

Multi-Period Forecasts of Volatility: Direct, Iterated, and Mixed-Data Approaches *

Eric Ghysels[†] Antonio Rubia[‡] Rossen Valkanov[§]

First Draft: May, 2007
This version: March 14, 2009

Abstract

Multi-period forecasts of stock market return volatilities are often used in asset pricing, portfolio allocation, risk-management and most other areas of finance where long-horizon measures of risk are necessary. Yet, very little is known about how to forecast volatility several periods ahead, as most of the focus has been on one-period-ahead forecasts. In this paper, we compare several approaches of producing multi-period ahead forecasts of volatility – iterated, direct, and mixed-data sampling (MIDAS) – as alternatives to the often-used “scaling” method. The comparison is conducted (pseudo) out-of-sample using returns data of the US stock market portfolio and a cross section of size, book-to-market, and industry portfolios. The results are surprisingly sharp. For the market and all other portfolios, we obtain the same ordering of the volatility forecasting methods. The direct approach provides the worse (in MSFE sense) forecasts; it is dominated even by the naive scaling method. Iterated forecasts are suitable for shorter horizons (5 to 10 days ahead), but their MSFEs deteriorate rapidly as the horizon increases. The MIDAS forecasts perform well at long horizons: they dominate all other approaches at horizons of 10-days ahead and longer. At 30-days ahead horizons, the MIDAS MSFE is about 20 percent lower than that of the next best volatility forecast. West (1996) and Giacomini and White (2006) tests show that the difference in predictive ability is statistically significant at conventional levels. In sum, this study dispels the notion that volatility is not forecastable at long horizons and offers an approach that delivers accurate out-of-sample predictions.

*We thank Yacine Aït-Sahalia, Andrew Patton, Allan Timmermann, as well as participants at the 2006 UNC Jackson Hole Finance Conference and the 2006 Risk Management conference, San Diego for helpful discussions. This project was initiated during Rubia’s 2006 visit to UCSD. All remaining errors are our own.

[†]Department of Finance, Kenan-Flagler Business School and Department of Economics, University of North Carolina, Gardner Hall CB 3305, Chapel Hill, NC 27599-3305, e-mail: eghysels@unc.edu.

[‡]University of Alicante, Campus de San Vicente, CP 03080, Spain. E-mail: antonio.rubia@ua.es.

[§]Corresponding Author. Rady School of Management, UCSD, Otterson Hall, 9500 Gilman Drive, MC 0093 La Jolla, CA 92093-0093, phone: (858) 534-0898, e-mail: rvalkanov@ucsd.edu.

– Preliminary: Comments Welcome –

Abstract

Multi-period forecasts of stock market return volatilities are often used in asset pricing, portfolio allocation, risk-management and most other areas of finance where long-horizon measures of risk are necessary. Yet, very little is known about how to forecast volatility several periods ahead, as most of the focus has been on one-period-ahead forecasts. In this paper, we compare several approaches of producing multi-period ahead forecasts of volatility – iterated, direct, and mixed-data sampling (MIDAS) – as alternatives to the often-used “scaling” method. The comparison is conducted (pseudo) out-of-sample using returns data of the US stock market portfolio and a cross section of size, book-to-market, and industry portfolios. The results are surprisingly sharp. For the market and all other portfolios, we obtain the same ordering of the volatility forecasting methods. The direct approach provides the worse (in MSFE sense) forecasts; it is dominated even by the naive scaling method. Iterated forecasts are suitable for shorter horizons (5 to 10 days ahead), but their MSFEs deteriorate rapidly as the horizon increases. The MIDAS forecasts perform well at long horizons: they dominate all other approaches at horizons of 10-days ahead and longer. At 30-days ahead horizons, the MIDAS MSFE is about 20 percent lower than that of the next best volatility forecast. West (1996) and Giacomini and White (2006) tests show that the difference in predictive ability is statistically significant at conventional levels. In sum, this study dispels the notion that volatility is not forecastable at long horizons and offers an approach that delivers accurate out-of-sample predictions.

Keywords: Volatility forecasting, multi-period forecasts, mixed-data sampling
JEL: G17, C53, C52, C22

1 Introduction

Financial decisions are often predicated on accurate multi-period-ahead forecasts of volatility. For instance, portfolio allocation, risk management, and regulator supervision often necessitate weekly, monthly, or quarterly volatility forecasts computed from return data available at, say, daily frequency. It is thus surprising that the extensive volatility literature has focused almost exclusively on the accuracy of one-period-ahead forecasts (Engle (1982), Bollerslev (1986), Andersen and Bollerslev (1998a), Hansen and Lunde (2005)) whereas long-horizon volatility forecasts have received much less attention. The dominant long-horizon volatility forecasting approach is still to scale the one-period-ahead forecasts by \sqrt{k} where k is the horizon of interest. Its popularity among practitioners stems mostly from its use in Riskmetrics.¹ While there are several alternative approaches to compute multi-period volatility forecasts, the common belief is that, in general, volatility is difficult to forecast at horizons longer than ten days or so (Christoffersen and Diebold (2000), West and Cho (1995)). This paper undertakes a comprehensive empirical examination of multi-period volatility forecasting approaches, beyond the simple \sqrt{k} -scaling rule. The perspective that we offer is markedly more optimistic: long-horizon volatility is much more forecastable than previously suggested at horizons as long as 60 trading days (about three months).

Long horizon volatility forecasts can be constructed in three fundamentally different ways. The first approach is to estimate a horizon-specific model of the volatility, such as a weekly, monthly, or quarterly GARCH, which can then be used to form direct predictions of volatility over the next week, month, or quarter. The second approach is to estimate a daily autoregressive volatility forecasting model and then iterate over the daily forecasts for the necessary number of periods to obtain weekly, monthly, or quarterly predictions of the volatility. The forecasting literature refers to the first approach as “direct” and the second as “iterated” (Marcellino, Stock, and Watson (2006)). A third method is the mixed-data sampling (MIDAS) approach introduced by Ghysels, Santa-Clara, and Valkanov ((2005), (2006)). A MIDAS model uses daily squared returns to produce directly multi-period volatility forecasts and can be viewed as a middle ground between the direct and the iterated approaches. These three methods have been extensively used in the empirical finance literature, yet little is known about their relative performance in the context of multi-period volatility forecasts.

¹See J.P.Morgan/Reuters (1996) Technical Report (pp. 84).

A systematic comparison of direct, iterated, and MIDAS multi-period volatility forecasts has, to our knowledge, not been carried out. A few notable exceptions are Diebold, Hickman, Inoue, and Schuermann (1997), Christoffersen and Diebold (2000), and Andersen, Bollerslev, and Lange (1999) but these studies are more limited in scope. Moreover, they do not consider MIDAS methods which, to preview the results, are particularly suitable for long-horizon volatility forecasting. Marcellino, Stock, and Watson (2006) compare direct and iterated forecasts, but their study focuses on the level of US macroeconomic data series, whereas our paper is about forecasts of volatility of asset returns. Also, Marcellino, Stock, and Watson (2006) do not investigate MIDAS models.

Perhaps a reason for the lack of papers on the subject is the theoretical difficulty of comparing multi-period forecasts, which can be summarized as follows. At a theoretical level, the trade-off between bias and estimation volatility that exists in multi-period forecasts has not been fully understood (Findley (1983), Findley (1985), Lin and Granger (1994), Clements and Hendry (1996), Bhanzali (1999), and Chevillon and Hendry (2005)). While the above cited papers do not consider volatility predictions per se, the general conclusion is that direct forecasts ought to dominate iterated forecasts, because of model uncertainty. In the realistic case of misspecification in the one-period model (model uncertainty), the direct method is more robust to biases arising from misspecification. The iterated model would dominate only if the one-period model is known with certainty (no bias) and we are only concerned with estimation uncertainty (efficiency).

Moreover, to assess the accuracy of the volatility forecasts, we need a loss function that penalizes deviations from the ex-post realization of the volatility. It is well known that the loss function plays an important role in forecast comparisons (Elliott and Timmermann (2008) and references therein). In our case, there is an additional complicating factor in assessing the forecast accuracy: the true volatility is not observable, even ex-post. We follow French, Schwert, and Stambaugh (1987a) and Andersen and Bollerslev (1998b) and compute realized volatility as a proxy for the true volatility, which is then used in the loss function. Because of the necessity to use a proxy, we need to make sure that our loss function is consistent, i.e. that it delivers the same forecast ranking with the proxy as it would with the true volatility. Patton (2007) shows that a loss function that has such a consistency property is the mean square forecasting error (MSFE) whereas loss functions such as mean absolute forecasting error would not be appropriate. Therefore, relying on Patton's (2007) results, we use the MSFE as the loss function in this study.

To correctly rank multi-period volatility forecasts, we need a test for predictive accuracy. Diebold and Mariano (1995) proposed one such test which was simple and, although failing to account for parameter estimation error, it gave impetus for further research on the topic. Estimation error is of particular concern in the volatility forecasting literature as there is no lack of competing predictive models. The approaches by Ghysels and Hall (1990), Hoffman and Pagan (1989) and West (1996) address explicitly parameter uncertainty. Therefore, we use West (1996) as one of the two tests in our forecast comparisons.

The second test we use was proposed by Giacomini and White (2006) and can be viewed as a generalization, or a conditional version of West’s (1996) test. Rather than comparing the difference in average performance, Giacomini and White (2006) consider the conditional expectation of the difference across forecasting models. This conditioning approach allows not only for parameter uncertainty (as in West (1996)) but also uncertainty in a number of implicit choices made by the researcher when formulating a forecast, such as what data to use, the windows of in-sample estimation period, the length of the out-of-sample forecast, among others. Since our volatility forecast comparisons would involve models that use data at different frequencies, different methods of constructing multi-period forecasts, and different forecasting horizons, the Giacomini and White (2006) test would be particularly appropriate. Hence, by using this procedure we can determine not only the best forecasting model, but rather the best *forecasting method*, given the sample series and the forecasting horizon we consider. This appears to be the most appropriate test for our purpose.

By addressing the above problems, we make the following contribution with respect to the previous volatility forecasting literature. First, we investigate whether multi-horizon forecasts of the volatility of US stock market returns are more accurate than the naive but widely-used scaling approach. While it might seem obvious that the well-documented predictability of volatility using one-period-ahead forecasts implies multi-period predictability, this is not necessarily the case.² In fact, Diebold, Hickman, Inoue, and Schuermann (1997) and Christoffersen and Diebold (2000) provide evidence that the opposite might be true in forecasting return volatility. We consider volatility forecasts of the US market portfolio returns as well as of five size, five book-to-market, and ten industry portfolio returns.

Second, we carry out an empirical comparison of the various multi-period forecasting

²Model uncertainty, parameter uncertainty and model instability are some of the reasons that might drive a wedge between one-period and multi-period forecasts.

approaches – direct, iterated, and MIDAS – using the same twenty one stock portfolio returns (market plus twenty size, book-to-market, and industry). Because of the lack of theoretical guidance on this topic, such a comparison would not only provide some stylized facts about the long-horizon forecastability of return volatility, but it would also allow us to gauge if one method produces clearly superior multi-horizon forecasts relative to the others. The results from such a data-driven comparison are ultimately conditional on the sample at hand and the design of the pseudo out-of-sample experiment. However, in our case, the findings (discussed below) are remarkably sharp. At the very least, they speak to the method that should be used in multi-period volatility forecasts. More generally, our results might provide guidance for future theoretical work on long-horizon volatility forecasting.

As a third contribution, building on recent work by Ghysels, Santa-Clara, and Valkanov (2006), we analyze more closely the performance of the MIDAS volatility forecasts. The MIDAS approach offers a natural middle-ground between the direct and iterated approaches. Indeed in a MIDAS, daily forecasting variables (for instance, squared returns) are used to produce a direct prediction of the long-horizon (proxy of) volatility. Similarly to an iterated approach, the MIDAS forecasts use information contained in the entire history of daily returns, which implies that they will be efficient. At the same time, the forecasted variable is the long-horizon volatility, which allows us to side-step the need of aggregating the forecasts and introducing bias. We analyze several parameterizations of MIDAS forecasting models.

Our study yields surprisingly sharp results. First, as expected, the scaling-up method performs poorly relative to the other methods across portfolios and horizons. This result is consistent with the findings of Diebold, Hickman, Inoue, and Schuermann (1997) and several others who have documented the poor performance of this approach. What is surprising, however, is that the direct method does not fair much better. At horizon longer than 10 days ahead, the scaling performs significantly better than the direct approach. At short and medium horizons (up to 10-days ahead), they are similar. Hence, if the direct method were the only alternative to the scaling approach, and since scaling is a poor forecaster of future volatility, one might come to the erroneous conclusion that the volatility is hard to forecast at long horizons.

While there is relatively more work on multi-horizon forecasting in linear models, it is clear that our paper reveals the important differences between conditional mean and conditional variance forecasts. Take for example the comparison between the \sqrt{k} -scaling rule and the direct method. As noted before the fact that the latter performs poorly at longer horizon in

comparison to the former is surprising in light of the conditional mean forecasting results of Marcellino, Stock, and Watson (2006). While the issue of one-step ahead model specification error is important for both conditional mean and variance forecasts, the case of volatility has two additional complications: (i) specification errors of short horizon models are less important, and (ii) the sampling frequency of returns in the information set plays a dominant role as well. Regarding (i), we know from Foster and Nelson (1994) that correct volatility forecasts can be made with the 'wrong' model. Regarding (ii), it is important to note that the \sqrt{k} -scaling rule relies on daily returns, whereas the direct method relies on long horizon past returns. Since the work of Merton (1980) we know that sampling frequency is a dominant factor in volatility measurement and forecasting.

Our second result dispels that notion. We find that for the volatility of the market portfolio, iterated and MIDAS forecasts perform significantly better than the scaling and the direct approaches. At relatively short horizons of 5- to 10-days ahead, the iterated forecasts are quite accurate. However, at horizons of 10 days ahead and higher, MIDAS forecasts have a significantly lower MSFE relative to the other forecasts. At horizons of 30- and 60-days ahead, the MSFE of MIDAS is more than 20 percent lower than that of the next best forecast. These differences are statistically significant at the one percent level according to the West (1996) and Giacomini and White (2006) tests. Hence, we find that suitable MIDAS models produce multi-period volatility forecasts that are significantly better than other widely used methods.

Third, the superior performance of MIDAS in multi-period forecasts is also observed in predicting the volatility of the size, book-to-market, and industry portfolios. Similarly to the market volatility results, the relative precision of the MIDAS forecasts improves with the horizon. At horizons of 10-periods and higher, the MIDAS forecasts of eight out of the ten size and book-to-market portfolios dominate the iterated and direct approaches. At horizons of 30-periods and higher, the MIDAS has the smallest MSFEs amongst all forecasting methods for all ten portfolios. We observe that the volatility of the size and book-to-market portfolios is significantly less predictable than that of the entire market. Also, the predictability of the volatility increases with the size of the portfolio. The volatility of the largest-cap stocks is the most predictable, albeit still less forecastable than the market's. We do not observe such a discernable pattern for the book-to-market portfolios.

The paper is organized as follows. In section 2, we introduce the direct, iterated, and MIDAS multi-period forecasts. The third section discusses the loss function and the West (1996) and

Giacomini and White (2006) test used to evaluate the forecasts at various horizons. Section 4 presents the empirical results. In section five, we conclude by offering directions for further research.

2 Multi-Period Volatility Forecasts

We use the following notation. Daily returns are indexed by d where $d = 1, 2, \dots, D$ and long-horizon returns, at horizon of k days, are indexed by $t = 1, 2, \dots, T_k$, where $T_k = \lfloor D/k \rfloor$ and $\lfloor \cdot \rfloor$ is the integer operator. For instance, in our data set, we have $D = 11202$ observations from 1963 to 2007 from which we can compute $T_5 = 2184$ 5-day (or weekly) non-overlapping returns, $T_{10} = 1092$ 10-day (or bi-weekly) non-overlapping returns, and so on. To keep the notation simple, we will henceforth drop the subscript from T but will keep in mind that the number of non-overlapping returns changes with the horizon of interest, k .

The daily return is defined as $r_d = \log(P_{d+1}/P_d)$ and the k -period, continuously compounded return is defined as $R_t^k = \log(P_{d+k}/P_d)$. Note that we use lower cases for daily returns and capitals for multi-period returns, indexed by the horizon k . All long-horizon returns are demeaned and are computed without overlap, to avoid mechanical serial correlation. The information set of daily returns at time t is $I_t = \{r_t, r_{t-1}, r_{t-2}, \dots, r_0\}$. Analogously, $I_t^k = \{R_t^k, R_{t-1}^k, R_{t-2}^k, \dots, R_0^k\}$ is the information set of the k -period, non-overlapping, continuously compounded returns. We denote by I_T and I_T^k the information sets based on the history of the entire one-period and k -period returns, respectively. Note that the two information sets are different: I_T^k contains T non-overlapping k -period returns whereas I_T contains $D = Tk$ daily returns, and $I_T^k \subset I_T$.

The various volatility forecasts will be denoted by $V_c(a, b, i)$, where: (i) c is the forecasting method, either *direct* (d), *iterated* (i), or *MIDAS* (m); (ii) a is the starting period of the forecast; (iii) b is the forecast horizon; (iv) i is the information set used. Often we will drop i , or even (a, b, i) in situations where it will be unambiguous. For example, $V_I(t, 1_k)$ versus $V_M(t, 1_k)$ are conditional forecasts, both using daily historical data I_t , to produce k -step ahead forecasts at time t with iterated and MIDAS methods. Finally, we will denote $V_P(t, 1_k)$ as the true - or population - conditional volatility given past daily data.

2.1 Direct Volatility Forecasts

The first method, perhaps the simplest to implement, is to use the multi-period returns R_t^k , and forecast the multiple horizon conditional volatility directly as one step ahead. For instance, we can model $V_D(T, 1_k)$ as a GARCH(p,q) that we estimate with k -period returns in I_T^k and then forecast the next k -period volatility. We call this a direct approach of forecasting and denote it by $V_D(T, 1_k, I_T^k)$, or more concisely $V_D(T, 1_k)$. One might expect this approach to yield accurate estimates on several grounds. First, the parsimony of the GARCH model makes it hard to beat in pseudo out-of-sample forecasts (Hansen and Lunde (2005)). Second, the direct approach would produce robust estimates in the sense that it does not display a bias. However, given that we use the multi-period returns R_t^k to formulate volatility forecasts, this estimator would not be as efficient as one using the information in set I_T .

In our comparison, we use a GARCH(1,1) model, or to more precise an AR(1)-GARCH(1,1) model to forecast volatility, both in the direct and the iterated approach. We also have results from more general GARCH(p,q) models, where p and q are chosen by the Akaike Information Criterion (AIC) and Bayes Information Criterion (BIC). However, the AIC- and BIC-chosen models fail to beat the GARCH(1,1) out-of-sample. This finding confirms that the Hansen and Lunde (2005) results hold at horizons longer than one-period ahead. Henceforth, we use the GARCH(1,1) exclusively in our analysis.

2.2 Iterated Volatility Forecasts

The second method is to use the daily returns r_t in I_T in estimating forecasts $V_I(T, 1_k)$. Hence, we form iterated forecasts of the daily volatility k period forward and, under the assumption that the conditional covariances are zero, we write that $V_I(T, 1_k, I_T) = \sum_{j=1}^k V_D(T+j, 1, I_T)$. Note that this forecast uses information at time T and the forecasts for days $T+1$, $T+2$, ..., $T+k$ would have to be iterated from the one-period daily forecasting model.

This iterated approach, seems viable because returns are serially uncorrelated (or close), but their volatilities are time-varying and persistent. Hence, it is an improvement over the simple scaling approach. This method has the advantage that we are using daily data to estimate the forecasting model and will hence be more efficient than the direct approach. However, since we are iterating the forecasts and summing them, then small errors due to

model misspecification will be amplified. Hence, in general this method is thought, at least theoretically to be bias-prone. But it will be efficient, because the data used is high-frequency. This is particularly important in volatility models.

2.3 MIDAS Volatility Forecasts

The third approach is to use the daily returns r_t and directly produce a multi-step ahead forecast using a mixed-data sampling (or MIDAS) approach. Since this approach is relatively new, we describe it more in detail. We start by formulating a MIDAS forecasting regression:

$$\tilde{V}_{t+1}^k = \mu_k + \phi_k \sum_{j=0}^{j^{max}} b_k(j, \theta) r_{t-j}^2 + \varepsilon_{k,t} \quad (2.1)$$

where \tilde{V}_{t+1}^k is a measure of (future) volatility such as realized volatility, e.g. $\tilde{V}_{t+1}^k = RV_{t+1}^k \equiv \sum_{j=1}^k r_{t+j}^2$ and $b_k(j, \theta)$ is a parsimonious weighting function parameterized by a low-dimensional parameter vector θ . The intercept μ_k , slope ϕ_k , and weighting scheme parameters θ are estimated with QMLE. The regression involves data sampled at different frequencies, since in this study, the realized volatility in equation (2.1) is measured at horizons ranging from one day ($k = 1$) to three months ($k = 60$), whereas the regressors are available at daily frequencies. For instance, equation (2.1) relates the realized volatility over the month of, say, December (measured from the close of the market during the last trading day of November to the close of the market during the last trading day of December) with daily squared returns up to the last day of November. The weights placed on the predictive lagged squared returns are estimated in-sample and used to form an pseudo out-of-sample forecast.

As noted before, the lag coefficients $b_k(j, \theta)$ are parameterized to be a low-dimensional function of underlying parameters θ . Without this parametric restriction, the number of parameters associated with the forecasters r_{t-j}^2 would proliferate significantly, leading to in-sample overfit and poor out-of-sample forecasts. A suitable parameterization of $b_k(j, \theta)$ circumvents the problem of parameter proliferation and is one of the most important ingredients in a MIDAS regression. We consider several parameterizations of $b_k(j, \theta)$, some of which have already been suggested in previous work. Specifically we consider the following five specifications: We postulate a flexible form for the weight given to the squared return on day $t - d$:

1. Exponential:

$$b_k(j, \theta_1, \theta_2) = \frac{\exp\{\theta_1 j + \theta_2 j^2\}}{\sum_{i=0}^{\infty} \exp\{\theta_1 i + \theta_2 i^2\}}. \quad (2.2)$$

This scheme guarantees that the weights are positive (which in turn ensures that the forecasted volatility is also positive) and that they add up to one. Also, the functional form in equation (2.2) can produce a wide variety of shapes for different values of the two parameters, and it is parsimonious, with only two parameters to estimate. Finally, as long as the coefficient κ_2 is negative, the weights go to zero as the lag length increases. The speed with which the weights decay controls the effective number of observations used to estimate the conditional volatility.

2. Beta:

$$b_k(j, \theta_1, \theta_2) = \frac{f(\frac{j}{j^{\max}}, \theta_1; \theta_2)}{\sum_{i=1}^{j^{\max}} f(\frac{i}{j^{\max}}, \theta_1; \theta_2)} \quad (2.3)$$

where: $f(z, a, b) = z^{a-1}(1-z)^{b-1}/\beta(a, b)$ and $\beta(a, b)$ is based on the Gamma function, or $\beta(a, b) = \Gamma(a)\Gamma(b)/\Gamma(a+b)$. Specification (2.3) was introduced in Ghysels, Santa-Clara, and Valkanov (2002) and further explored in Ghysels, Sinko, and Valkanov (2006). One appealing feature is positivity of the coefficients, which is necessary for *a.s.* positive definiteness of the forecasted volatility. For $\theta_1 = 1$ and $\theta_2 > 1$ one has a slowly decaying pattern typical of volatility filters, which means that only one parameter is left to determine the shape, whereas in the case of $\theta_1 = \theta_2 = 1$ we obtain equal weights, which corresponds to a rolling estimator of the volatility (French, Schwert, and Stambaugh (1987a), (2005)). The flexibility of the Beta function is well known and it is often used in Bayesian econometrics to impose flexible, yet parsimonious prior distributions. The function can take many shapes, including flat weights, gradually declining weights as well as hump-shaped patterns.

3. Linear:

$$b_k(j) = 1/j^{\max} \quad (2.4)$$

where j^{\max} is the truncation point specified above. This simple decay functional form has the advantage that there are no parameters to estimate in the lagged weight function and might offer good out-of-sample forecasts.

4. Hyperbolic:

$$b_k(j, \theta) = \frac{g(\frac{j}{j^{max}}, \theta)}{\sum_{i=1}^{j^{max}} g(\frac{i}{j^{max}}, \theta)} \quad (2.5)$$

where $g(j, \theta) = \Gamma(j+\theta) / (\Gamma(j+1)\Gamma(\theta))$ which can be written equivalently as $g(0, \theta) = 1$ and $g(j, \theta) = (j + \theta - 1) g(j-1, \theta) / j$, for $j \geq 1$. The above Gamma functional form has only one parameter to estimate. While it is not as flexible as the Beta specification, it has been extensively used in the volatility modeling literature particularly in the context of ARFIMA long memory specifications (see e.g. Campbell, Lo, and MacKinlay (1997) and Andersen and Bollerslev (1998a)). The weights in (2.5) decay hyperbolically rather than exponentially (Hosking (1981)). The weights are the normalized (or proportional to the) impulse response of a truncated ARFIMA model. They are the impulse response of a ARFIMA model (see, Hosking (1981) and Tanaka (1999) and references therein).

5. Geometric:

$$b_k(j, \theta) = \frac{\theta^j}{\sum_{i=0}^{\infty} \theta^i}. \quad (2.6)$$

where $|\theta| \leq 1$, and $b_k(j, \theta)$ are normalized so that they sum up to one as in the previous specifications. This specification is almost identical to a GARCH(1,1) model.

We also estimate two restricted versions, denoted by “Exp. Rest.” and “Beta Rest.,” with $\theta_2 = 0$ (Exp. Rest.) and $\theta_1 = 1$ (Beta Rest.). These restrictions are to ensure a slowly decaying pattern of the weighting functions. Moreover, the linear scheme is the simplest MIDAS-type model that sets all weights equal to each other. It therefore reduces to the rolling-window estimator in the spirit of French, Schwert, and Stambaugh (1987a). The interest on this “naive” specification is that we can address whether more sophisticated specifications are able to yield better out-of-sample estimates. In other words, we can address whether the forecasting power of the MIDAS approach comes from simply using a rolling window (i.e., a conditional estimator), or also from using a weighting function that acknowledges that remote observations may be less important to forecast future volatility. Obviously, in this comparison, the forecasting horizon is expected to play an important role. Our statistical analysis will reveal that, at short-horizons, any of the MIDAS filter outperforms easily the linear specification.

The MIDAS forecasts are denoted by $V_M(T, 1_k, I_T)$, or $V_M(T, 1_k)$. Whenever necessary, we

will specify the weights used in the forecasts, which implicitly capture the dynamics of the conditional volatility. Larger weights on distant past returns induce more persistence on the volatility process. The weighting function also determines the statistical precision of the estimator by controlling the amount of data used to estimate the conditional volatility. There is a tension between capturing the dynamics of volatility and minimizing measurement error. Because volatilities change through time, we would like to use more recent observations to forecast the level of volatility in the next month. However, to the extent that measuring volatility precisely requires a large number of daily observations, the estimator could still place significant weight on more distant observations.

Some of the above weight specifications have been used in the previous literature, while others are new to the MIDAS framework but not to the volatility forecasting literature. The exponential lag structure has been suggested by Ghysels, Santa-Clara, and Valkanov (2005) to study the risk return tradeoff, while the beta lag has been used in Ghysels, Santa-Clara, and Valkanov (2006) in comparing short-horizon forecasts using different predictors. The linear lag is a simple natural benchmark, with no parameters to estimate. Hence, it may prove robust out-of-sample. The hyperbolic weights are similar to the impulse responses of ARFIMA models which have been successfully used in the volatility literature (see e.g. Andersen, Bollerslev, and Diebold (2003)). The geometric weights provide a MIDAS model that is the closest and easiest to compare to a GARCH(1,1) model. The weights can also be specified as a step function with a predetermined number and relative magnitude of the steps. The step function specification presents the most data-mining problems, but is also the simplest to implement. Corsi (2004) and Forsberg and Ghysels (2006) use this specification to model the volatility of stock returns.

We can think of the mixed-data regression (2.1) as combining the attractive features of the iterated and direct forecasts. Notice that we can vary the forecast horizon by changing k , whereas the predictive variables remain the same and allow us to explore the richer information set I_T . This is not true for the direct approach, where the predictive variables change with the horizon and estimation and forecasts are formed using information set $I_T^k \subset I_T$. In the MIDAS forecasts, it is not the regressors that change but the estimated shape of the lag function b_k , thus changing the weights placed on the lagged daily squared returns. Moreover, we form direct forecasts of future volatility at the horizon of interest without having to iterate over forecasts. This is in contrast with the iterated GARCH forecasts. Therefore, we use (2.1) to side-step the iteration and aggregation issues associated

with iterated forecasts as well as the inefficient use of lagged returns that is characteristic of the direct approach.

While the MIDAS approach to formulate forecasts is quite general, we focused on regression (2.1) for several reasons. First, we could have extended the number of regression to include not only daily squared returns but other volatility forecasts such as daily absolute returns, daily range measures (high-low), and others, as done in Forsberg and Ghysels (2006) and Ghysels, Santa-Clara, and Valkanov (2006). We use daily squared returns only in order to make this forecast as directly comparable with the GARCH forecasts as possible. Moreover, the comparison of the squared daily return predictors to other predictors at shorter horizons have already been investigated extensively in Forsberg and Ghysels (2006) and Ghysels, Santa-Clara, and Valkanov (2006). MIDAS regressions typically do *not* exploit an autoregressive scheme, so that r_{t-j}^2 is not necessarily related to lags of the left hand side variable. Instead, MIDAS regressions are first and foremost regressions and therefore the selection of r_{t-j}^2 amounts to choosing the best predictor of future volatility from the set of several possible measures of past fluctuations in returns. In other words, MIDAS is a reduced-form forecasting device rather than a model of conditional volatility.

2.4 Other Forecasting Approaches: Scaling and Integrated

There are many other approaches, some of which we have tried. In this section, we discuss these approaches. The first one is the scaling approach. It involves estimating a daily volatility forecasting model up to time t , using it to form a forecast of the volatility at day $t + 1$ and scaling this forecast by \sqrt{k} , where k is the length of the horizon of interest. This forecasting method assumes that log returns are i.i.d. It has been documented of not being appropriate in forecasting long-horizon volatility by Christoffersen and Diebold (2000), Diebold, Hickman, Inoue, and Schuermann (1997), among others. However, its prominence in applied work is undisputed, largely due to J.P.Morgan/Reuters's (1996) widely adopted Riskmetrics approach and the suggestions in the Basel II agreement.

An alternative approach is to compute volatility forecasts using an AR(p) model based on passed realized volatility estimates. Such an AR(p) approach has been used by French, Schwert, and Stambaugh (1987b), Schwert (1989) with daily returns and more recently by Andersen and Bollerslev (1998a), Andersen, Bollerslev, Diebold, and Labys (2001), and others with higher frequency data. Based on these papers, we have also considered the

following specifications. First, we have estimated a general AR(p) model of realized volatility, where the lag selection is done in sample using the AIC and BIC. From an out-of-sample MSE criterion perspective, the selected AR(p) model fails to significantly outperform the AR(1) model for most assets. Second, Andersen, Bollerslev, Diebold, and Labys (2001) and Andersen, Bollerslev, and Diebold (2003) provide compelling evidence that realized volatility has long memory. However, those papers are conducted with high frequency, 5-minute returns to forecast daily volatility. Since we have only daily returns, we fail to find long memory and ARFIMA(p,d,0) dynamics are rejected in favor of a simple AR(p) model in-sample. Also, out-of-sample forecasts of the best ARFIMA(p,d,0) model are dominated by those of the AR(1). These results are reminiscent of those in Hansen and Lunde (2005).

After a rather extensive preliminary search, we use the AR(1) model. We compute realized volatility at horizon k to directly forecast the realized volatility over the horizon of interest. From that perspective, this is a direct approach. It can also be viewed as a restricted version of a MIDAS model, where the weights on lagged returns are not estimated but are rather equally weighted.

Another natural alternative is to use the unconditional volatility, either computed from daily or k -period returns, to forecast as a forecast of future volatility.

3 Comparing the Forecasts

Once we have the forecasts $V_D(T, 1_k)$, $V_I(T, 1_k)$ and $V_M(T, 1_k)$, we need to decide which one is the closest to the true volatility. To answer this question, we need to tackle three related issues. First, since the true k -period volatility is unobservable, we will need to proxy for it. Second, in evaluating the forecasts, we have to agree on an appropriate loss function. Given the first issue, we require a loss function that produces robust rankings of the forecasts even in the absence of true volatility. Finally, to gauge the statistical significance of the predictive gains or losses, we need a test that takes into account the uncertainty involved in producing the forecasts. We address these issues below.

3.1 Proxy for Unobservable Long-Horizon Volatility

The pseudo out-of-sample forecast error is

$$e_{F,T+1}^k \equiv V_P(T, 1_k) - V_F(T, 1_k)$$

where $V_F(T, 1_k)$ is the forecasted volatility (either $V_D(T, 1_k)$, $V_I(T, 1_k)$, or $V_M(T, 1_k)$,) and $V_P(T, 1_k)$ is the true population volatility. The forecasting error $e_{F,T+1}^k$ is indexed by the forecasting method (subscript) and by the horizon (superscript). However, we cannot obtain $e_{F,T+1}^k$ because the true volatility is unobservable. Hence, we use the realized volatility RV_{T+1}^k as a proxy for V_P . Andersen and Bollerslev (1998a) and subsequent work show that the realized volatility is a good proxy for the true volatility, or at least much better than squared returns. The realized volatility is computed using high-frequency returns. The consensus in that literature is that we need data that is very frequent. Unfortunately, we do not have access to high-frequency data for our sample period, nor for the size, book-to-market, and industry portfolios in the cross-section. Hence, we use the highest frequency data that is available to us – daily returns – to compute the realized volatility at horizon k . Given that we don't have high frequency data, the estimated RV_{T+1}^k will be a noisy proxy of the true underlying volatility. We have to keep that in mind when ranking the forecasts which makes choosing the appropriate loss function that much more important. We turn to that issue next.

3.2 Ranking the Forecasts: Appropriate Loss Function

Using RV_{T+1}^k , we compute the feasible out-of-sample forecast error

$$u_{F,T+1}^k = RV_{T+1}^k - V_F(T, 1_k, i)$$

and the sample MSFE at the k -horizon:

$$MSFE_F^k = \frac{1}{T_2 - T_1 + 1} \sum_{t=T_1}^{T_2} (u_{F,t}^k)^2.$$

The sample MSFE is computed for each forecasting horizon k and for each forecasting method F . For a horizon k , the empirical efficiency of the forecasts is assessed by comparing

the respective MSFEs. The ranking that we thus obtain will be consistent in the sense of Patton (2007). He showed that when we used the MSFE function, a forecast that dominates using the feasible errors u_{T+1}^k will also dominate using the infeasible e_{t+1}^k . In other words, the error introduced from using a proxy rather than the true volatility will not change the ranking of our forecasting methods. This robustness property is not shared by some other popular forecasting evaluation methods, such as the mean absolute forecasting error (MAFE). Hence, we focus on the MSFE as a loss function to evaluate the forecasts.

To summarize, since all forecasts are compared against the same volatility proxy and the MSFE is a consistent loss function, our conclusions should not be affected by the noisy measurement of volatility.

3.3 Tests for Predictive Ability

A test for predictive ability of two competing forecasting methods, $F1$ and $F2$, can be formulated as the following null hypothesis:

$$H_0 : E[(u_{F1,t}^k)^2 - (u_{F2,t}^k)^2] = 0 \quad (3.7)$$

While this setup was first proposed by Diebold and Mariano (1995), their testing procedure was valid under the assumption that the parameters of the model are known. West (1996) derived an asymptotic test of the above null under parameter uncertainty. Since parameter uncertainty is of great concern in volatility forecasting, we use the West (1996) as one of the two tests in our forecast comparisons.

An alternative hypothesis of predictability was proposed by Giacomini and White (2006), who argue that a more practically relevant hypothesis is:

$$H_0 : E[(u_{F1,t}^k)^2 - (u_{F2,t}^k)^2 | I_{t-1}] = 0 \quad (3.8)$$

where I_{t-1} is the information set available at time $t - 1$. The rationale for this test of conditional predictive ability is that the unconditional test of West (1996), while accounting for parameter uncertainty, fails to capture more general forms of model uncertainty. To see that, notice that the unconditional test for (3.7) will tend to choose a forecast based on a correctly specified model. However, as argued by Giacomini and White (2006), “even a

model that well approximates the data-generating process may forecast poorly, for example in the case that its parameters are imprecisely estimated.” In our multi-horizon volatility forecasting setup, the focus is squarely on volatility forecasting. Since our goal is to choose the most accurate forecasting method, the uncertainty is not only in the parameters but also in the model that we need to choose. For instance, we will be comparing iterated versus direct, versus MIDAS approaches. Also, we will be comparing the forecasts at various horizons. Hence, the conditional test of predictive ability of Giacomini and White (2006) is quite appropriate in our application and we will use it in conjunction with the West (1996) test.

In the empirical section, we will use both the West (1996) and Giacomini and White (2006) tests, because the difference in these tests hinges on our view of the source of uncertainty. We can see that by noticing that these are tests of two different null hypotheses. We cannot ultimately say which test would be more appropriate in our application (i.e., accurate size and higher power). Moreover, we don’t want our results to be predicated on the choice of the test. Hence, the statistical significance will be computed under both (3.7) and (3.8).

4 Data and Results

We present in this section the empirical results, starting with a data description.

4.1 Data

We have daily CRSP log returns for the period July 1, 1963 to December 31, 2007. Using these returns, we compute k -period continuously compounded, non-overlapping returns R_t^k , for $k = 1$ (daily), $k = 5$ (weekly), 10 (bi-weekly), 15, 20 (monthly), 25, 30, 60 (quarterly). We also have data of five daily size, five daily book-to-market, and ten industry portfolio returns obtained from Kenneth French’s website. In sum, we will forecast the volatilities of 21 portfolio returns (market plus 20 portfolios) at various horizons. The dataset is standard in empirical finance and, in the interest of conciseness, we do not provide returns summary statistics³.

The log daily returns are used to estimate the GARCH and MIDAS forecasting models.

³They are available upon request.

They are also used to compute the realized volatility RV_t^k for each horizon k as a proxy for the true (unobservable) volatility. The long-horizon returns R_t^k are used in the direct GARCH forecasts.

4.2 Results

Table 1 provides the summary statistics of the multi-period forecasts for the various models. All forecasts are out-of-sample as described above. We consider seven MIDAS models, an iterated GARCH (1,1), a direct GARCH (1,1), and an integrated GARCH. As a reference, we also provide the statistics of a scaled GARCH(1,1) as well as the realized volatility that is used as proxy for the true next-period volatility. We have also tried numerous higher order GARCH(p,q) models for $p = 1, \dots, 6$ and $q = 1, \dots, 6$ as well as the asymmetric GJR-GARCH of Glosten, Jagannathan, and Runkle (1993), for iterated and direct forecasts. However, these forecasts were almost always dominated by the GARCH(1,1) model in terms of multi-period, out-of-sample MSFE. This is a multi-period version of the results of Hansen and Lunde (2005) who show that a simple GARCH(1,1) has a smaller one-period out-of-sample MSFE than more richly parameterized GARCH models. Hence, we focus on the direct and iterated GARCH(1,1) models from this point onward in the interest of brevity.⁴

The statistics that we display—the annualized mean, annualized volatility, skewness, kurtosis, first-order autocorrelation, and probability of a forecast being higher than the realized volatility—can be summarized as follows. First, all forecasts are upward biased, as their means are higher than the averaged realized volatility. Moreover, the probability of all models to predict a volatility higher than the next period volatility is about 65 percent across horizons and most models. The direct GARCH is an exception as its bias is higher than that of the other approaches and its probability of over-prediction is about 70 percent. This upward bias is undoubtedly due to outliers such as October 1987 that non of the above models can accommodate. When we exclude the three highest volatile months from the sample, the bias is significantly reduced.

Second, the multi-period forecasts are much less volatile than the multi-period realized volatility. The square-root rule is extremely smooth at all horizons, which is the main reason for its poor forecasting performance. The GARCH and MIDAS models are much more time-varying, but their volatility is still about 50 percent lower than that of the realized volatility,

⁴The results from the other GARCH models are available upon request.

at all horizons. Third, all multi-period forecasts are significantly more skewed than the realized volatility. The exception is the square-root rule whose skewness is less than that of the realized volatility. Finally, while the multi-period realized volatility is quite persistent as are the GARCH forecasts, some of the MIDAS models, such as the exponential and restricted exponential MIDAS, exhibit significantly less serial correlation.

In Table 2, we compare the direct, iterated, and MIDAS forecasts to the scaling-up approach. Scaling the one-period volatility by the horizon k is admittedly a naive approach and is not directly comparable to the other three methods. The summary statistics in Table 1 also suggest that it may not be the best forecasting approach. However, despite the evidence against this method (Diebold, Hickman, Inoue, and Schuermann (1997)), it is still widely used in practice.⁵ We use the scaling approach as a benchmark not because we think it is a particularly hard forecast to beat, but because of its widespread use in the profession. The MIDAS forecasts are computed using the hyperbolic specification (2.5), which as we will see below has some desirable properties.

In panel A, we report the level of the MSFE, 5, 10, 15, 20, 25, 30, and 60 days ahead in addition to the one-period ahead forecast. Not surprisingly, the MSFEs increase but at a rate slower than \sqrt{k} . The one-period ahead forecasts of the iterated, direct forecasts are the same, by construction. However, they differ as the horizon k increases. At five periods ahead, the iterated, direct, MIDAS, and integrated forecasts have comparable MSFEs. The k-rule forecasts are the exception with significantly higher MSFEs. As the horizon increases, the MSFE of the k-rule is similar to that of the direct forecasts. In fact, at horizons of 20 periods and higher, the k-rule forecasts are better than the direct forecasts, while the integrated forecasts are slightly better.

This pattern is perhaps best observed in Panel B, which reports the MSFE's of the forecasts relative to that of the square-root rule. All other forecasts dominate the scaling rule at shorter horizons of 5 to 10 days. As horizons longer than 10 days, the relative forecasting performance of the iterated and integrated methods subside and the direct forecast actually has higher MSFE by as much as 23 percent at 60 days. The rapid deterioration of the long-horizon direct forecast is consistent with the findings of Christoffersen and Diebold (2000). The relatively lower advantage of the iterated forecast is also consistent with the theoretical papers that have emphasized the bias of the iterated forecasts. As the iterations increase,

⁵This method has been mentioned in the Basel II agreement, which might explain to a great extent its use by professionals.

so does the bias. Hence, our results suggest that the iterated method is clearly suitable for shorter horizon forecasts in the range of 5-period (one week) to 20-periods (one month).

The MIDAS method produces the best forecasts at long horizons. In Panel B, its relative MSFE is better than the scaling rule at 5-period ahead forecasts and higher. More importantly, its forecasting performance does not deteriorate nearly as rapidly as that of the other methods. At all horizons, it produces the best forecast and its advantage relative to the other methods increases steadily. At 60 days, its MSFE is about 23 percent lower than that of the scaling rule. The long-horizon forecastability of the market volatility with the MIDAS is a new finding.

To judge the statistical significance of the forecasting improvements, we use the West (1996) and Giacomini and White (2006) tests for predictive ability. As discussed in section 3.3, the first method is unconditional whereas the second is conditional. Henceforth, we report asymptotic p-values from both tests in Panels C and D of Table 2, as they test for different null hypotheses. If the p-value is smaller than 1 percent, we display 0.00. The following results are worth noting. The direct forecasts are insignificantly better than the scaling rule at a the 20-day horizon using the West (1996) test. At longer horizons, the scaling rule is actually significantly better than the direct forecast using both tests. This suggests that the direct method should not be used for multi-period volatility forecasting.

The other forecasting methods are significantly better than the scaling approach. However, their superior performance deteriorates rapidly with k . At the longest forecasting horizon we consider, $k = 60$ days, the iterated forecasts are not statistically better than the scaling rule, according to the West (1996) test whose p-value is 0.08. The MIDAS forecasts stand out in Table 2 as providing the best multi-period ahead predictions. Even at 60 days, their superior performance is statistically significant at one percent under the West (1996) and Giacomini and White (2006) tests.

As a final comparison, we test whether the best performing test, at any horizon, is significantly better than the second best forecast. Significance at the five percent level using the West (1996) and Giacomini and White (2006) tests is denoted by bold the numbers in the table. For the one-period forecasts, the direct and iterated GARCH methods dominate all other approaches. At 10 day horizons and longer, however, the MIDAS forecast is statistically significantly better than the other forecasting methods which confirms its superior multi-period forecasting performance.

It might be argued that the MIDAS approach has an unfair advantage in this out-of-sample exercise, because we have chosen the hyperbolic specification (2.5) which is known from previous work to produce good results. While the same comment can be levied against the other methods, it is interesting to see whether alternative polynomial specifications produce vastly different forecasts. In Table 3, we compute the MSFE of various MIDAS specifications. The polynomial weights that we use in the MIDAS are the hyperbolic (2.5), the linear (2.4), the beta (2.3), the exponential (2.2), and various specifications of step functions (2.6). We also display the MSFE of the iterated forecast (as a reference) which was shown to produce the best prediction among the non-MIDAS approaches.

Panel A of Table 3 provides the MSFEs which are directly comparable with Table 3, while Panel B provides the MSFEs relative to the iterated approach. Focusing our attention on Panel B, we note that exponential weights are the worst overall. While these weights have been used successfully by Ghysels, Santa-Clara, and Valkanov (2005) to estimate the risk return tradeoff, in the context of volatility forecasting, they are dominated by the other methods. It is not surprising to find that different weights will be appropriate in different applications. The suitability of the weighting function will be determined by the stochastic properties of the predicted variable and it is not reasonable to expect one functional form of $b_k(j, \theta)$ to dominate across applications.

The hyperbolic specification produces the best multi-period forecasts. This is not surprising, given that it is very similar to an ARFIMA model (whose impulse response function is also hyperbolically decaying) and ARFIMAs produce good out-of-sample forecasts of future volatility (cite). However, it is interesting to note that the beta and most of the step specifications produce very good results, as well. With the exception of the exponential MIDAS, all other models dominate the iterated forecasts at long-horizons.

Tables 4, 5, and 6 display results similar to those of Table 2 for the five size, five book-to-market, and ten industry portfolios, respectively. In the interest of conciseness, we have omitted the k-rule forecasts, as they are inferior to the alternative models for all portfolios and all horizons. First, we notice that the portfolio MSFE across forecasting methods are larger than the corresponding market MSFE in Table 2. In other words, the volatilities of the portfolio returns are less predictable than the volatility of the market portfolio.

Looking at Panel B of Table 4, the MSFE of the direct forecasts are markedly higher than those of the MIDAS and iterated forecasts. For the smallest cap stocks, they are about 369

percent larger at 60-period ahead. For the largest cap stocks, they are about 37 percent larger and the decrease is monotonic. Hence, the direct approach is particularly inappropriate to use for volatile, small-cap stocks. Turning to the significance results in Panels C and D, the MSFEs of the MIDAS are significantly lower than those of the iterated forecasts, especially at long horizons. Moreover, the large cap stocks are significantly more predictable than are smaller cap stocks. Tables 5 and 6 contains similar results for the book-to-market and industry portfolio returns. A slight difference is that volatility of these portfolio returns seems to be more forecastable than that of the size portfolios as the levels of the MSFE in Tables 5 and 6 are lower than those in Table 4.

5 Conclusion

We consider two widely used methods to forecast volatility at long horizons: iterated and direct forecasts. In addition, we use a relatively new third approach – MIDAS. All three approaches yield multi-step ahead volatility forecasts without relying on the restrictive assumption of i.i.d. returns that is implicit in the often-used scaling approach. We compare these forecasting methods in terms of their average forecasting accuracy—using the MSFE. Since no general analytic results are possible, the comparison is carried out using daily stock market returns from 1963 to 2007 for the US stock market as well as five size, five book-to-market, and ten industry portfolios. All forecasts are (pseudo) out-of-sample.

The MIDAS forecasts are significantly more precise than the direct and iterated forecasts according to the West (1996) and Giacomini and White (2006) tests of predictive ability. We also document sizeable differences in the two tests. We conjecture that the gains in forecasting power in the MIDAS approach are due to the ability of the approach to take advantage of the bias-efficiency trade-off that exists in multi-period forecasts. While the other two approaches are either efficient but biased (iterated) or unbiased but inefficient (direct), the MIDAS strikes a balance between the two.

References

- Andersen, Torben, and Tim Bollerslev, 1998a, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885–905.
- , 1998b, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885–905.
- , and Francis Diebold, 2003, Parametric and nonparametric volatility measurement, in Yacine Aït-Sahalia, and Lars P. Hansen, ed.: *Handbook of Financial Econometrics*.
- Andersen, Torben, Tim Bollerslev, Frank X. Diebold, and P. Labys, 2001, The distribution of exchange rate volatility, *Journal of American Statistical Association* 96, 42–55.
- Andersen, Torben, Tim Bollerslev, and S. Lange, 1999, Forecasting Financial Market Volatility: Sample Frequency vis-a-vis Forecast Horizon, *Journal of Empirical Finance* 6, 457–477.
- Bhanzali, R.J., 1999, Parameter estimation and model selection for multistep prediction of a time series: A review, in S. Ghosh, ed.: *Asymptotics, Nonparametrics, and Time Series* (Marcel Dekker).
- Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307–327.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press: Princeton).
- Chevillon, Chillaume, and David Hendry, 2005, Non-parametric direct multi-step estimation for forecasting economic processes, *International Journal of Forecasting* 21, 201–218.
- Christoffersen, Peter, and Francis Diebold, 2000, How relevant is volatility forecasting for financial risk management?, *Review of Economics and Statistics* 82, 12–22.
- Clements, M., and David Hendry, 1996, Mutli-step estimation for forecasting, *Oxford Bulletin of Economics and Statistics* 58, 657–684.
- Corsi, F., 2004, A simple long memory model of realized volatility, Tech. rep., University of Southern Switzerland.

- Diebold, Francis, Andrew Hickman, Atsushi Inoue, and Til Schuermann, 1997, Converting 1-day volatility to h-day volatility: Scaling by \sqrt{h} is worse than you think, mimeo, University of Pennsylvania.
- Diebold, Francis, and R. Mariano, 1995, Comparing predictive accuracy, *Journal of Business and Economic Statistics* 13, 253–263.
- Elliott, Graham, and Allan Timmermann, 2008, Economic forecasting, *Journal of Economic Literature* 46, 3–56.
- Engle, R.F., 1982, Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation, *Econometrica* 50, 987–1008.
- Findley, D., 1983, On the use of multiple models of multi-period forecasting, *Proceedings of the Business and Statistics Section, American Statistical Association* pp. 528–531.
- , 1985, Model selection for multi-step-ahead forecasting, in S. Ghosh, and P. Young, ed.: *Proceedings of the 7th Symposium on Identification and System Parameter Estimation* . pp. 1039–1044 (Pergamon, Oxford).
- Forsberg, Lars, and Eric Ghysels, 2006, Why do absolute returns predict volatility so well?, *Journal of Financial Econometrics* 5, 31–67.
- Foster, Dean P., and Daniel B. Nelson, 1994, Asymptotic filtering theory for univariate arch models, *Econometrica* 62, 1–41.
- French, Kenneth R., William Schwert, and Robert F. Stambaugh, 1987a, Expected stock returns and volatility, *Journal of Financial Economies* 19, 3–29.
- , 1987b, Expected stock returns and volatility, *Journal of Financial Economies* 19, 3–29.
- Ghysels, Eric, and Alastair Hall, 1990, A test for structural stability of euler conditions parameters estimated via the generalized method of moments estimator, *International Economic Review* pp. 355–364.
- Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2002, The MIDAS touch: Mixed data sampling regression models, Working paper, UNC and UCLA.

- , 2005, There is a risk-return tradeoff after all, *Journal of Financial Economics* 76, 509–548.
- , 2006, Predicting volatility: getting the most out of return data sampled at different frequencies, *Journal of Econometrics* 131, 59–95.
- Ghysels, Eric, Arthur Sinko, and Rossen Valkanov, 2006, MIDAS Regressions: Further Results and New Directions, *Econometric Reviews* 26, 53–90.
- Giacomini, Raffaella, and Halbert White, 2006, Tests of conditional predictive ability, *Econometrica* 15, 549–582.
- Glosten, Lawrence R., Ravi Jagannathan, and David 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.
- Hansen, Peter, and Asger Lunde, 2005, A forecast comparison of volatility models: Does anything beat a garch(1,1)?, *Journal of Applied Econometrics* 20, 873–889.
- Hoffman, D., and Adrian Pagan, 1989, Post-sample prediction tests for generalized method of moments estimators, *Oxford Bulletin of Economics and Statistics* pp. 333–343.
- Hosking, J., 1981, Fractional differencing, *Biometrika* 68, 165–176.
- J.P.Morgan/Reuters, 1996, Riskmetrics–technical document, J.P. Morgan/Reuters, New York.
- Lin, J., and Clive Granger, 1994, Model selection for multiperiod forecasts, *Journal of Forecasting* 13, 1–9.
- Marcellino, Massimiliano, James Stock, and Mark Watson, 2006, A comparison of direct and iterated multistep ar methods for forecasting macroeconomic time series, *Journal of Econometrics* 135, 499–526.
- Merton, Robert C., 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.
- Patton, Andrew, 2007, Volatility forecast comparison using imperfect volatility proxies, Working Paper, London School of Economics.

- Schwert, William G., 1989, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115–1153.
- Tanaka, Katsuto, 1999, The nonstationary fractional unit root, *Econometric Theory* 15, 549–582.
- West, Kenneth, 1996, Asymptotic inference about predictive ability, *Econometrica* 64, 1067–1084.
- , and D. Cho, 1995, The predictive ability of several models of exchange rate volatility, *Journal of Econometrics* 69, 367–391.

Tables

Table 1: Sample Statistics of Multi-period Volatility Forecasts – Market Portfolio.

The table reports standard descriptive statistics for the volatility forecasts time-series at different horizons from different forecasting methods. The out-of-sample forecasts obtain by re-estimating the model at each step and using it to formulate a k-period ahead prediction ($k=1, \dots, 60$). The data is from July 1, 1963 to December 31, 2007. The first 4,000 daily observations are used to estimate the first forecast. The forecasting models include the MIDAS method with the different weighting-functions discussed in the main text, the AR(1)-GARCH(1,1) iterated method, the AR(1)-GARCH(1,1) direct method, integrated volatility method from an autoregression, and the scaling-up (or square-root rule) method. The entry Proxy shows the values of the sample statistics for the true volatility process proxied by the realized volatility. The sample statistics include the annualized mean value of the volatility forecasts, the annualized standard deviation, skewness, kurtosis, first-order correlation and the sample frequency of overprediction in relation to the proxy considered.

Method	Horizon								Horizon							
	1	5	10	15	20	25	30	60	1	5	10	15	20	25	30	60
Annualized Mean (%)																
V_M HYPERB	14.06	14.04	14.10	14.13	14.19	14.23	14.24	14.31	5.27	5.02	4.34	4.14	4.17	4.16	4.03	3.44
V_M LINEAR	14.34	14.31	14.28	14.27	14.27	14.27	14.27	14.27	4.05	3.72	3.38	3.21	3.11	3.05	2.99	2.72
V_M BETA	13.87	13.84	14.21	14.24	14.22	14.48	14.14	14.15	5.67	5.30	6.44	4.68	6.27	5.98	4.63	3.08
V_M BETA REST	13.94	13.86	13.99	14.11	14.18	14.19	14.22	14.38	5.45	5.24	4.47	4.29	4.76	4.69	4.51	3.83
V_M EXP	13.91	13.85	14.06	14.15	14.20	14.23	14.20	14.36	5.71	5.23	4.49	4.00	5.24	4.98	4.64	3.60
V_M EXP REST	14.14	13.86	14.01	14.09	14.13	14.14	14.16	14.24	4.07	4.72	3.92	3.30	3.64	3.18	2.90	2.84
V_M GEOMET	13.94	13.84	13.88	14.06	14.15	14.18	14.20	14.35	5.72	5.32	4.61	4.47	4.90	4.77	4.50	3.83
V_I	13.84	13.88	13.93	13.97	14.00	14.03	14.06	14.17	6.24	6.09	5.93	5.78	5.64	5.52	5.40	4.82
V_D	13.84	15.25	15.59	15.74	15.64	15.62	15.57	15.30	6.24	5.47	5.34	4.62	4.16	3.92	3.99	4.11
V_{int}	14.08	14.05	14.08	14.10	14.12	14.16	14.17	14.19	1.18	4.55	4.32	4.13	4.02	3.88	3.74	2.73
$\sqrt{k}V_t$	14.08	14.08	14.08	14.08	14.08	14.08	14.08	14.08	1.18	1.18	1.18	1.18	1.18	1.18	1.18	1.18
PROXY	10.52	12.85	13.23	13.39	13.50	13.58	13.64	13.88	11.04	8.20	7.58	7.29	7.10	6.95	6.83	6.39
Skewness																
V_M HYPERB	6.24	4.81	5.65	5.05	5.70	5.83	5.41	3.74	116.80	56.37	71.63	56.44	71.53	74.41	64.84	28.38
V_M LINEAR	3.18	3.12	3.01	2.99	3.00	3.03	3.10	3.31	17.29	17.04	16.40	16.55	16.70	16.69	17.39	19.41
V_M BETA	22.80	5.69	6.66	5.86	9.47	8.33	8.96	2.43	1010.30	74.42	67.20	55.20	130.00	96.59	136.42	13.11
V_M BETA REST	19.21	6.20	6.83	6.40	7.06	7.36	6.68	3.16	796.37	82.91	91.11	80.13	91.52	103.12	86.27	19.73
V_M EXP	23.71	5.18	6.51	6.45	7.22	8.82	7.44	5.56	1094.86	58.92	92.64	79.92	87.57	128.47	103.41	55.30
V_M EXP REST	57.65	4.33	5.31	5.13	9.04	7.24	6.48	6.10	4237.64	49.53	73.68	62.01	194.13	123.04	110.91	81.65
V_M GEOMET	23.08	5.62	6.18	5.85	7.17	6.97	6.70	3.23	1079.60	68.75	82.23	70.35	97.39	94.55	87.23	21.45
V_I	4.19	4.31	4.46	4.60	4.74	4.87	5.00	5.64	35.34	37.20	39.49	41.74	43.94	46.10	48.21	59.61
V_D	4.19	3.28	4.58	2.41	2.29	2.48	2.10	3.74	35.34	22.71	48.44	14.38	15.21	20.97	11.98	34.97
V_{int}	0.32	16.14	14.20	11.54	9.96	8.91	7.88	4.55	2.60	485.73	331.36	210.38	152.33	120.68	95.55	37.16
$\sqrt{k}V_t$	0.32	0.32	0.32	0.32	0.32	0.32	0.32	0.32	2.60	2.60	2.60	2.60	2.60	2.60	2.60	2.60
PROXY	4.77	4.71	4.42	4.04	3.72	3.47	3.26	2.53	78.72	58.30	46.32	36.91	30.44	25.91	22.58	13.25
First-order Autocorrelation																
V_M HYPERB	0.89	0.90	0.92	0.92	0.92	0.92	0.93	0.93	0.72	0.67	0.68	0.67	0.67	0.67	0.67	0.65
V_M LINEAR	1.00	1.00	0.99	0.99	1.00	1.00	1.00	1.00	0.74	0.70	0.70	0.70	0.69	0.69	0.68	0.65
V_M BETA	0.73	0.68	0.86	0.94	0.94	0.96	0.93	0.96	0.72	0.64	0.66	0.67	0.65	0.67	0.66	0.64
V_M BETA REST	0.73	0.73	0.80	0.85	0.88	0.90	0.92	0.91	0.72	0.65	0.66	0.67	0.65	0.65	0.65	0.65
V_M EXP	0.69	0.66	0.48	0.57	0.87	0.91	0.92	0.67	0.72	0.64	0.65	0.65	0.65	0.66	0.66	0.63
V_M EXP REST	0.14	0.20	0.21	0.25	0.22	0.24	0.27	0.27	0.74	0.65	0.66	0.65	0.64	0.64	0.64	0.63
V_M GEOMET	0.68	0.63	0.71	0.83	0.80	0.87	0.91	0.85	0.72	0.64	0.65	0.66	0.65	0.65	0.66	0.64
V_I	0.98	0.98	0.98	0.98	0.98	0.98	0.98	0.98	0.71	0.65	0.64	0.64	0.64	0.64	0.64	0.63
V_D	0.98	0.83	0.77	0.85	0.84	0.85	0.90	0.93	0.71	0.73	0.75	0.75	0.74	0.73	0.72	0.67
V_{int}	1.00	0.86	0.92	0.95	0.97	0.98	0.98	0.99	0.74	0.67	0.67	0.66	0.66	0.66	0.65	0.64
$\sqrt{k}V_t$	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.74	0.68	0.66	0.66	0.64	0.64	0.64	0.62
PROXY	0.17	0.91	0.97	0.98	0.99	0.99	0.99	1.00	-	-	-	-	-	-	-	-
Prob. Overprediction																
V_M HYPERB	0.89	0.90	0.92	0.92	0.92	0.92	0.93	0.93	0.72	0.67	0.68	0.67	0.67	0.67	0.67	0.65
V_M LINEAR	1.00	1.00	0.99	0.99	1.00	1.00	1.00	1.00	0.74	0.70	0.70	0.70	0.69	0.69	0.68	0.65
V_M BETA	0.73	0.68	0.86	0.94	0.94	0.96	0.93	0.96	0.72	0.64	0.66	0.67	0.65	0.67	0.66	0.64
V_M BETA REST	0.73	0.73	0.80	0.85	0.88	0.90	0.92	0.91	0.72	0.65	0.66	0.67	0.65	0.65	0.65	0.65
V_M EXP	0.69	0.66	0.48	0.57	0.87	0.91	0.92	0.67	0.72	0.64	0.65	0.65	0.65	0.66	0.66	0.63
V_M EXP REST	0.14	0.20	0.21	0.25	0.22	0.24	0.27	0.27	0.74	0.65	0.66	0.65	0.64	0.64	0.64	0.63
V_M GEOMET	0.68	0.63	0.71	0.83	0.80	0.87	0.91	0.85	0.72	0.64	0.65	0.66	0.65	0.65	0.66	0.64
V_I	0.98	0.98	0.98	0.98	0.98	0.98	0.98	0.98	0.71	0.65	0.64	0.64	0.64	0.64	0.64	0.63
V_D	0.98	0.83	0.77	0.85	0.84	0.85	0.90	0.93	0.71	0.73	0.75	0.75	0.74	0.73	0.72	0.67
V_{int}	1.00	0.86	0.92	0.95	0.97	0.98	0.98	0.99	0.74	0.67	0.67	0.66	0.66	0.66	0.65	0.64
$\sqrt{k}V_t$	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.74	0.68	0.66	0.66	0.64	0.64	0.64	0.62
PROXY	0.17	0.91	0.97	0.98	0.99	0.99	0.99	1.00	-	-	-	-	-	-	-	-

Table 2: Multi-Period MSFEs of Volatility Forecasts – Market Portfolio

The table reports mean square forecasting errors (MSFE) of the market volatility from the iterated method, V_b direct method, V_{D_t} MIDAS with hyperbolic weights, V_M HYPERB, integrated V_{int} , and scaling-up $\sqrt{k}v_1$ methods. The out-of-sample forecasts obtain by re-estimating the model at each step and using it to formulate a k-period ahead prediction ($k=1, \dots, 60$). The data is from July 1, 1963 to December 31, 2007. The first 4,000 daily observations are used to estimate the first forecast.

Panel A reports the level of the MSFE's (the smallest MSFE given k is in bold letters).

Panel B reports the values of the Improvement Ratio, which is defined as $100 * [\text{MSFE}(\text{Benchmark}) - \text{MSFE}(\text{Alternative})] / \text{MSFE}(\text{Benchmark})$, where the benchmark model is the scaling-up method. Panel C reports the p -values of the West's tests for equal forecasting ability, whereas Panel D reports the p -values of the test for equal conditional forecasting ability of Giacomini and White.

[illegible]

Table 3: Multi-Period MSFEs of Volatility Forecasts – Market Portfolio.

The table reports mean square forecasting errors (MSFE) of the market return volatility for various MIDAS specifications. The specifications include the hyperbolic model, denoted V_M HYPERB; the linear model, denoted V_M LINEAR; the unrestricted and restricted versions of the beta model, denoted V_M BETA and V_M BETA REST, respectively; the unrestricted and restricted versions of the exponential model, denoted V_M EXP and V_M EXP REST, respectively, and the geometric model, denoted V_M GEOMET. As a reference, we also report the MSFE of the iterated model. The out-of-sample forecasts obtain by re-estimating the model at each step and using it to formulate a k-period ahead prediction. The data is from July 1, 1963 to December 31, 2007. The first 4,000 daily observations are used to estimate the first forecast. Panels A,B,C, and D report the same statistical information as in Table 2 above, using V_M HYPERB (the model with smallest MSFE) as a benchmark.

Method	Horizon								Horizon							
	1	5	10	15	20	25	30	60	1	5	10	15	20	25	30	60
	Panel A: MSFE of ALL models (x10000)								Panel B: Improvement Ratio over V_M HYPERB (%)							
V_M LINEAR	0.52	1.17	1.94	2.67	3.38	4.05	4.72	8.60	-5.33	-22.19	-16.07	-11.48	-8.31	-7.26	-7.52	-8.16
V_M BETA	0.52	1.12	2.32	2.83	4.42	5.77	5.48	9.57	-6.63	-16.97	-39.35	-18.12	-41.84	-52.81	-24.78	-20.35
V_M BETA REST	0.51	1.09	1.86	2.59	3.54	4.18	4.83	8.59	-4.91	-14.68	-11.77	-8.23	-13.49	-10.79	-10.12	-8.00
V_M EXP	0.53	1.12	2.19	3.02	3.82	4.49	4.90	10.94	-7.80	-17.61	-31.39	-26.31	-22.69	-18.87	-11.58	-37.59
V_M EXP REST	0.57	1.34	2.22	3.02	4.02	4.65	5.35	9.72	-17.23	-40.01	-33.36	-25.99	-28.91	-23.06	-21.90	-22.31
V_M GEOMET	0.52	1.13	1.93	2.62	3.60	4.21	4.79	8.61	-7.23	-18.79	-15.52	-9.27	-15.43	-11.41	-9.20	-8.33
V_I	0.48	0.99	1.71	2.44	3.15	3.85	4.53	8.62	1.02	-3.62	-2.39	-1.89	-1.17	-1.93	-3.30	-8.42
V_M HYPERB	0.49	0.95	1.67	2.39	3.12	3.78	4.39	7.95	-	-	-	-	-	-	-	-
	Panel C: West test p-values								Panel D: GW test p-values							
V_M LINEAR	0.05	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00
V_M BETA	0.18	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.40	0.00	0.00	0.00	0.00	0.00	0.00	0.00
V_M BETA REST	0.19	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.48	0.00	0.00	0.00	0.00	0.00	0.00	0.00
V_M EXP	0.16	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.38	0.00	0.00	0.00	0.00	0.00	0.00	0.00
V_M EXP REST	0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.09	0.00	0.00	0.00	0.00	0.00	0.00	0.00
V_M GEOMET	0.18	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.45	0.00	0.00	0.00	0.00	0.00	0.00	0.00
V_I	0.34	0.16	0.33	0.39	0.44	0.39	0.34	0.24	0.81	0.02	0.02	0.02	0.02	0.03	0.02	0.15

Table 4: Multi-Period MSFEs of Volatility Forecasts – Size Portfolios.

The table reports mean square forecasting errors (MSFE) of the return volatility of five size-sorted portfolios. The forecasts are obtained using the iterated method, $V_{\hat{r}}$, the direct method, $V_{\hat{D}}$, the MIDAS with hyperbolic weights, V_M , and the integrated method, V_{int} . The out-of-sample forecasts obtain by re-estimating the model at each step and using it to formulate a k -period ahead prediction. The data is from July 1, 1963 to December 31, 2007. The first 4,000 daily observations are used to estimate the first forecast. Panel A1 reports the level of the MSFE's. Panels A,B,C, and D report the same statistical information as in Table 2 above, using V_M (the model with smallest MSFE) as benchmark.

Panel A: MSFE of ALL models (x10000)										Panel B: Improvement Ratio over VM (%)									
Size	Method	1	5	10	15	20	25	30	60	1	5	10	15	20	25	30	60		
1	Vi	0.33	0.84	1.53	2.24	2.96	3.85	4.33	8.19	5.92	-0.12	-0.83	-0.91	-1.48	-2.09	-3.44	-5.47		
2	Vi	0.42	0.91	1.59	2.29	3.00	3.70	4.37	8.10	5.64	-0.36	-0.47	0.69	0.46	-0.11	-0.61	-1.84		
3	Vi	0.41	0.86	1.52	2.18	2.84	3.49	4.11	7.76	3.87	-0.23	-0.46	0.70	0.23	-0.22	-0.96	-4.06		
4	Vi	0.44	0.90	1.59	2.31	3.03	3.74	4.43	8.54	1.78	-1.38	-0.75	-0.42	-1.15	-1.77	-3.51	-8.16		
5	Vi	0.56	1.16	1.98	2.81	3.61	4.40	5.17	9.86	-1.72	-4.54	-2.87	-3.27	-1.52	-3.02	-4.88	-12.41		
1	VD	0.33	1.48	3.99	6.09	9.77	13.86	16.86	36.47	5.92	-77.48	-163.93	-173.81	-235.52	-288.13	-303.18	-369.66		
2	VD	0.42	1.41	3.64	5.31	7.50	10.03	11.99	23.09	5.64	-55.87	-129.69	-130.21	-148.83	-171.59	-176.34	-190.15		
3	VD	0.41	1.29	3.05	4.44	6.16	8.09	9.52	18.95	3.87	-50.33	-102.06	-102.37	-116.51	-132.58	-133.81	-154.15		
4	VD	0.44	1.26	2.71	3.86	5.18	6.43	7.69	15.78	1.78	-41.13	-72.04	-68.28	-72.72	-74.82	-79.50	-99.87		
5	VD	0.56	1.18	2.12	2.85	3.77	4.57	5.53	12.02	-1.72	-6.57	-10.57	-5.05	-6.23	-7.11	-12.21	-37.09		
1	Vint	0.49	0.86	1.64	2.41	3.11	3.74	4.36	8.16	-39.52	-2.45	-8.08	-8.23	-6.73	-4.77	-4.32	-5.14		
2	Vint	0.67	0.95	1.76	2.60	3.34	4.02	4.62	8.58	-50.56	-4.36	-11.13	-12.55	-10.83	-8.82	-6.49	-7.82		
3	Vint	0.66	0.90	1.70	2.50	3.18	3.83	4.39	8.10	-55.88	-4.71	-12.94	-13.98	-11.80	-9.94	-7.73	-8.70		
4	Vint	0.72	0.96	1.79	2.65	3.39	4.06	4.65	8.66	-62.15	-7.28	-13.84	-15.29	-12.83	-10.37	-8.45	-9.70		
5	Vint	0.90	1.31	2.37	3.38	4.25	5.00	5.58	9.55	-64.09	-18.07	-23.34	-24.36	-19.62	-17.18	-13.26	-8.96		
1	VM	0.35	0.83	1.51	2.22	2.91	3.57	4.18	7.76	-	-	-	-	-	-	-	-		
2	VM	0.45	0.91	1.59	2.31	3.02	3.69	4.34	7.96	-	-	-	-	-	-	-	-		
3	VM	0.43	0.86	1.51	2.19	2.85	3.48	4.07	7.46	-	-	-	-	-	-	-	-		
4	VM	0.45	0.89	1.57	2.30	3.00	3.68	4.28	7.90	-	-	-	-	-	-	-	-		
5	VM	0.55	1.11	1.92	2.72	3.55	4.27	4.93	8.77	-	-	-	-	-	-	-	-		
Panel C: West test p-values										Panel D: GW test p-values									
Size	Method	1	5	10	15	20	25	30	60	1	5	10	15	20	25	30	60		
1	Vi	0.00	0.48	0.44	0.45	0.43	0.41	0.35	0.29	0.00	0.00	0.00	0.01	0.01	0.00	0.01	0.06		
2	Vi	0.00	0.42	0.43	0.44	0.46	0.49	0.45	0.38	0.00	0.00	0.00	0.01	0.01	0.01	0.03	0.13		
3	Vi	0.00	0.45	0.44	0.44	0.48	0.48	0.41	0.20	0.00	0.00	0.00	0.01	0.01	0.02	0.04	0.22		
4	Vi	0.05	0.32	0.43	0.47	0.43	0.38	0.30	0.16	0.26	0.00	0.01	0.01	0.03	0.03	0.07	0.44		
5																			

Table 5: Multi-Period MSFEs of Volatility Forecasts – Book-to-Market Portfolios.

The table reports mean square forecasting errors (MSFE) of the return volatility of five book-to-market-sorted portfolios. The forecasts are obtained using the iterated method, V_I , the direct method, V_D , the MIDAS with hyperbolic weights, V_M , and the integrated method, V_{int} . The out-of-sample forecasts obtain by re-estimating the model at each step and using it to formulate a k-period ahead prediction. The data is from July 1, 1963 to December 31, 2007. The first 4,000 daily observations are used to estimate the first forecast. Panel A1 reports the level of the MSFE's. Panels A,B,C, and D report the same statistical information as in Table 2 above, using V_M (the model with smallest MSFE) as benchmark.

		Horizon								Horizon							
		1	5	10	15	20	25	30	60	1	5	10	15	20	25	30	60
Panel A: MSFE of ALL models (x10000)										Panel B: Improvement Ratio over V_M (%)							
BTM	Method																
1	V_I	0.60	1.17	1.94	2.71	3.44	4.16	4.87	9.24	1.35	-2.63	-1.60	-0.29	1.68	1.39	0.34	-2.28
2	V_I	0.50	1.03	1.77	2.54	3.29	4.02	4.72	8.86	1.25	-3.58	-3.40	-3.71	-2.69	-3.89	-6.18	-12.13
3	V_I	0.45	0.91	1.57	2.27	2.95	3.61	4.25	8.16	1.02	-4.47	-4.46	-4.71	-4.95	-5.49	-7.42	-14.30
4	V_I	0.40	0.92	1.60	2.28	2.95	3.58	4.18	7.57	-1.17	-4.66	-5.12	-6.19	-5.64	-10.77	-11.46	-12.28
5	V_I	0.41	0.86	1.47	2.10	2.71	3.29	3.83	7.06	0.86	-3.05	-2.71	-2.99	-2.18	-2.04	-3.27	-6.93
1	V_D	0.60	1.35	2.58	3.61	4.69	5.80	7.07	14.57	1.35	-19.00	-34.68	-33.62	-34.04	-37.68	-44.54	-61.27
2	V_D	0.50	1.22	2.46	3.48	4.39	5.50	6.47	13.37	1.25	-21.83	-44.04	-41.99	-36.75	-42.31	-45.80	-69.20
3	V_D	0.45	1.06	2.10	2.93	3.68	4.43	5.37	11.26	1.02	-21.63	-39.30	-35.16	-31.13	-29.51	-35.68	-57.69
4	V_D	0.40	0.99	1.75	2.39	3.13	3.90	4.64	9.12	-1.17	-13.16	-15.07	-11.22	-12.29	-20.55	-23.65	-35.38
5	V_D	0.41	1.02	2.04	2.95	4.27	5.27	6.47	13.51	0.86	-22.07	-41.98	-44.77	-61.23	-63.74	-74.47	-104.56
1	V_{int}	1.02	1.27	2.28	3.23	4.05	4.77	5.36	9.68	-66.87	-11.36	-19.48	-19.78	-15.90	-13.29	-9.74	-7.19
2	V_{int}	0.83	1.12	2.02	2.92	3.70	4.36	4.89	8.48	-63.26	-12.16	-18.29	-18.86	-15.23	-12.87	-10.06	-7.33
3	V_{int}	0.74	0.97	1.76	2.56	3.26	3.89	4.40	7.72	-63.67	-10.65	-16.93	-18.08	-16.15	-13.65	-10.98	-8.10
4	V_{int}	0.62	1.02	1.79	2.56	3.18	3.72	4.17	7.08	-56.28	-16.61	-17.87	-19.03	-14.09	-15.06	-11.18	-4.98
5	V_{int}	0.67	0.93	1.67	2.40	3.03	3.58	4.04	7.08	-63.11	-11.33	-16.00	-17.87	-14.30	-11.27	-8.87	-7.17

Table 6: Multi-Period MSFEs of Volatility Forecasts – Industries Portfolios.

The table reports mean square forecasting errors (MSFE) of the return volatility of ten industries portfolios. The forecasts are obtained using the iterated method, V_b , the direct method, V_D , the MIDAS with hyperbolic weights, V_M , and the integrated method, V_{int} . The out-of-sample forecasts obtain by re-estimating the model at each step and using it to formulate a k -period ahead prediction. The data is from July 1, 1963 to December 31, 2007. The first 4,000 daily observations are used to estimate the first forecast. Panel A,B,C, and D reports the level of the MSFE's. Panels A1, A2, B1 and B2 report the same statistical information as in Table 2 above, using V_M (the model with smallest MSFE) as benchmark.

		Horizon								Horizon							
		1	5	10	15	20	25	30	60	1	5	10	15	20	25	30	60
		Panel A: MSFE of ALL models (x10000)								Panel B: Improvement Ratio over VM (%)							
Industry	Method																
NoDur	Vf	0.44	0.95	1.57	2.20	2.81	3.40	3.97	7.35	-0.18	-4.24	-4.41	-4.87	-1.10	-3.28	-3.92	-6.27
Durbl	Vf	0.79	1.50	2.36	3.20	4.03	4.84	5.62	10.06	-0.66	-3.26	-3.32	-4.14	-2.90	-4.22	-5.01	-8.58
Manuf	Vf	0.53	1.14	1.98	2.83	3.64	4.41	5.17	9.58	2.45	-3.05	-2.60	-1.96	-0.56	-0.59	-1.09	-4.41
Enrgy	Vf	0.82	1.60	2.51	3.42	4.33	5.25	6.15	11.91	-2.14	-4.70	-5.28	-5.82	-7.04	-9.44	-11.09	-17.50
HiTec	Vf	1.12	1.99	3.26	4.52	5.73	6.91	8.05	15.76	1.21	-1.70	-2.36	-2.07	-1.33	-1.76	-3.51	-6.05
Telcm	Vf	0.65	1.32	2.13	2.95	3.75	4.51	5.22	9.34	-2.95	-4.02	-2.78	-2.97	-1.88	-3.11	-4.99	-0.42
Shops	Vf	0.60	1.19	1.93	2.69	3.41	4.10	4.79	9.00	1.75	-0.07	0.30	1.68	3.73	3.19	2.46	1.81
Hlth	Vf	0.64	1.30	2.10	2.90	3.64	4.36	5.07	9.15	1.17	-1.77	-2.38	-1.41	0.96	0.94	0.71	-0.07
Utils	Vf	0.36	0.88	1.52	2.19	2.85	3.47	4.07	7.30	-4.14	-7.82	-10.81	-13.71	-16.39	-19.50	-19.83	-24.39
Other	Vf	0.53	1.09	1.85	2.66	3.43	4.17	4.88	8.81	0.12	-2.89	-3.92	-4.08	-4.11	-4.85	-5.39	-5.04
NoDur	VD	0.44	1.07	1.96	2.83	3.66	4.65	5.48	11.62	-0.18	-17.65	-30.42	-34.92	-31.89	-41.29	-43.28	-67.94
Durbl	VD	0.79	1.67	2.78	3.83	4.94	5.89	7.09	13.48	-0.66	-14.77	-21.89	-24.59	-26.24	-26.81	-32.51	-45.53
Manuf	VD	0.53	1.40	2.89	3.94	5.21	6.34	7.53	14.15	2.45	-26.84	-49.81	-42.01	-44.06	-44.45	-47.17	-54.20
Enrgy	VD	0.82	1.65	2.69	3.77	5.00	6.37	7.70	16.17	-2.14	-7.79	-12.89	-16.77	-23.43	-32.89	-39.13	-59.55
HiTec	VD	1.12	2.25	4.15	6.00	7.88	9.60	11.85	24.62	1.21	-15.18	-30.12	-35.51	-39.41	-41.34	-52.33	-65.64
Telcm	VD	0.65	1.36	2.16	3.12	4.11	5.05	6.11	14.06	-2.95	-6.83	-4.22	-8.67	-11.79	-15.31	-22.73	-51.29
Shops	VD	0.60	1.45	2.82	4.18	5.62	7.03	8.55	21.40	1.75	-21.76	-45.15	-52.73	-58.63	-65.75	-74.09	-133.45
Hlth	VD	0.64	1.47	2.60	3.65	4.48	5.26	6.30	12.00	1.17	-14.89	-26.55	-27.72	-21.90	-19.56	-23.43	-31.23
Utils	VD	0.36	1.04	1.79	2.47	3.20	3.99	4.67	9.26	-4.14	-27.24	-30.23	-28.17	-30.64	-37.52	-37.45	-57.66
Other	VD	0.53	1.40	2.78	3.93	5.36	6.67	8.36	18.41	0.12	-32.71	-56.11	-54.16	-62.53	-67.51	-80.54	-119.41
NoDur	Vint	0.69	1.03	1.81	2.58	3.27	3.87	4.32	7.35	-55.22	-13.34	-20.09	-22.74	-17.88	-17.39	-12.98	-6.25
Durbl	Vint	1.31	1.59	2.59	3.51	4.38	5.19	5.84	9.83	-67.65	-8.90	-13.62	-14.28	-12.08	-11.77	-9.14	-6.12
Manuf	Vint	0.87	1.25	2.30	3.31	4.15	4.89	5.49	9.60	-59.38	-12.70	-19.14	-19.29	-14.82	-11.50	-7.39	-4.63
Enrgy	Vint	1.31	1.73	2.84	3.91	4.79	5.58	6.26	11.33	-63.07	-12.89	-19.00	-21.13	-18.33	-16.46	-13.10	-11.81
HiTec	Vint	1.98	2.25	3.90	5.46	6.78	7.95	8.93	16.49	-75.39	-15.24	-22.34	-23.30	-19.93	-17.06	-14.72	-10.98
Telcm	Vint	1.03	1.50	2.53	3.63	4.59	5.41	6.08	10.64	-62.69	-17.74	-22.31	-26.61	-24.81	-23.73	-22.27	-14.46
Shops	Vint	1.00	1.30	2.28	3.26	4.05	4.73	5.30	9.82	-62.75	-9.33	-17.66	-19.14	-14.43	-11.51	-7.85	-7.17
Hlth	Vint	1.04	1.41	2.39	3.32	4.11	4.81	5.40	9.53	-60.18	-10.12	-16.16	-16.14	-11.87	-9.27	-5.78	-4.26
Utils	Vint	0.52	0.88	1.54	2.21	2.79	3.30	3.72	6.40	-47.97	-7.77	-11.85	-14.95	-14.23	-13.75	-9.61	-8.94
Other	Vint	0.88	1.15	2.05	2.96	3.37	4.47	5.04	9.05	-65.24	-8.62	-14.97	-16.23	-14.37	-12.29	-8.82	-7.80
NoDur	VM	0.44	0.91	1.51	2.10	2.78	3.29	3.82	6.92	-	-	-	-	-	-	-	-
Durbl	VM	0.78	1.46	2.28	3.07	3.91	4.64	5.35	9.26	-	-	-	-	-	-	-	-
Manuf	VM	0.54	1.11	1.93	2.77	3.62	4.39	5.11	9.17	-	-	-	-	-	-	-	-
Enrgy	VM	0.80	1.53	2.38	3.23	4.05	4.79	5.53	10.14	-	-	-	-	-	-	-	-
HiTec	VM	1.13	1.95	3.19	4.43	5.66	6.79	7.78	14.86	-	-	-	-	-	-	-	-
Telcm	VM	0.63	1.27	2.07	2.87	3.68	4.38	4.97	9.30	-	-	-	-	-	-	-	-
Shops	VM	0.61	1.19	1.94	2.74	3.54	4.24	4.91	9.17	-	-	-	-	-	-	-	-
Hlth	VM	0.65	1.28	2.06	2.86	3.67	4.40	5.10	9.14	-	-	-	-	-	-	-	-
Utils	VM	0.35	0.82	1.37	1.93	2.45	2.90	3.40	5.87	-	-	-	-	-	-	-	-
Other	VM	0.53	1.05	1.78	2.55	3.30	3.98	4.63	8.39	-	-	-	-	-	-	-	-

Table 6: Multi-Period MSFEs of Volatility Forecasts – Industries Portfolios (Continuation)

[illegible]