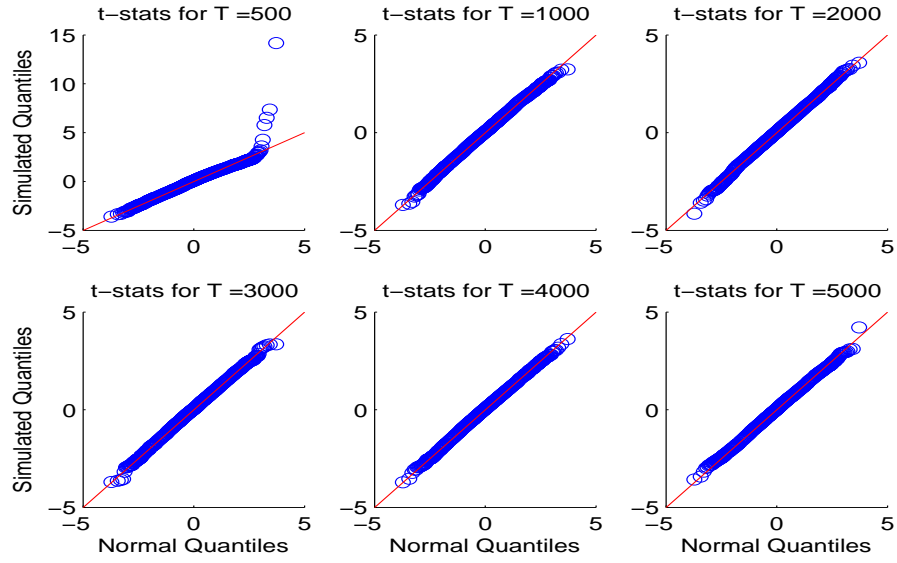


Figure 14: QQ-Plots of Simulated t -Statistics in the Basic Monthly GARCH Model for the test $\gamma = \gamma_0$, $\gamma_0 = 2$.

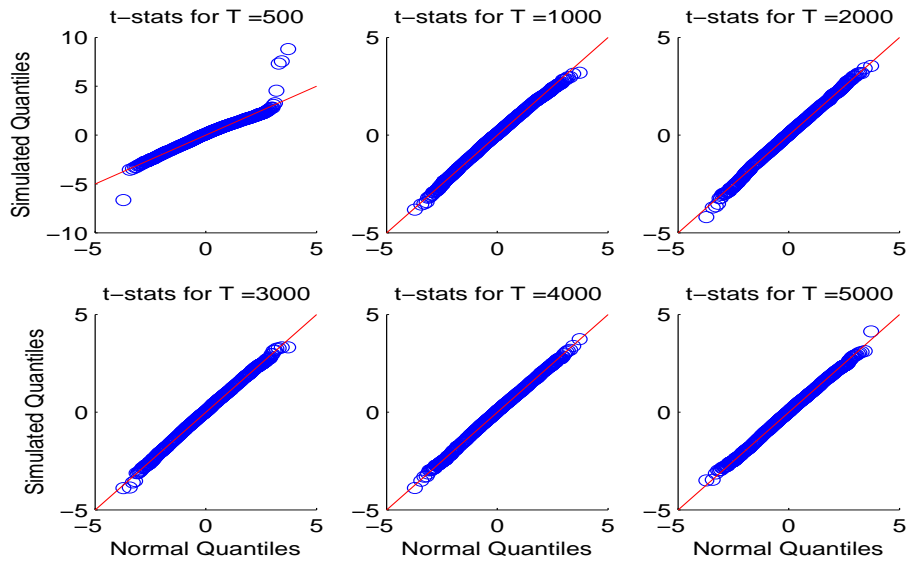


Figure 15: QQ-Plots of Simulated t -Statistics in the Basic Monthly GARCH Model for the test $\gamma = \gamma_0$, $\gamma_0 = 3$.

C.2 Basic Monthly MIDAS model

This section presents QQ-plots of the t -statistics for the test $\gamma = \gamma_0$, $\gamma_0 = 0, 1, 2,$ and 3 for the basic monthly MIDAS model. For small sample sizes, the distributions of the t -statistics have slightly thinner tails than the normal distribution.

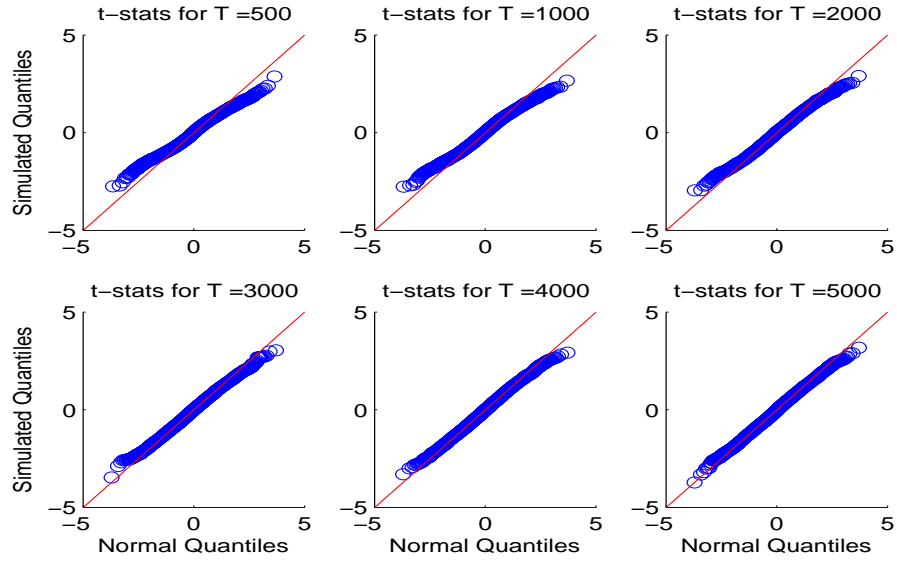


Figure 16: QQ-Plots of Simulated t -Statistics in the Basic Monthly MIDAS Model for the test $\gamma = \gamma_0, \gamma_0 = 0$.

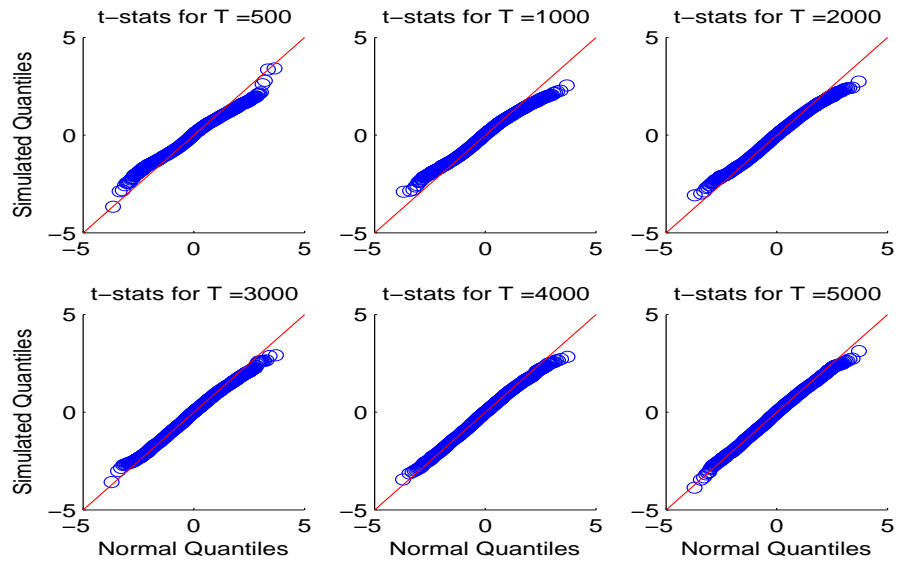


Figure 17: QQ-Plots of Simulated t -Statistics in the Basic Monthly MIDAS Model for the test $\gamma = \gamma_0, \gamma_0 = 1$.

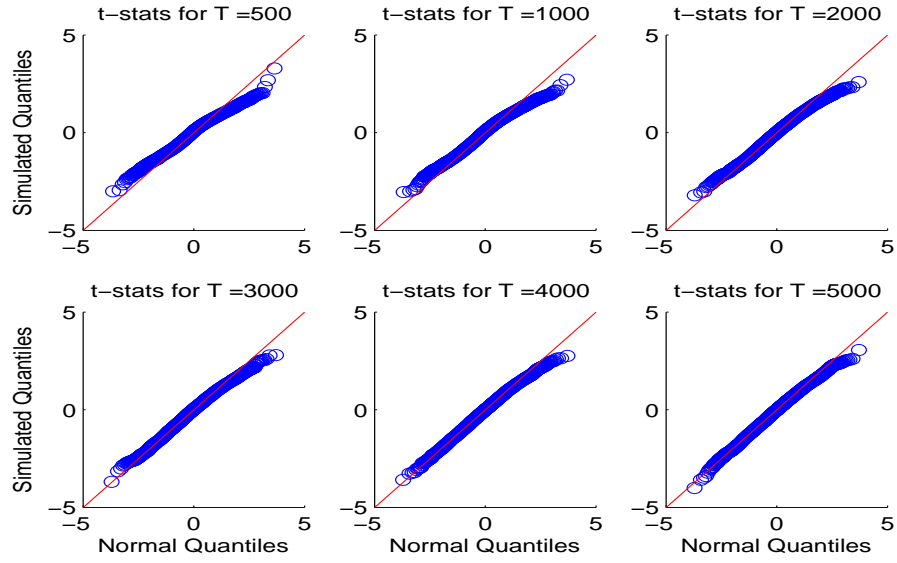


Figure 18: QQ-Plots of Simulated t -Statistics in the Basic Monthly MIDAS Model for the test $\gamma = \gamma_0$, $\gamma_0 = 2$.

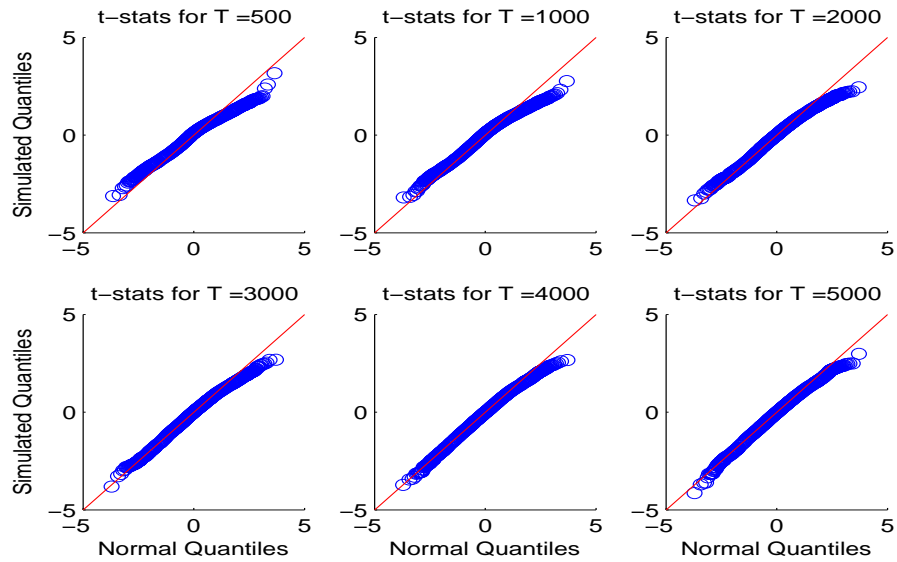


Figure 19: QQ-Plots of Simulated t -Statistics in the Basic Monthly MIDAS Model for the test $\gamma = \gamma_0$, $\gamma_0 = 3$.

C.3 The Equivalence of the Standard Errors of ODIN and Average of Individual Estimates

Section 6 considers the estimator that is the average of the individual estimates from the correlated non-overlapping samples. Here we prove that the standard error of this estimator is the same as the standard error of the ODIN estimator.

C.3.1 The ODIN Estimator

Let y_{t+k} represent a dependent variable realized at time $t + k$; let x_t represent a vector of exogenous or predetermined variables that are in the time t information set; and let θ be the vector of parameters to be estimated. Then, for some vector of functions of the data and the parameters, $g(y_{t+k}, x_t; \theta)$, the orthogonality conditions of the model can be represented as

$$E_t [g(y_{t+k}, x_t; \theta_0)] = 0$$

where θ_0 is the true parameter. Observations on y_{t+k} and x_t are assumed to be available for a sample $t = 1, 2, \dots, T$.

In the ODIN-GMM version of MLE, the model is just identified, so the dimension of $g(y_{t+k}, x_t; \theta)$ is the same as the dimension of θ , which is taken to be n . We simplify the presentation by letting g_t represent $g(y_{t+k}, x_t; \theta)$. The sample average of the orthogonality conditions is

$$G_T(\theta) = \frac{1}{T} \sum_{t=1}^T g_t.$$

The GMM estimator of θ , denoted $\hat{\theta}$, sets $G_T(\hat{\theta}) = 0$ because the system is just identified.

Let

$$D_T(\hat{\theta}) = \nabla_{\theta} G_T(\hat{\theta})$$

be the gradient of $G_T(\hat{\theta})$ with respect to θ evaluated at $\hat{\theta}$. From Hansen's (1982) GMM it

can be demonstrated that the asymptotic distribution of the estimator is characterized by

$$\sqrt{T} \left(\hat{\theta} - \theta_0 \right) \rightarrow N \left(0, D_T(\hat{\theta})^{-1} S D_T^{\top}(\hat{\theta})^{-1} \right) \quad (21)$$

where

$$S = \sum_{j=-k+1}^{k-1} C \left(g_t, g_{t-j}^{\top} \right) \quad (22)$$

and the matrixes, $C \left(g_t, g_{t-j}^{\top} \right)$, are the unconditional autocovariances of g_t . For ease of exposition in what follows, let $C(j)$ denote $C \left(g_t, g_{t-j}^{\top} \right)$. These covariances may be estimated with their sample counterparts. As in Hansen and Hodrick (1980), the covariances are non-zero until the realization of the dependent variable enters the information set that forms the conditional expectation. Thus, the g_t process is a $k - 1$ order vector moving average. This approach uses all of the data.

C.3.2 The Sampled Estimators

Now, consider sampling the data to form a new process g_{t^*} that is serially uncorrelated. There are k different such processes depending on the starting date, so we index them by their starting date, $g_{t^*}(j)$, $j = 1, \dots, k$. Given that there are T total observations, there are $T^* = T/k$ observations on each of the $g_{t^*}(j)$, where for simplicity and without loss of generality we assume that T^* is an integer. The notation indicates that $g_{t^*}(j)$ is observed at $t^* = 1, 2, \dots, T^*$ which corresponds to observing g_t at $t = j, k + j, 2k + j, \dots, T - k + j$. For any particular non-overlapping sample, it remains true that

$$E_{t^*} [g_{t^*}(j)] = 0.$$

Using these orthogonality conditions, one would estimate $\hat{\theta}(j)$, and for any individual

estimator, its asymptotic distribution would be

$$\sqrt{T} \left(\hat{\theta}(j) - \theta_0 \right) \rightarrow N \left(0, k D_{T^*}(\hat{\theta}(j))^{-1} S(j) D_{T^*}^{\top}(\hat{\theta}(j))^{-1} \right)$$

where

$$S(j) = C(0)$$

for each $j = 1, \dots, k$, and the probability limit of $D_{T^*}(\hat{\theta}(j))$ is the same as the probability limit of $D_T(\hat{\theta})$, which we denote D . Hansen and Hodrick (1980) demonstrate that $kC(0)$ exceeds S in equation (22) by a positive definite matrix except in the rare situation in which the parameters of the covariance generating function of g_t are equal. This is the sense in which the ODIN estimator has a smaller asymptotic covariance matrix than any given sampling of the data.

C.3.3 The Average Estimator

Now, let the average of the k different non-overlapping samples be

$$\bar{\theta} = \frac{\sum_{j=1}^k \hat{\theta}(j)}{k}$$

which is clearly consistent for θ_0 . To derive the asymptotic distribution of the average estimator, consider a GMM estimation which stacks the k vectors of parameters into an $nk \times 1$ vector

$$\hat{\Theta} = (\hat{\theta}(1)^{\top}, \dots, \hat{\theta}(k)^{\top})^{\top}$$

and stacks the k sets of orthogonality conditions into

$$G_{t^*} = (g_{t^*}(1)^{\top}, \dots, g_{t^*}(k)^{\top})^{\top}.$$

To derive the asymptotic distribution of $\hat{\Theta}$, define a to be the k -dimensional vector of ones, $a = (1, \dots, 1)^\top$ and define A to be the matrix that has k identity matrixes each of dimension n stacked vertically, $A = (I, I, \dots, I)^\top$. Then, define $\Theta_0 = a \otimes \theta_0$; define the $nk \times nk$ matrix of gradients

$$D^* = (I \otimes D)$$

and define

$$S^* = \sum_{j=-1}^1 E(G_{t^*} G_{t^*-j}^\top)$$

which is also $nk \times nk$. In defining S^* we sum one lagged covariance in each direction because although each of the elements of G_{t^*} is not individually serially correlated, lagged cross-covariances are present. The S^* matrix is composed of three matrixes:

$$E(G_{t^*} G_{t^*}^\top) = \begin{bmatrix} C(0) & C(-1) & C(-2) & \cdot & \cdot & C(-k+1) \\ C(1) & C(0) & C(-1) & \cdot & \cdot & C(-k+2) \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ C(k-1) & C(k-2) & \cdot & \cdot & \cdot & C(0) \end{bmatrix}$$

$$E(G_{t^*} G_{t^*+1}^\top) = \begin{bmatrix} 0 & 0 & \cdot & \cdot & \cdot & 0 \\ C(-k+1) & 0 & \cdot & \cdot & \cdot & 0 \\ C(-k+2) & C(-k+1) & 0 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 0 & \cdot \\ C(-1) & C(-2) & \cdot & \cdot & C(-k+1) & 0 \end{bmatrix}$$

$$E(G_{t^*} G_{t^*-1}^\top) = \begin{bmatrix} 0 & C(k-1) & C(k-2) & \cdot & \cdot & C(1) \\ 0 & 0 & C(k-1) & \cdot & \cdot & C(2) \\ \cdot & \cdot & 0 & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & 0 & C(k-1) \\ 0 & 0 & \cdot & \cdot & \cdot & 0 \end{bmatrix}$$

Then, the joint asymptotic distribution of the k sampled estimators is

$$\sqrt{T}(\hat{\Theta} - \Theta_0) \rightarrow N(0, kD^{*-1}S^*D^{\top*-1}).$$

The average of the k non-overlapping estimators is

$$\bar{\theta} = \frac{A^\top \hat{\Theta}}{k}$$

in which case the asymptotic distribution of $\bar{\theta}$ is

$$\sqrt{T}(\bar{\theta} - \theta_0) \rightarrow N\left(0, \frac{kA^\top D^{*-1}S^*D^{\top*-1}A}{k^2}\right).$$

Multiplying out the matrixes implies that

$$\sqrt{T}(\bar{\theta} - \theta_0) \rightarrow N(0, D^{-1}SD^{\top-1})$$

which is the same asymptotic variance as the asymptotic variance of the ODIN estimator in equation (21).

C.4 Asymmetric Specifications

It has long been noted that volatility seems to increase more in response to negative shocks to returns than to positive shocks. Modeling this asymmetry was introduced by Nelson (1991) and Engle and Ng (1993). While there are many possible asymmetric models, we adopt the straightforward specification of Glosten, Jagannathan, and Runkle (1993) for the asymmetric GARCH model by allowing for a different response parameter controlled by the dummy variable, $1_{(\varepsilon_{t+22,t} < 0)}$, which takes the value 1 if $\varepsilon_{t+22,t} < 0$, and is 0 otherwise:

$$\sigma_{t+22,t+44}^2 = \omega + \alpha \varepsilon_{t+22,t}^2 + \beta \sigma_{t,t+22}^2 + \rho \varepsilon_{t+22,t}^2 1_{(\varepsilon_{t+22,t} < 0)}.$$

The parameter α is thus the base impact of positive squared innovations in returns, and ρ is the additional effect of the squared innovation when the innovation is negative. The estimate of ω is now variance-targeted as $\omega = Var(R)(1 - \alpha - \beta - 0.5\rho)$ where $Var(R)$ is the sample variance of the monthly data.

The specification for the asymmetric MIDAS model is

$$V_{t,t+22}^{MIDAS} = 22 \left(\omega + \phi_1 \sum_{d=0}^D w_{t-d}(\kappa) r_{t-d}^2 + \phi_2 \sum_{d=0}^D w_{t-d}(\kappa) r_{t-d}^2 1_{(r_{t-d} < 0)} \right).$$

The estimate of ω is also variance-targeted. Note that the two weight functions are parameterized by the same κ . Although Ghysels, Santa Clara, and Valkanov (2005) use different weight functions for positive and negative shocks, we find it is difficult to identify different weight functions.

Table ?? presents estimation results for the asymmetric ODIN-GARCH and asymmetric ODIN-MIDAS models. Note that in both models the conditional variances are primarily driven by negative returns. With these asymmetric specifications, none of the $\hat{\gamma}$'s are significant. If the $\hat{\gamma}$'s are positive, they are smaller than their standard errors, and one of the

Table 8: Estimation Results for asymmetric ODIN Models.

Asymmetric ODIN-GARCH							
Period	μ	γ	α	β	ρ	Obs	Sum LLF
1927:9:30–2011:12:8	0.005	0.935	0.048	0.813	0.169	22220	35684.62
Standard error	(0.003)	(1.325)	(0.023)	(0.052)	(0.077)		
<i>p</i> -value	.133	.480	.038	.000	.028		
1927:9:30–1952:12:12	0.008	0.095	0.043	0.871	0.122	7370	10427.29
Standard error	(0.006)	(1.973)	(0.038)	(0.061)	(0.067)		
<i>p</i> -value	.178	.962	.258	.000	.070		
1955:1:26–2011:12:13	0.005	0.447	−0.026	0.717	0.333	14300	24349.28
Standard error	(0.003)	(1.134)	(0.037)	(0.068)	(0.094)		
<i>p</i> -value	.099	.693	.477	.000	.000		
Asymmetric ODIN-MIDAS							
Period	μ	γ	ϕ_1	ϕ_2	κ	Obs	Sum LLF
1927:9:30–2011:12:8	0.005	0.289	0.261	0.818	19.473	22220	35798.07
Standard error	(0.002)	(1.209)	(0.186)	(0.396)	(14.359)		
<i>p</i> -value	.026	.811	.159	.039	.175		
1927:9:30–1952:12:12	0.012	−1.358	0.833	−0.057	3.320	7370	10562.52
Standard error	(0.005)	(1.745)	(0.630)	(1.244)	(1.356)		
<i>p</i> -value	.010	.436	.186	.963	.014		
1955:1:26–2011:12:13	0.004	0.453	0.160	0.743	53.709	14300	24392.73
Standard error	(0.003)	(1.557)	(0.105)	(0.215)	(18.834)		
<i>p</i> -value	.103	.771	.129	.001	.004		

Note: The table present estimates for the asymmetric ODIN-GARCH and ODIN-MIDAS models. Standard errors are in parenthesis.

MIDAS estimates is negative.

C.5 The Continuous Time GARCH Simulations

C.5.1 Strong and Weak GARCH Processes

Consider the stochastic error process $\varepsilon_t, t \in \mathbb{Z}$ and let the sequence $h_t, t \in \mathbb{Z}$ be given by

$$h_t = \omega + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1}. \quad (23)$$

We follow Drost and Nijman (1993) and define strong and weak GARCH processes as follows:

Definition 1. $\varepsilon_t, t \in \mathbb{Z}$ is a strong GARCH process if

$$\frac{\varepsilon_t}{\sqrt{h_t}} \sim i.i.d. D(0, 1) \quad (24)$$

where $D(0, 1)$ specifies a distribution with mean 0 and variance 1.

Definition 2. $\varepsilon_t, t \in \mathbb{Z}$ is a weak GARCH process if the best linear predictor of ε_t and ε_t^2 are, respectively

$$P(\varepsilon_t | \varepsilon_{t-1}, \varepsilon_{t-2}, \dots) = 0 \quad (25)$$

$$P(\varepsilon_t^2 | \varepsilon_{t-1}, \varepsilon_{t-2}, \dots) = h_t, \quad i \geq 0, r = 0, 1, 2 \quad (26)$$

where $P(x_t | \varepsilon_{t-1}, \varepsilon_{t-2}, \dots)$ denotes the best linear predictor of x_t in terms of $1, \varepsilon_{t-1}, \varepsilon_{t-2}, \dots, \varepsilon_{t-1}^2, \varepsilon_{t-2}^2, \dots$.

Drost and Nijman (1993) show that the classical strong GARCH definition is not closed under temporal aggregation (Example 3), but that the class of symmetric weak GARCH models is (Example 1).

C.5.2 The Continuous Time Limit

Alexander and Lazar (2012) discuss several continuous time limits of the GARCH(1,1) model. They argue that to sensibly take the continuous-time limit, the original model must be time-aggregating, and that when sampling the continuous-time process the original discrete-time process should be reproduced. They refer to Drost and Nijman (1993), who show that the weak GARCH process and the continuous-time process defined by

$$d\sigma_t^2 = \theta(\omega - \sigma_t^2)dt + \sqrt{2\lambda\theta}\sigma_t^2 dW_{\sigma,t}. \quad (27)$$

satisfy these requirements. The weak GARCH process is time-aggregating, and when sampling the above continuous-time process, the discretely sampled observations satisfy a weak GARCH process.

C.5.3 Simulations

We start with the continuous-time GARCH-in-mean process

$$dp_t = \lambda\sigma_t^2 dt + \sigma_t dW_{p,t}. \quad (28)$$

where the volatility process is given by

$$d\sigma_t^2 = \theta(\omega - \sigma_t^2)dt + \sqrt{2\lambda\theta}\sigma_t^2 dW_{\sigma,t}. \quad (29)$$

Here $W_{p,t}$ and $W_{\sigma,t}$ are independent Brownian motions (see Andersen and Bollerslev (1998) as well as Alexander and Lazar (2012) who show that the Brownian motions are uncorrelated in the case of a symmetric return distribution). Letting m denote the intraday sampling

frequency, consider the intraday return

$$r_{(m),t} = p_t - p_{t-1/m}. \quad (30)$$

The transformation between the parameters $(\delta_0, \alpha, \beta)$ in the discrete-time GARCH(1,1) model, estimated based on m intradaily returns, and the parameters $(\theta, \omega, \lambda)$ in the continuous time diffusion in (29) is given by equations (10)-(12) in Andersen and Bollerslev (1998) (see Section C.5.5). This transformation is a limiting result of a relationship between the parameters of a weak GARCH process at two frequencies, as one of the frequencies approaches continuous sampling. Since we simulate in 5-minute increments below, we use the limiting result for the parameters.

Since the GARCH(1,1) model in Lundblad (2007) is estimated based on monthly observations, this corresponds to $m = 1/22$ (less than one observation per day). Using monthly GARCH(1,1) parameters of $\delta_0 = 0.0002, \alpha = 0.1, \beta = 0.85$, and using formulas (10)-(12) in Andersen and Bollerslev (1998) (note that there is a typo in equation (12)), the continuous time parameters become

$$\theta = 0.0023, \quad \omega = 1.8182 \cdot 10^{-4}, \quad \lambda = 0.459 \quad (31)$$

which is close to what Lundblad (2007) reports ($\theta = 0.0023, \omega = 0.0001, \lambda = 0.459$).

We then simulate from the continuous time model given by (28) and (29) using a standard Euler scheme:

$$\Delta p_t = \lambda \sigma_t^2 \Delta + \sigma_t \sqrt{\Delta} W_{p,t} \quad (32)$$

$$\Delta \sigma_t^2 = \theta(\omega - \sigma_t^2) \Delta + \sqrt{2\lambda\theta} \sigma_t^2 \sqrt{\Delta} W_{\sigma,t} \quad (33)$$

where $W_{p,t}, W_{\sigma,t} \sim N(0, 1)$. As in Lundblad (2007), we simulate 5 minute returns, hence

$\Delta = 1/288$.

Finally, sample the process to get daily log prices: $p_{\text{daily}} = p(288 : 288 : \text{end})$, i.e., we use every 288th observations. These are the daily log-prices, so daily returns are e.g. $r(1) = p_{\text{daily}}(2) - p_{\text{daily}}(1)$. Summing these gives monthly returns for any given start date, and these returns satisfy a weak GARCH model.

C.5.4 Estimation

As mentioned above, Andersen and Bollerslev (1998) and Drost and Nijman (1993) show how discretely sampling the continuous time GARCH process results in observations that satisfy a weak GARCH(1,1) model. We estimate the parameters using the usual QMLE procedure which is based on the likelihood function for the strong GARCH model with normal innovations. The first-order conditions arising from this likelihood function are not necessarily satisfied for the weak GARCH model, and as a result, the usual QMLE theory does not apply: There is no guarantee that the parameter estimates converge in probability to the values used in the data generating process as the sample size is increased. Drost and Nijman (1993) note that they've conducted extensive simulation experiments showing that the asymptotic bias of QMLE, if any, is small. This is in agreement with our findings that when simulating from a continuous-time GARCH model, and subsequently estimating Basic Monthly GARCH and ODIN GARCH on discretely sampled observations, there is a small bias in the estimate of γ . Note that this caveat applies to both the basic and the ODIN GARCH models.

C.5.5 The Transformation Equations

Suppose the econometrician estimates a GARCH model based on m observations per day (so $m = 1/22$ for monthly observations). Andersen and Bollerslev (1998) show that the exact one-to-one relationship between the discrete-time weak GARCH(1,1) parameters and

the continuous-time stochastic volatility parameters are (note the typo in the second line for λ : it should be $+\alpha_{(m)} [1 - \beta_{(m)} (\alpha_{(m)} + \beta_{(m)})]$) as here)

$$\theta = -m \log(\alpha_{(m)} + \beta_{(m)}) \quad (34)$$

$$\omega = m\omega_{(m)} (1 - \alpha_{(m)} - \beta_{(m)})^{-1} \quad (35)$$

$$\lambda = 2\alpha_{(m)} \log^2 (\alpha_{(m)} + \beta_{(m)}) [1 - \beta_{(m)} (\alpha_{(m)} + \beta_{(m)})] \quad (36)$$

$$\left\{ \left[1 - (\alpha_{(m)} + \beta_{(m)})^2 \right] (1 - \beta_{(m)})^2 + \alpha_{(m)} [1 - \beta_{(m)} (\alpha_{(m)} + \beta_{(m)})] \right. \quad (37)$$

$$\left. \left[6 \log (\alpha_{(m)} + \beta_{(m)}) + 2 \log^2 (\alpha_{(m)} + \beta_{(m)}) + 4 (1 - \alpha_{(m)} - \beta_{(m)}) \right] \right\}^{-1} \quad (38)$$

That is, to transform from daily GARCH parameters to continuous-time parameters, set $m = 1$. To transform from monthly parameters, set $m = 1/22$.

C.6 ODIN Estimations

C.6.1 Point Estimates: Box Plots

Figure 20, 21, 22, and 23 show box-plots of $\hat{\gamma}$ for $\gamma = 0, 1, 2$, and 3, respectively. For each sampling frequency, the left blue box-plot shows the results for the basic GARCH, and the right green box-plot shows the results for the ODIN GARCH. In the right figure, showing sampling frequencies 33-66, the values have been truncated at 50.

Beneath each box-plot is the number of estimations that did not converge (out of 10,000). For instance, for $\gamma = 0$ and a 33 day sampling frequency, 88 of the basic estimations did not converge, and 5 of the ODIN estimations did not converge (note that they're estimated on exactly the same simulated data).

It is clear from the figure that the precision of the ODIN GARCH estimates is much higher than the precision of the basic GARCH estimates for all sampling frequencies and all values of γ . Further, the increase in precision increases with the sampling frequency.

Looking at the figures for the different values of γ , it is also clear that the estimates from the basic model become increasingly skewed with fat right tails. For instance, compare the box-plots for sampling frequency 66. When $\gamma = 0$, the distribution of $\hat{\gamma}$ is roughly symmetric around 0. But for $\gamma = 3$, there are no estimates below 30, and many more estimates above 50.




Figure 20: Box-plots of $\hat{\gamma}$ when $\gamma = 0$, for different sampling frequencies. For each sampling frequency, the left box-plots shows the results for the basic GARCH, and the right box-plot shows the results for the ODIN GARCH. In the right, the values have been truncated at 50. Beneath each box-plot is the number of estimations that did not converge (out of 10,000).




Figure 21: Box-plots of $\hat{\gamma}$ when $\gamma = 1$, for different sampling frequencies. For each sampling frequency, the left box-plots shows the results for the basic GARCH, and the right box-plot shows the results for the ODIN GARCH. In the right, the values have been truncated at 50. Beneath each box-plot is the number of estimations that did not converge (out of 10,000).




Figure 22: Box-plots of $\hat{\gamma}$ when $\gamma = 2$, for different sampling frequencies. For each sampling frequency, the left box-plots shows the results for the basic GARCH, and the right box-plot shows the results for the ODIN GARCH. In the right, the values have been truncated at 50. Beneath each box-plot is the number of estimations that did not converge (out of 10,000).


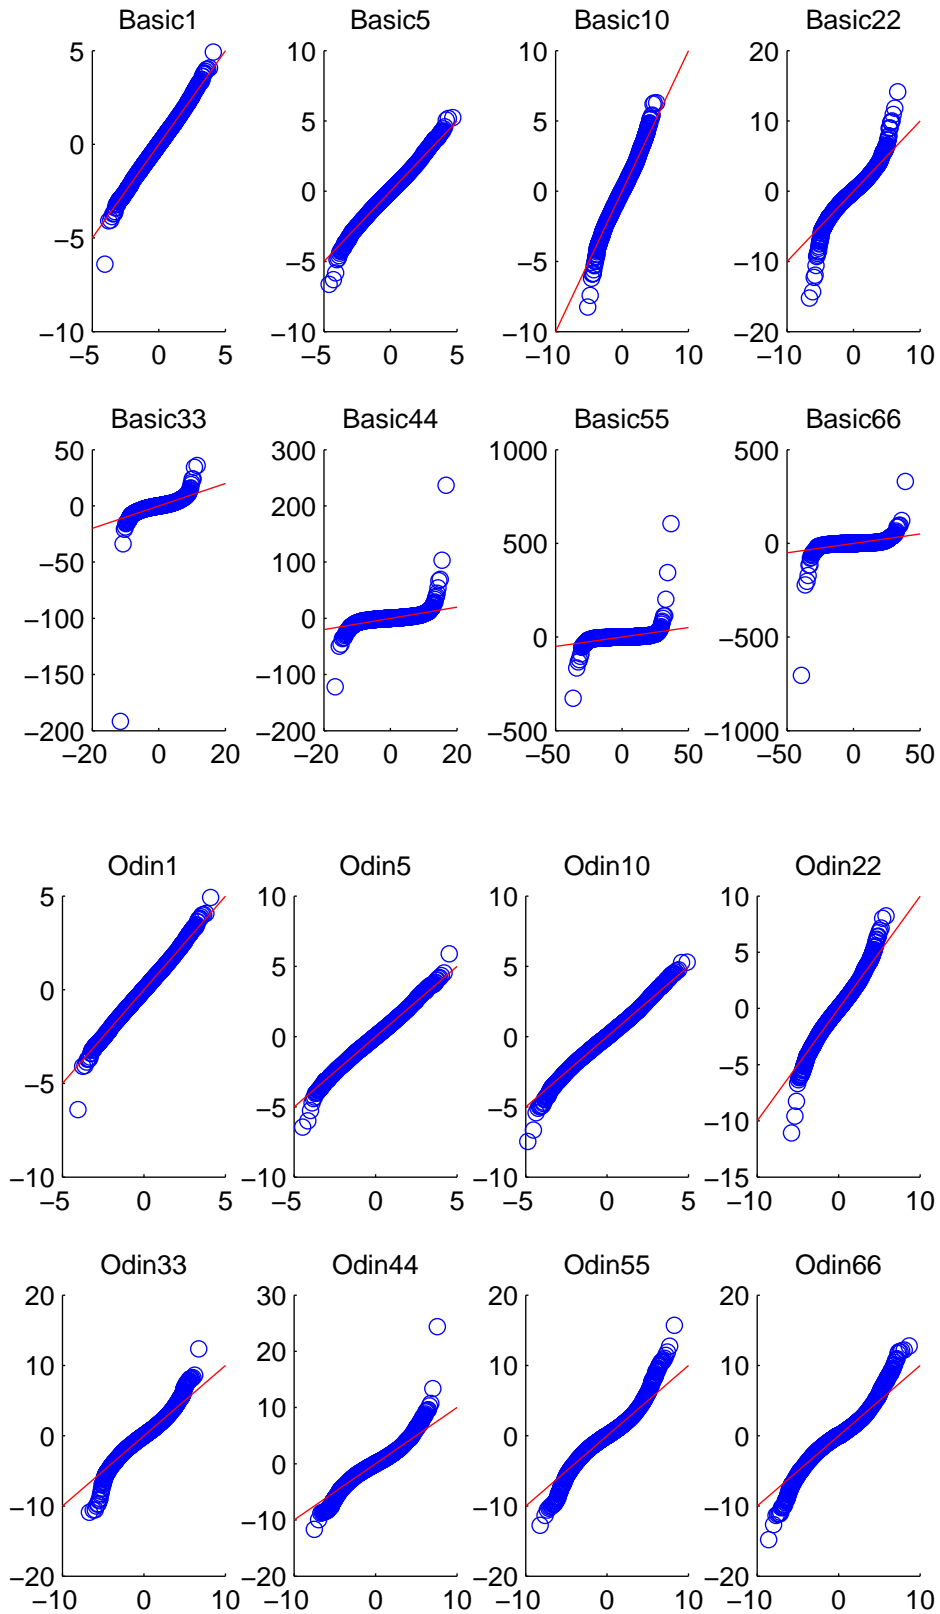


Figure 23: Box-plots of $\hat{\gamma}$ when $\gamma = 3$, for different sampling frequencies. For each sampling frequency, the left box-plots shows the results for the basic GARCH, and the right box-plot shows the results for the ODIN GARCH. In the right, the values have been truncated at 50. Beneath each box-plot is the number of estimations that did not converge (out of 10,000).

C.6.2 Point Estimates: QQ-Plots

Figure 24–27 show QQ-plots of $\hat{\gamma}$ for the basic model and the ODIN model, for $\gamma = 0, 1, 2,$ and 3, respectively. On each page, the top figure shows QQ-plots for the basic model, and the bottom figure shows QQ-plots for the ODIN model.

Comparing the QQ-plots for basic GARCH and ODIN GARCH for a given sampling frequency, it's clear that the distribution of $\hat{\gamma}$ is closer to a normal distribution for the ODIN model. The distributions of the estimates for the basic model have much fatter tails.



76
 Figure 24: QQ-plots of $\hat{\gamma}$ when $\gamma = 0$, for different sampling frequencies.

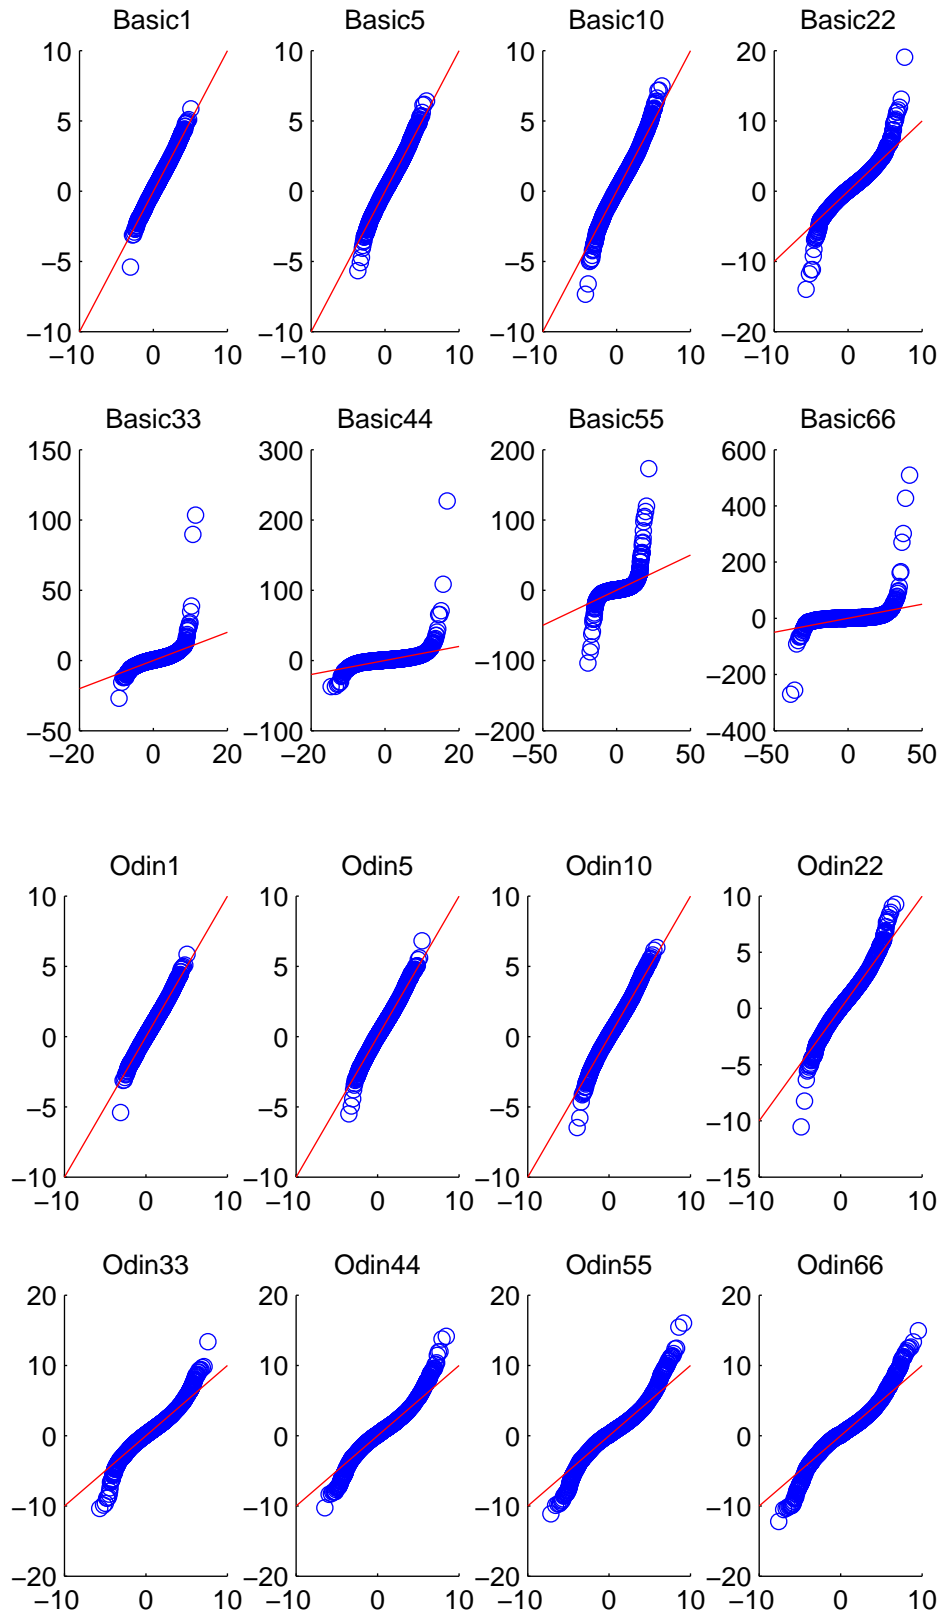
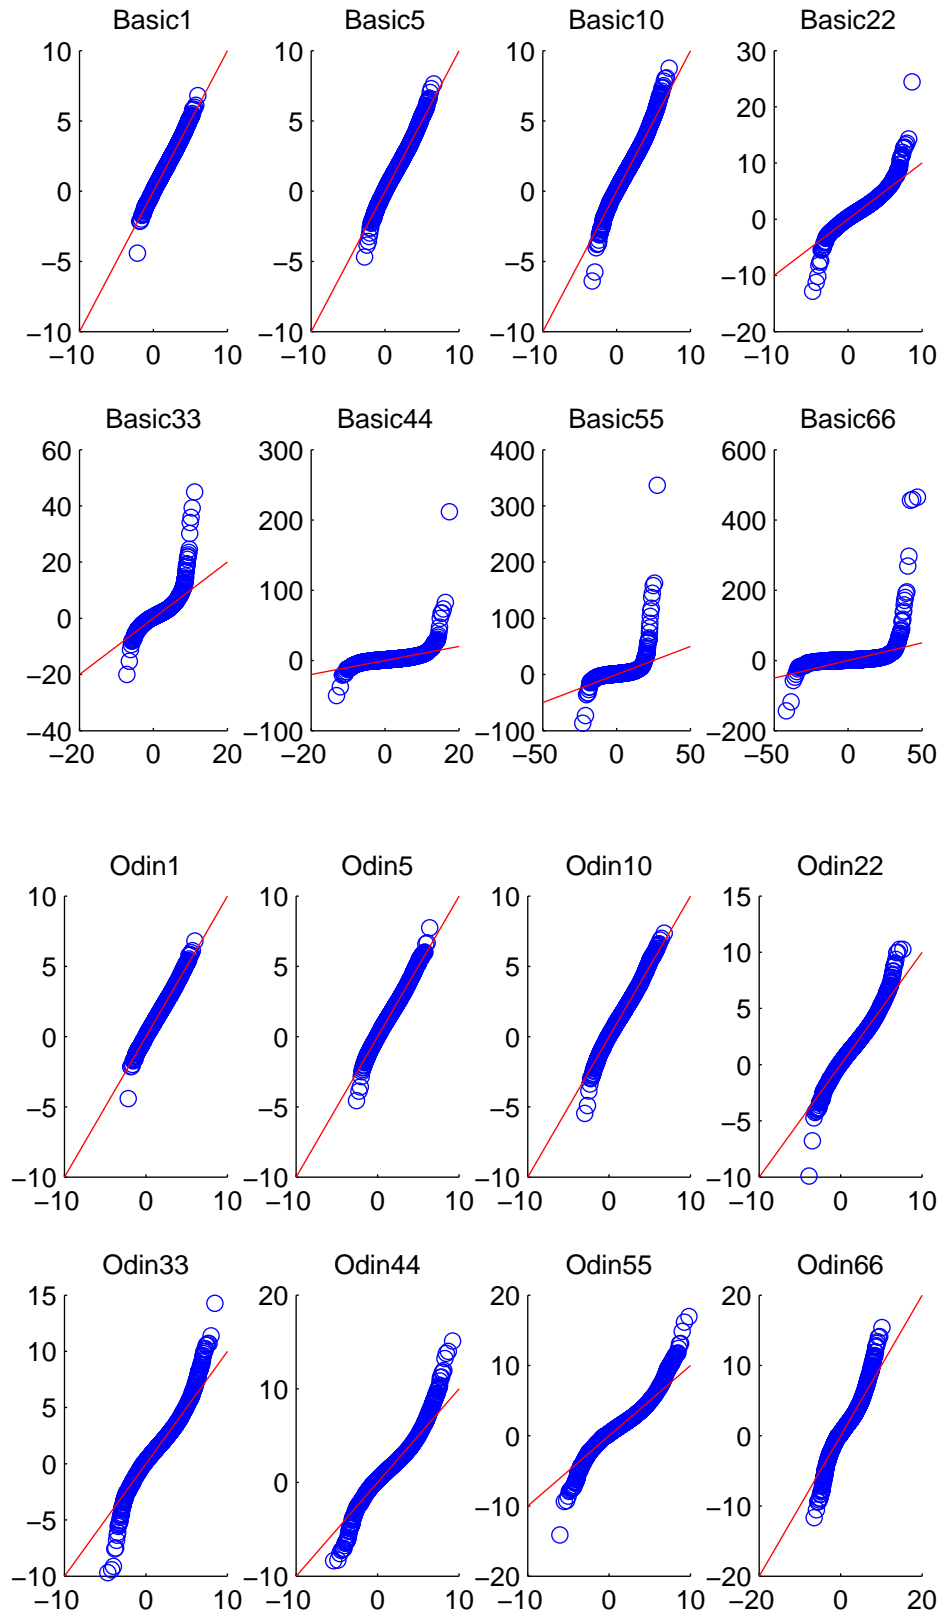
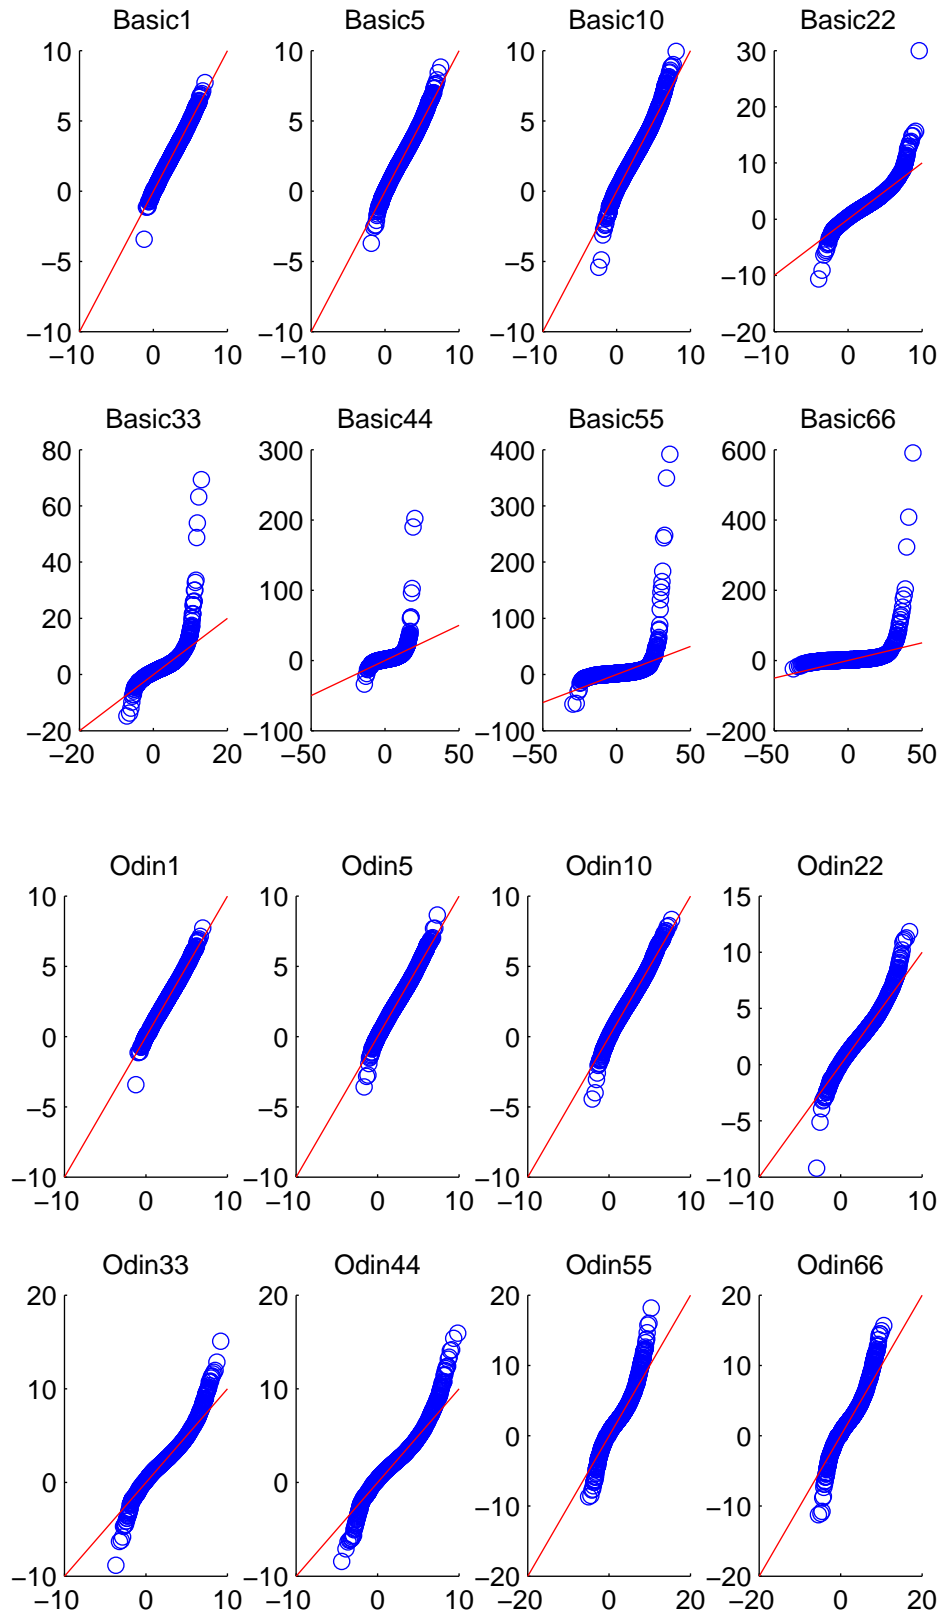


Figure 25: QQ-plots of $\hat{\gamma}$ when $\gamma = 0$, for different sampling frequencies.



78
 Figure 26: QQ-plots of $\hat{\gamma}$ when $\gamma = 0$, for different sampling frequencies.



79
 Figure 27: QQ-plots of $\hat{\gamma}$ when $\gamma = 0$, for different sampling frequencies.

C.6.3 Standard Errors: Precision

Here, we compare the estimated standard errors to the empirical standard deviations of $\hat{\gamma}$. Table 9 shows the results for the standard error estimates. The first two lines of the table show the empirical standard deviation of $\hat{\gamma}$ for the basic and ODIN models. The next two lines show the average of the 10,000 estimated standard errors. The last two rows show the bias in the estimated standard errors. As we saw above, the standard deviation of the 10,000 estimates of γ is much smaller for the ODIN model which we now see in rows 1 and 2. Next, the estimated standard errors, $\hat{\sigma}(\hat{\gamma})$, are also smaller for the ODIN model, as seen in rows 3 and 4 where I calculate the average estimated standard error. Finally, rows 5 and 6 show that the standard error estimates are generally too small, but more so for the basic model (the basic model has a higher bias).

Table 10-12 show the results for $\gamma = 1, 2, 3$. These tables show that the bias of the ODIN standard errors are generally smaller than the bias of the standard errors for the basic model.

Table 9: Table for $\gamma = 0$. The first two lines of the table shows the empirical standard deviation of $\hat{\gamma}_i$ for the basic and ODIN models, based on all 10,000 estimates. The next two lines show the average of the 10,000 estimated standard errors. The last two rows shows the bias in the estimated standard errors.

Sampling Frequency	1	5	10	22	33	44	55	66
Empirical Std of $\hat{\gamma}$, Basic	1.05	1.19	1.33	1.71	2.98	4.34	9.78	10.66
Empirical Std of $\hat{\gamma}$, ODIN	1.05	1.15	1.26	1.49	1.72	1.94	2.13	2.25
Average of $\hat{\sigma}(\hat{\gamma})$, Basic	1.02	1.15	1.28	1.61	2.28	3.86	7.70	9.40
Average of $\hat{\sigma}(\hat{\gamma})$, ODIN	1.02	1.11	1.20	1.41	1.60	1.78	1.93	2.06
Basic Bias	-0.03	-0.04	-0.04	-0.05	-0.24	-0.11	-0.21	-0.12
ODIN Bias	-0.03	-0.04	-0.04	-0.06	-0.07	-0.08	-0.09	-0.09

Table 10: Table for $\gamma = 1$. The first two lines of the table shows the empirical standard deviation of $\hat{\gamma}_i$ for the basic and ODIN models, based on all 10,000 estimates. The next two lines show the average of the 10,000 estimated standard errors. The last two rows shows the bias in the estimated standard errors.

Sampling Frequency	1	5	10	22	33	44	55	66
Empirical Std of $\hat{\gamma}$, Basic	1.05	1.20	1.33	1.71	2.67	4.08	5.46	10.97
Empirical Std of $\hat{\gamma}$, ODIN	1.05	1.16	1.26	1.49	1.71	1.92	2.10	2.23
Average of $\hat{\sigma}(\hat{\gamma})$, Basic	1.01	1.15	1.28	1.61	2.26	3.73	5.16	14.45
Average of $\hat{\sigma}(\hat{\gamma})$, ODIN	1.02	1.11	1.20	1.40	1.59	1.77	1.91	2.01
Basic Bias	-0.03	-0.04	-0.04	-0.06	-0.15	-0.09	-0.05	0.32
ODIN Bias	-0.03	-0.04	-0.05	-0.06	-0.07	-0.08	-0.09	-0.10

Table 11: Table for $\gamma = 2$. The first two lines of the table shows the empirical standard deviation of $\hat{\gamma}_i$ for the basic and ODIN models, based on all 10,000 estimates. The next two lines show the average of the 10,000 estimated standard errors. The last two rows shows the bias in the estimated standard errors.

Sampling Frequency	1	5	10	22	33	44	55	66
Empirical Std of $\hat{\gamma}$, Basic	1.05	1.20	1.34	1.73	2.36	3.96	6.65	11.97
Empirical Std of $\hat{\gamma}$, ODIN	1.05	1.16	1.26	1.48	1.69	1.87	2.04	2.15
Average of $\hat{\sigma}(\hat{\gamma})$, Basic	1.01	1.15	1.27	1.61	2.07	3.02	5.98	13.68
Average of $\hat{\sigma}(\hat{\gamma})$, ODIN	1.01	1.11	1.20	1.39	1.56	1.73	1.86	1.95
Basic Bias	-0.04	-0.04	-0.05	-0.07	-0.12	-0.24	-0.10	0.14
ODIN Bias	-0.04	-0.04	-0.05	-0.06	-0.07	-0.08	-0.09	-0.09

Table 12: Table for $\gamma = 3$. The first two lines of the table shows the empirical standard deviation of $\hat{\gamma}_i$ for the basic and ODIN models, based on all 10,000 estimates. The next two lines show the average of the 10,000 estimated standard errors. The last two rows shows the bias in the estimated standard errors.

Sampling Frequency	1	5	10	22	33	44	55	66
Empirical Std of $\hat{\gamma}$, Basic	1.05	1.20	1.35	1.75	2.59	4.42	8.63	10.82
Empirical Std of $\hat{\gamma}$, ODIN	1.05	1.16	1.25	1.47	1.65	1.82	1.97	2.05
Average of $\hat{\sigma}(\hat{\gamma})$, Basic	1.01	1.14	1.27	1.58	2.14	3.24	7.24	14.40
Average of $\hat{\sigma}(\hat{\gamma})$, ODIN	1.01	1.10	1.19	1.36	1.52	1.67	1.78	1.85
Basic Bias	-0.04	-0.05	-0.06	-0.09	-0.17	-0.27	-0.16	0.33
ODIN Bias	-0.04	-0.05	-0.05	-0.07	-0.08	-0.09	-0.10	-0.10

C.6.4 t -Statistics: QQ-plots

This section presents QQ-plots of the t -statistics for the test used to construct Figure 4. Each Figure shows 8 QQ-plots for the sampling frequencies 1, 5, 10, 22, 33, 44, 55, and 66 days. The data generating process is a continuous time GARCH process. For each value of $\gamma = 0, 1, 2, 3$, there are two figures. The first shows QQ-plots for the basic (non-overlapping) GARCH model which samples the data at the above frequencies. The second shows QQ-plots for the ODIN model that uses all of the available data.

Figure 28 shows a slight asymmetry in the t -statistics for the basic GARCH model, and the 97.5th percentile is actually lower than 1.96. When constructing Figure 4 we use the empirical quantiles, and using 1.96 instead would therefore *decrease* the power of the basic GARCH model and thus *increase* the improvement of the ODIN model over the basic GARCH model.

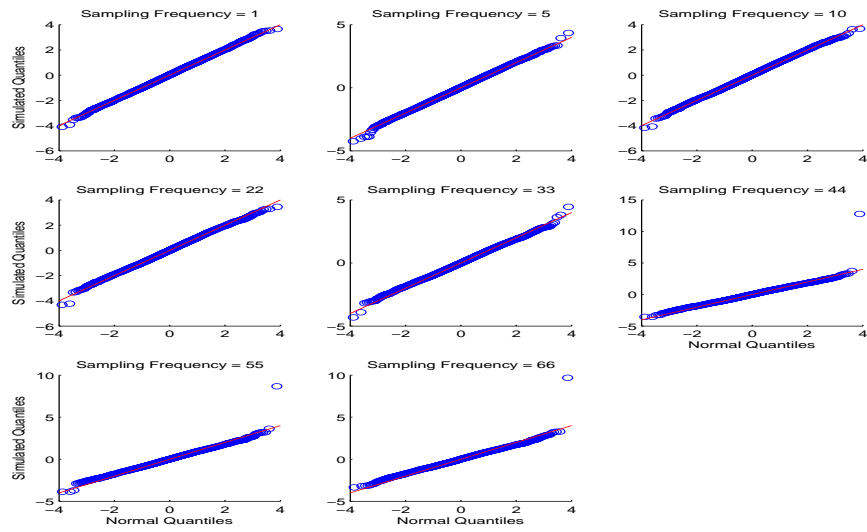


Figure 28: QQ-Plots of Simulated t -Statistics in the Basic GARCH Model for different sampling frequencies. The model is estimated to data simulated from a continuous time GARCH model as described in the text. The t -statistics are for the test $\gamma = \gamma_0$, $\gamma_0 = 0$.

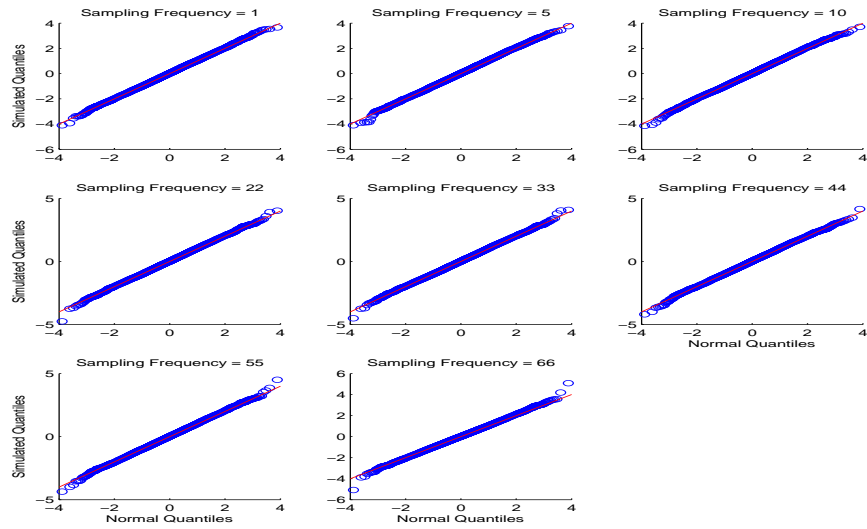


Figure 29: QQ-Plots of Simulated t -Statistics in the ODIN GARCH Model for different sampling frequencies. The model is estimated to data simulated from a continuous time GARCH model as described in the text. The t -statistics are for the test $\gamma = \gamma_0$, $\gamma_0 = 0$.

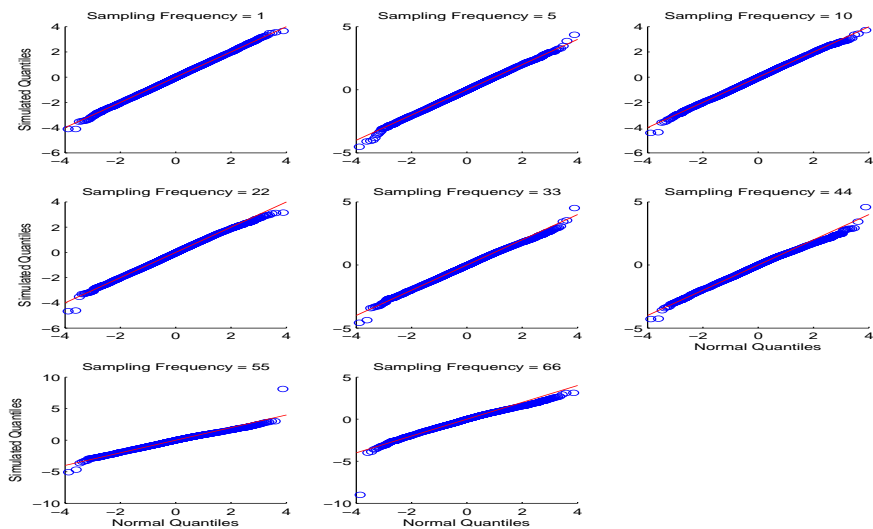


Figure 30: QQ-Plots of Simulated t -Statistics in the Basic GARCH Model for different sampling frequencies. The model is estimated to data simulated from a continuous time GARCH model as described in the text. The t -statistics are for the test $\gamma = \gamma_0$, $\gamma_0 = 1$.

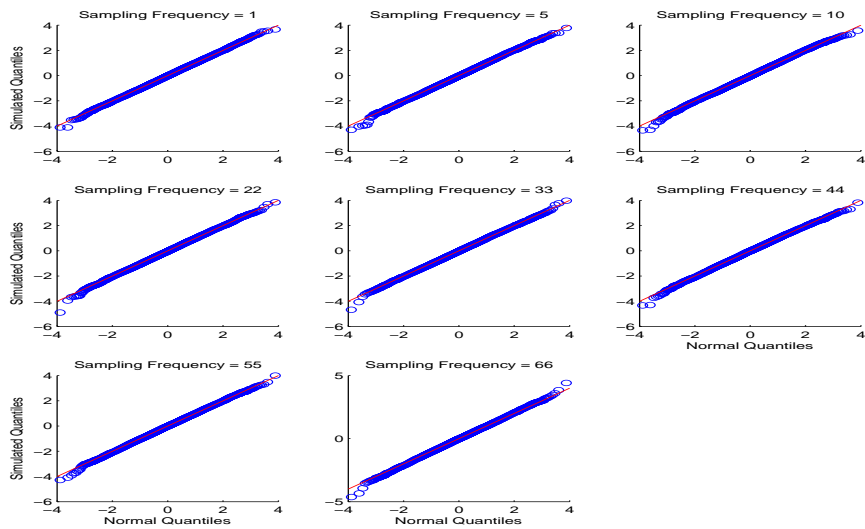


Figure 31: QQ-Plots of Simulated t -Statistics in the ODIN GARCH Model for different sampling frequencies. The model is estimated to data simulated from a continuous time GARCH model as described in the text. The t -statistics are for the test $\gamma = \gamma_0$, $\gamma_0 = 1$.

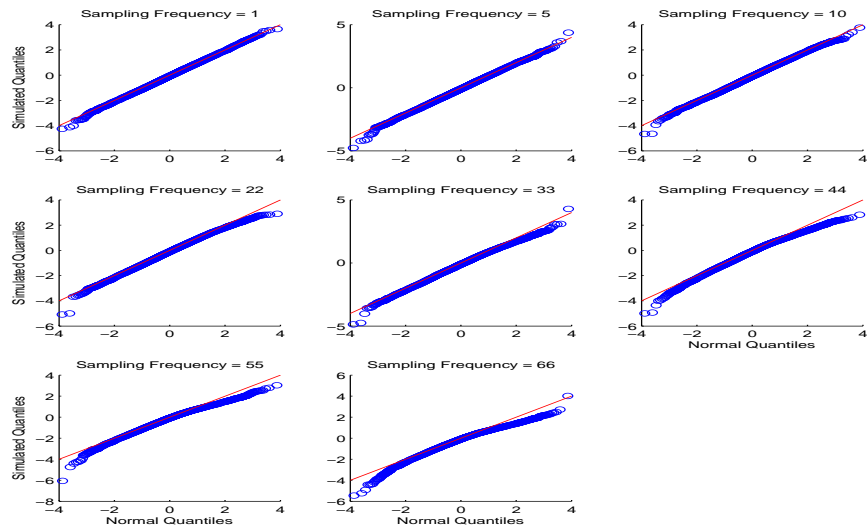


Figure 32: QQ-Plots of Simulated t -Statistics in the Basic GARCH Model for different sampling frequencies. The model is estimated to data simulated from a continuous time GARCH model as described in the text. The t -statistics are for the test $\gamma = \gamma_0$, $\gamma_0 = 2$.

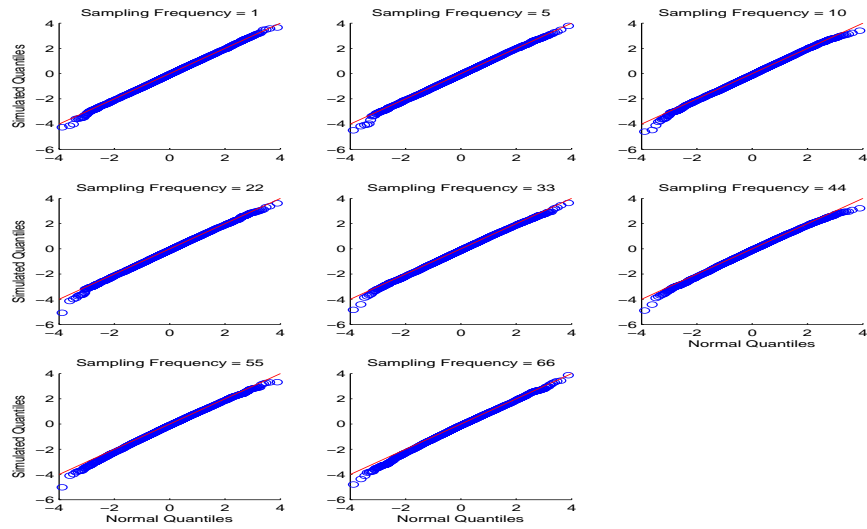


Figure 33: QQ-Plots of Simulated t -Statistics in the ODIN GARCH Model for different sampling frequencies. The model is estimated to data simulated from a continuous time GARCH model as described in the text. The t -statistics are for the test $\gamma = \gamma_0$, $\gamma_0 = 2$.

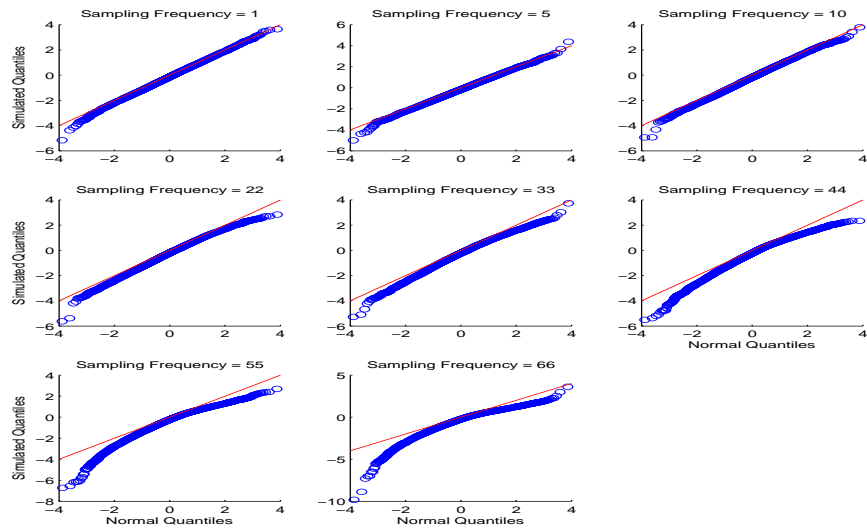


Figure 34: QQ-Plots of Simulated t -Statistics in the Basic GARCH Model for different sampling frequencies. The model is estimated to data simulated from a continuous time GARCH model as described in the text. The t -statistics are for the test $\gamma = \gamma_0$, $\gamma_0 = 3$.

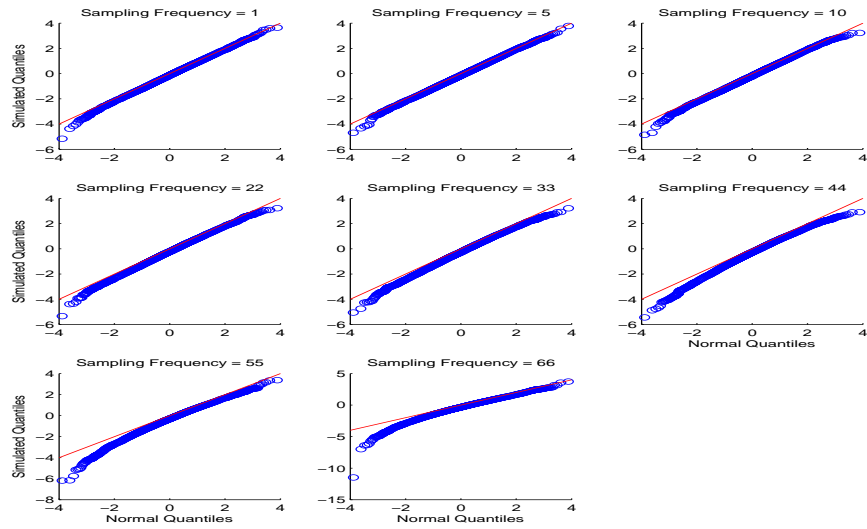


Figure 35: QQ-Plots of Simulated t -Statistics in the ODIN GARCH Model for different sampling frequencies. The model is estimated to data simulated from a continuous time GARCH model as described in the text. The t -statistics are for the test $\gamma = \gamma_0$, $\gamma_0 = 3$.

C.6.5 Large Sample Properties: 5,000 Months

Table 13: Empirical Means.

Forecast Horizon	1	5	10	22	33	44	55	66
Panel A: 1,000 Months								
$\gamma = 0$								
Basic GARCH	0.0273	0.0353	0.0371	0.0215	0.0042	0.0739	0.1118	-0.0222
ODIN GARCH	0.0273	0.0273	0.0258	0.0228	0.0229	0.0360	0.0393	0.0434
$\gamma = 1$								
Basic GARCH	0.9808	0.9777	0.9791	0.9889	1.0621	1.1657	1.2817	1.5190
ODIN GARCH	0.9808	0.9752	0.9686	0.9702	0.9795	1.0003	1.0149	1.0031
$\gamma = 2$								
Basic GARCH	1.9305	1.9139	1.9100	1.9302	2.0185	2.1966	2.4449	2.8619
ODIN GARCH	1.9305	1.9189	1.9023	1.8943	1.8997	1.9182	1.9195	1.8949
$\gamma = 3$								
Basic GARCH	2.8722	2.8379	2.8194	2.8235	2.9354	3.1225	3.4697	3.7735
ODIN GARCH	2.8722	2.8542	2.8185	2.7750	2.7535	2.7412	2.7100	2.6651
Panel B: 5,000 Months								
$\gamma = 0$								
Basic GARCH	0.0103	0.0073	0.0076	0.0129	0.0086	0.0486	0.0302	0.0256
ODIN GARCH	0.0103	0.0102	0.0124	0.0213	0.0277	0.0319	0.0336	0.0388
$\gamma = 1$								
Basic GARCH	0.9772	0.9516	0.9415	0.9434	0.9396	0.9791	0.9686	0.9747
ODIN GARCH	0.9772	0.9558	0.9475	0.9460	0.9493	0.9524	0.9553	0.9605
$\gamma = 2$								
Basic GARCH	1.9429	1.8914	1.8652	1.8474	1.8272	1.8480	1.8275	1.8288
ODIN GARCH	1.9429	1.8974	1.8727	1.8445	1.8284	1.8139	1.8020	1.7924
$\gamma = 3$								
Basic GARCH	2.9060	2.8224	2.7691	2.7013	2.6357	2.6116	2.5575	2.5244
ODIN GARCH	2.9060	2.8310	2.7788	2.6936	2.6300	2.5711	2.5190	2.4708

Note: The table shows the empirical means of $\hat{\gamma}$ in the basic and ODIN models when estimating the models on simulated data from a continuous time GARCH model. In Panel A, the models are estimated on 1,000 months of simulated data, and in Panel B they are estimated on 5,000 months of data. The value of γ used in the simulations varies from 0 to 3, and the forecast horizon varies from 1 to 66 days.

C.7 Quarterly Individual vs. ODIN Estimates

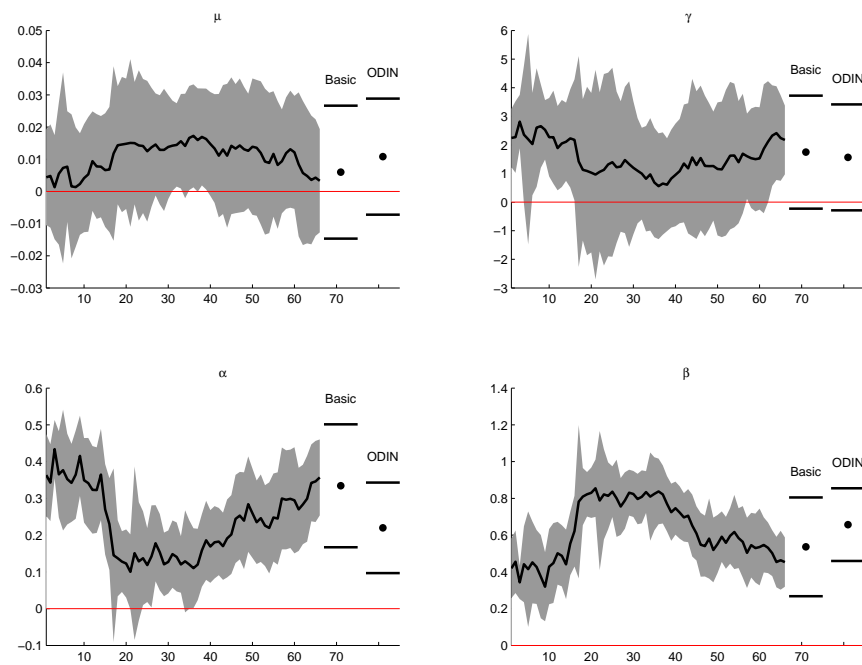


Figure 36: GARCH, Sample A

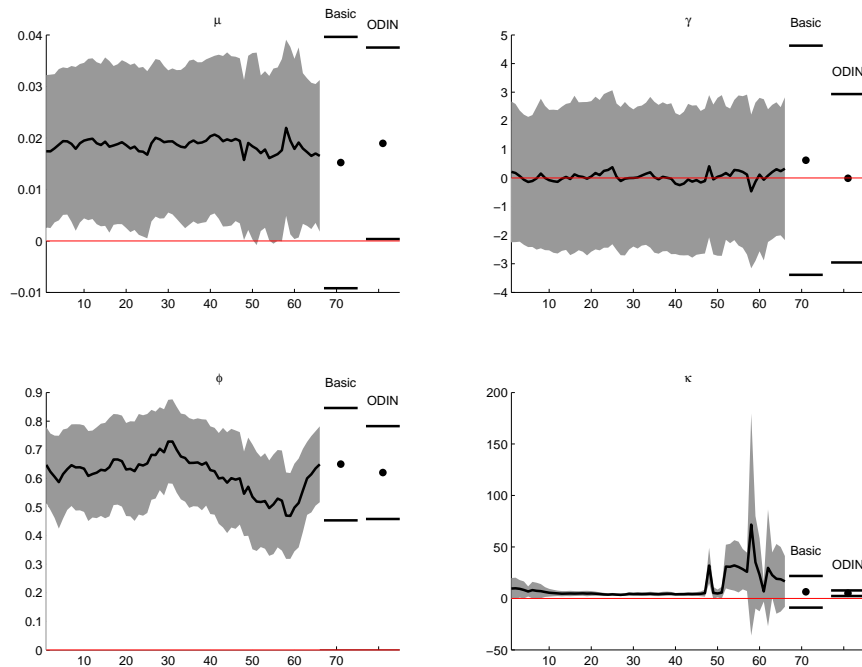


Figure 37: MIDAS, Sample A

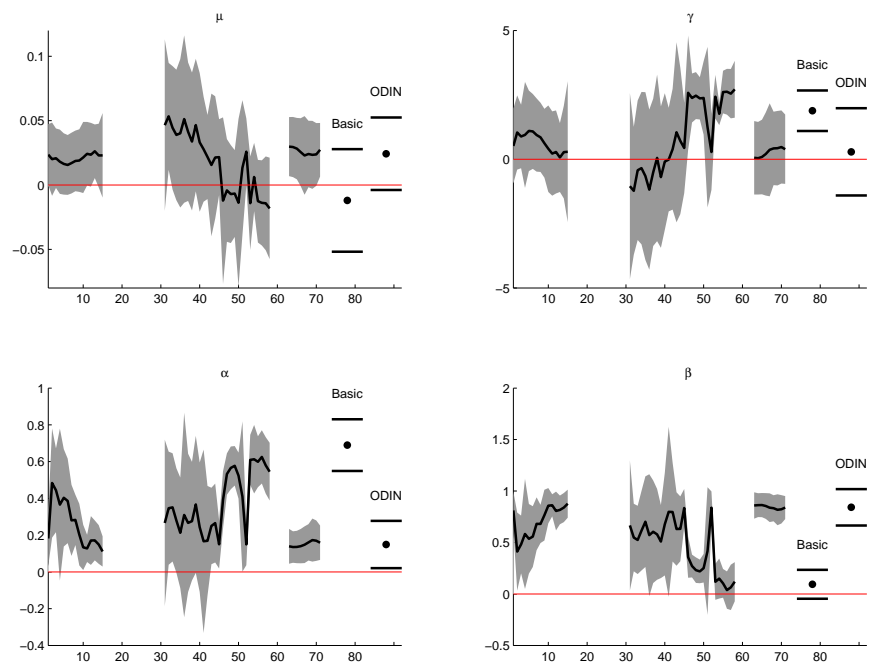


Figure 38: GARCH, Sample B

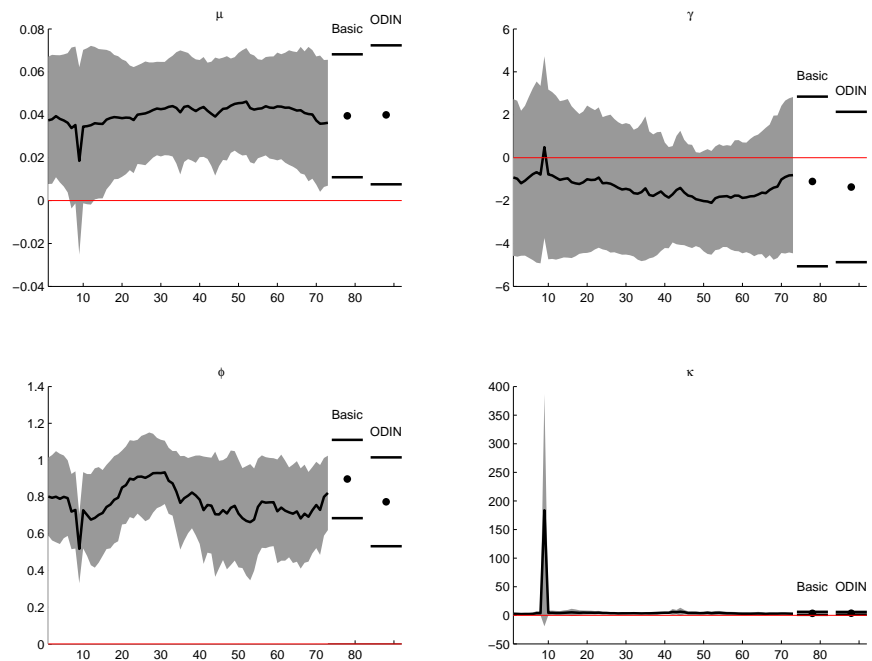


Figure 39: MIDAS, Sample B

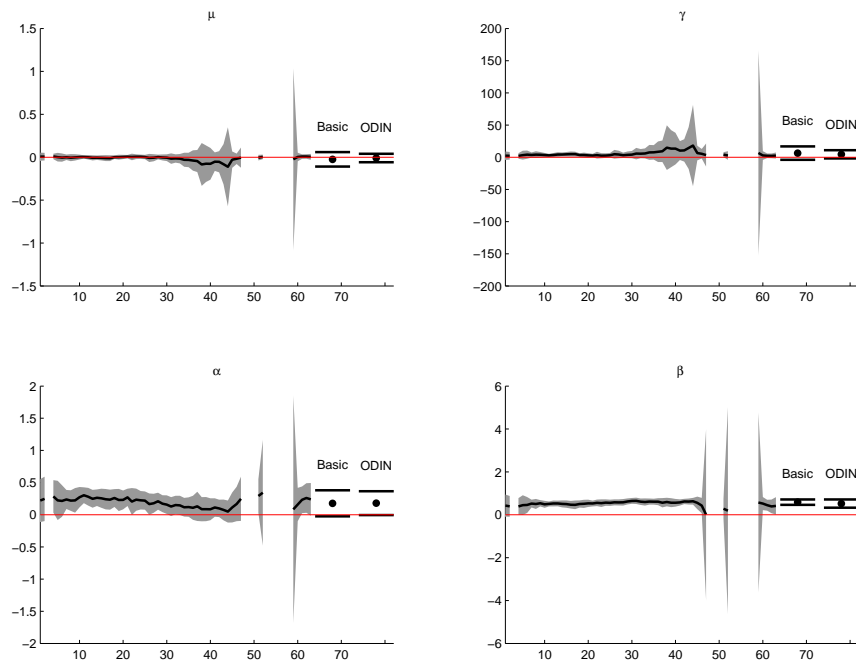


Figure 40: GARCH, Sample C

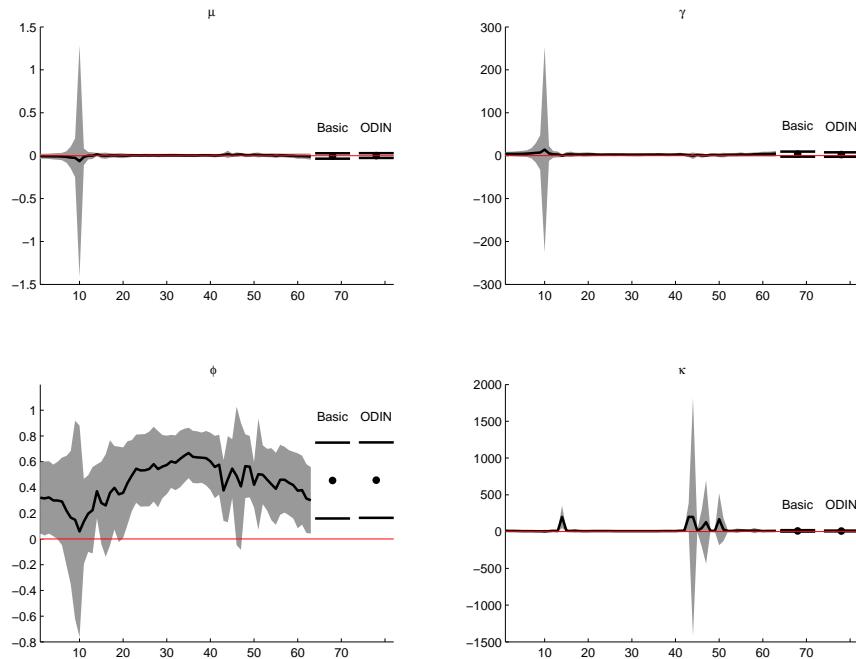


Figure 41: MIDAS, Sample C