

NBER WORKING PAPER SERIES

THERE IS A RISK-RETURN TRADEOFF AFTER ALL

Eric Ghysels
Pedro Santa-Clara
Rossen Valkanov

Working Paper 10913
<http://www.nber.org/papers/w10913>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2004

We thank Michael Brandt, Tim Bollerslev, Mike Chernov, Rob Engle, Shingo Goto, Amit Goyal, Campbell Harvey, David Hendry, Francis Longstaff, Nour Meddahi, Eric Renault, Matt Richardson, Neil Shephard, and seminar participants at Barclays Global Investors, CEMFI (Madrid), Emory University, GARP, ITAM (Mexico City), ISCTE (Lisbon), CIRANO-CIREQ Conference on Financial Econometrics (Montreal), Lehman Brothers, London School of Economics, Morgan Stanley, New York University, Oxford University, University of Cyprus, University of North Carolina, and University of Southern California for helpful comments. We especially thank an anonymous referee whose suggestions greatly improved the paper. Arthur Sinko provided able research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2004 by Eric Ghysels, Pedro Santa-Clara, and Rossen Valkanov. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Jump and Volatility Risk and Risk Premia: A New Model and Lessons from S&P 500 Options
Eric Ghysels, Pedro Santa-Clara, and Rossen Valkanov
NBER Working Paper No. 10913
November 2004
JEL No. G1

ABSTRACT

This paper studies the ICAPM intertemporal relation between the conditional mean and the conditional variance of the aggregate stock market return. We introduce a new estimator that forecasts monthly variance with past daily squared returns – the Mixed Data Sampling (or MIDAS) approach. Using MIDAS, we find that there is a significantly positive relation between risk and return in the stock market. This finding is robust in subsamples, to asymmetric specifications of the variance process, and to controlling for variables associated with the business cycle. We compare the MIDAS results with tests of the ICAPM based on alternative conditional variance specifications and explain the conflicting results in the literature. Finally, we offer new insights about the dynamics of conditional variance.

Eric Ghysels
University of North Carolina, Chapel Hill
Gardner Hall, CB 3305
Chapel Hill, NC 27599-3305
eghysels@unc.edu

Pedro Santa-Clara
Anderson School
UCLA
110 Westwood Plaza, Suite C4.21
Los Angeles, CA 90095-1481
and NBER
pedro.santa-clara@anderson.ucla.edu

Rossen Valkanov
Anderson School
UCLA
110 Westwood Plaza, Suite C4.21
Los Angeles, CA 90095-1481
rossen.valkanov@anderson.ucla.edu

1 Introduction

Merton’s (1973) ICAPM suggests that the conditional expected excess return on the stock market should vary positively with the market’s conditional variance:

$$E_t[R_{t+1}] = \mu + \gamma \text{Var}_t[R_{t+1}], \quad (1)$$

where γ is the coefficient of relative risk aversion of the representative agent and, according to the model, μ should be equal to zero. The expectation and the variance of the market excess return are conditional on the information available at the beginning of the return period, time t . This risk-return tradeoff is so fundamental in financial economics that it could well be described as the “first fundamental law of finance.”¹ Unfortunately, the tradeoff has been hard to find in the data. Previous estimates of the relation between risk and return often have been insignificant and sometimes even negative.

Baillie and DeGennaro (1990), French, Schwert, and Stambaugh (1987), and Campbell and Hentschel (1992) do find a positive albeit mostly insignificant relation between the conditional variance and the conditional expected return. In contrast, Campbell (1987) and Nelson (1991) find a significantly negative relation. Glosten, Jagannathan, and Runkle (1993), Harvey (2001), and Turner, Startz, and Nelson (1989) find both a positive and a negative relation depending on the method used.² The main difficulty in testing the ICAPM relation is that the conditional variance of the market is not observable and must be filtered from past returns.³ The conflicting findings of the above studies are mostly due to differences in the approach to modeling the conditional variance.

In this paper, we take a new look at the risk-return tradeoff by introducing a new estimator of the conditional variance. Our *Mixed Data Sampling*, or MIDAS, estimator forecasts the monthly variance with a weighted average of lagged daily squared returns. We use a flexible functional form to parameterize the weight given to each lagged daily squared

¹However, Abel (1988), Backus and Gregory (1993), and Gennotte and Marsh (1993) offer models where a *negative* relation between return and variance is consistent with equilibrium. Campbell (1993) discusses general conditions under which the risk-return relation holds as an approximation.

²See also Chan, Karolyi, and Stulz (1992), Lettau and Ludvigson (2002), Merton (1980), and Pindyck (1984). Goyal and Santa-Clara (2002) find a positive tradeoff between market return and *average stock variance*.

³We could think of using option implied volatilities as do Santa-Clara and Yan (2001) to make variance “observable.” Unfortunately, option prices are only available since the early 1980’s which is insufficient to reliably make inferences about the conditional mean of the stock market.

return and show that a parsimonious weighting scheme with only two parameters works quite well. We estimate the coefficients of the conditional variance process jointly with μ and γ from the expected return equation (1) with quasi-maximum likelihood.

Using monthly and daily market return data from 1928 to 2000 and, with MIDAS as a model of the conditional variance, we find a *positive and statistically significant* relation between risk and return. The estimate of γ is 2.6, which lines up well with economic intuition about a reasonable level of risk aversion. The MIDAS estimator explains about 40 percent of the variation of realized variance in the subsequent month and its explanatory power compares favorably to that of other models of conditional variance such as GARCH. The estimated weights on the lagged daily squared returns decay slowly, thus capturing the persistence in the conditional variance process. More impressive still is the fact that, in the ICAPM risk-return relation, the MIDAS estimator of conditional variance explains about two percent of the variation of next month's stock market returns (and five percent in the period since 1964). This is quite substantial given previous results about forecasting the stock market return.⁴ Finally, the above results are qualitatively similar when we split the sample into two subsamples of approximately equal sizes, 1928-1963 and 1964-2000.

To better understand MIDAS and its success in testing the ICAPM risk-return tradeoff, we compare our approach to previously used models of conditional variance. French, Schwert, and Stambaugh (1987) propose a simple and intuitive rolling window estimator of the monthly variance. They forecast monthly variance by the sum of daily squared returns in the previous month. Their method is similar to ours in that it uses daily returns to forecast monthly variance. However, when French, Schwert, and Stambaugh use that method to test the ICAPM, they find an insignificant (and sometimes negative) γ coefficient. We replicate their results but also find something rather interesting and new. When the length of the rolling window is increased from one month to three or four months, the magnitude of the estimated γ increases and the coefficient becomes statistically significant. This result nicely illustrates the point that the window length plays a crucial role in forecasting variances and detecting the tradeoff between risk and return. By optimally choosing the weights on lagged squared returns, MIDAS implicitly selects the optimal window size to estimate the variance, and that in turn allows us to find a significant risk-return tradeoff.

The ICAPM risk-return relation has also been tested using several variations of

⁴For instance, the forecasting power of the dividend yield for the market return does not exceed 1.5 percent (see Campbell, Lo, and MacKinlay (1997) and references therein).

GARCH-in-mean models. However, the evidence from that literature is inconclusive and sometimes conflicting. Using simple GARCH models, we confirm the finding of French, Schwert, and Stambaugh (1987) and Glosten, Jagannathan, and Runkle (1993), among others, of a positive but insignificant γ coefficient in the risk-return tradeoff. The absence of statistical significance comes both from GARCH's use of monthly return data in estimating the conditional variance process and the inflexibility of the parameterization. The use of daily data and the flexibility of the MIDAS estimator provides the power needed to find statistical significance in the risk-return tradeoff.

A comparison of the time series of conditional variance estimated according to MIDAS, GARCH, and rolling windows reveals that while the three estimators are correlated, there are some differences that affect their ability to forecast returns in the ICAPM relation. We find that the MIDAS variance process is more highly correlated with both the GARCH and the rolling windows estimates than these last two are with each other. This suggests that MIDAS combines some of the unique information contained in the other two estimators. We also find that MIDAS is particularly successful at forecasting realized variance both in high and low volatility regimes. These features explain the superior performance of MIDAS in finding a positive and significant risk-return relation.

It has long been recognized that volatility tends to react more to negative returns than to positive returns. Nelson (1991) and Engle and Ng (1993) show that GARCH models that incorporate this asymmetry perform better in forecasting the market variance. However, Glosten, Jagannathan, and Runkle (1993) show that when such asymmetric GARCH models are used in testing the risk-return tradeoff, the γ coefficient is estimated to be negative (sometimes significantly so). This stands in sharp contrast with the positive and insignificant γ obtained with symmetric GARCH models and remains a puzzle in empirical finance. To investigate this issue, we extend the MIDAS approach to capture asymmetries in the dynamics of conditional variance by allowing lagged positive and negative daily squared returns to have different weights in the estimator. Contrary to the asymmetric GARCH results, we still find a large positive estimate of γ that is statistically significant. This discrepancy between the asymmetric MIDAS and asymmetric GARCH tests of the ICAPM turns out to be quite interesting.

We find that what matters for the tests of the risk-return tradeoff is not so much the asymmetry in the conditional variance process but rather its persistence. In this respect, asymmetric GARCH and asymmetric MIDAS models prove to be very different. Consistent

with the GARCH literature, negative shocks have a larger immediate impact on the MIDAS conditional variance estimator than do positive shocks. However, we find that the impact of negative returns on variance is only temporary and lasts no more than one month. Positive returns, on the other hand, have an extremely persistent impact on the variance process. In other words, while short-term fluctuations in the conditional variance are mostly due to negative shocks, the persistence of the variance process is primarily driven by positive shocks. This is an intriguing finding about the dynamics of the variance process. Although asymmetric GARCH models allow for a different response of the conditional variance to positive and negative shocks, they constrain the persistence of both types of shocks to be the same. Since the asymmetric GARCH models “load” heavily on negative shocks and these have little persistence, the estimated conditional variance process shows little to no persistence.⁵ In contrast, by allowing positive and negative shocks to have different persistence, the asymmetric MIDAS model still obtains high persistence for the overall conditional variance process. Since only persistent variables can capture variation in expected returns, the difference in persistence between the asymmetric MIDAS and the asymmetric GARCH conditional variances explains their success and lack thereof in finding a risk-return tradeoff.

Campbell (1987) and Scruggs (1998) point out that the difficulty in measuring a positive risk-return relation may be due to misspecification of equation (1). Following Merton (1973), they argue that if changes in the investment opportunity set are captured by state variables in addition to the conditional variance itself, then those variables must be included in the equation of expected returns. In parallel, an extensive literature on the predictability of the stock market finds that variables that capture business cycle fluctuations are also good forecasters of market returns (see Campbell (1991), Campbell and Shiller (1988), Fama (1990), Fama and French (1988, 1989), Ferson and Harvey (1991), and Keim and Stambaugh (1986), among many others). We include business cycle variables together with both the symmetric and asymmetric MIDAS estimators of conditional variance in the ICAPM equation and find that the tradeoff between risk and return is virtually unchanged. Indeed, the explanatory power of the conditional variance for expected returns is orthogonal to the other predictive variables.

We conclude that the ICAPM is alive and well.

⁵The only exception is the two-component GARCH model of Engle and Lee (1999) who report findings similar to our asymmetric MIDAS model. They obtain persistent estimates of conditional variance while still capturing an asymmetric reaction of the conditional variance to positive and negative shocks.

The rest of the paper is structured as follows. Section 2 explains the MIDAS model and details the main results. Section 3 offers a comparison of MIDAS with rolling window and GARCH models of conditional variance. In Section 4, we discuss the asymmetric MIDAS model and use it to test the ICAPM. In Section 5, we include several often-used predictive variables as controls in the risk-return relation. Section 6 concludes.

2 MIDAS Tests of the Risk-Return Tradeoff

In this section, we introduce the *Mixed Data Sampling*, or MIDAS, estimator of conditional variance and use it to test the ICAPM relation between risk and return of the stock market.

2.1 Methodology

The MIDAS approach mixes daily and monthly data to estimate the conditional variance of the stock market. The returns on the left-hand side of equation (1) are measured at monthly intervals since, as argued below, higher frequency returns may be too noisy to use in a study of conditional means. On the right-hand side of the equation, we use daily returns in the variance estimator to exploit the advantages of high-frequency data in the estimation of second moments explained by the well-known continuous-record argument of Merton (1980).⁶ We allow the variance estimator to load on a large number of past daily squared returns with optimally chosen weights.

The MIDAS estimator of the conditional variance of *monthly* returns, $\text{Var}_t[R_{t+1}]$, is based on prior *daily* squared return data:

$$V_t^{\text{MIDAS}} = 22 \sum_{d=0}^{\infty} w_d r_{t-d}^2 \quad (2)$$

where w_d is the weight given to the squared return of day $t - d$. We use the lower case r to denote *daily* returns, which should be distinguished from the upper case R used for *monthly* returns; the corresponding subscript $t - d$ stands for the date t minus d days. R_{t+1}

⁶Recently, several authors, including Andersen, Bollerslev, Diebold, and Ebens (2001), Andreou and Ghysels (2002), Barndorff-Nielsen and Shephard (2002), and Taylor and Xu (1997) suggest various methods using high-frequency data to estimate variances. Alizadeh, Brandt, and Diebold (2002) propose an alternative measure of realized variance using the daily range of the stock index.

is the monthly return from date t to date $t + 1$ and r_{t-d} is the daily return d days before date t . Although this notation is slightly ambiguous, it has the virtue of not being overly cumbersome. With weights that sum up to one, the factor 22 ensures that the variance is expressed in monthly units since there are typically 22 trading days in a month.

We postulate a flexible form for the weight given to the squared return on day $t - d$:

$$w_d(\kappa_1, \kappa_2) = \frac{\exp\{\kappa_1 d + \kappa_2 d^2\}}{\sum_{i=0}^{\infty} \exp\{\kappa_1 i + \kappa_2 i^2\}}. \quad (3)$$

This scheme has several advantages. First, it guarantees that the weights are positive which in turn ensures that the conditional variance in (2) is also positive. Second, the weights add up to one. Third, the functional form in (3) can produce a wide variety of shapes for different values of the two parameters. Fourth, the specification is parsimonious, with only two parameters to estimate. Fifth, 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 variance. Finally, we can increase the order of the polynomial in (3) or consider other functional forms. For instance, all the results shown below are robust to parameterizing the weights as a Beta function instead of the exponential form in (3).⁷ As a practical matter, the infinite sum in (2) and (3) needs to be truncated at a finite lag. In all the results that follow, we use 252 days (which corresponds to roughly one year of trading days) as the maximum lag length. The results are not sensitive to increasing the maximum lag length beyond one year.

The weights of the MIDAS estimator implicitly capture the dynamics of the conditional variance. A larger weight on distant past returns induces more persistence on the variance process. The weighting function also determines the statistical precision of the estimator by controlling the amount of data used to estimate the conditional variance. When the function decays slowly, a large number of observations effectively enter in the forecast of the variance and the measurement error is low. Conversely, a fast decay corresponds to using a small number of daily returns to forecast the variance with potentially large measurement error. To some extent, there is a tension between capturing the dynamics of variance and minimizing measurement error. Since variance changes through time, we would like to use more recent observations to forecast the level of variance in the next month. However, to

⁷See Ghysels, Santa-Clara, and Valkanov (2003) for a general discussion of the functional form of the weights.

the extent that measuring variance precisely requires a large number of daily observations, the estimator may still place significant weight on more distant observations. The Appendix offers a more formal treatment of the MIDAS estimator.

To estimate the parameters in the weight function, we maximize the likelihood of monthly returns. We use the variance estimator (2) with the weight function (3) in the ICAPM relation (1) and estimate the parameters κ_1 and κ_2 jointly with μ and γ by maximizing the likelihood function, assuming that the conditional distribution of returns is normal:⁸

$$R_{t+1} \sim N(\mu + \gamma V_t^{\text{MIDAS}}, V_t^{\text{MIDAS}}). \quad (4)$$

In this way, the conditional mean and the conditional variance of the monthly return in April (from the close of the last day of March to the close of the last day of April) depends on daily returns up to the last day of March. Since the true conditional distribution of returns may depart from normality, our estimator really is only quasi-maximum likelihood. The parameter estimators are nevertheless consistent and asymptotically normally distributed. Their covariance matrix is estimated using the Bollerslev and Wooldridge (1992) approach to account for heteroscedasticity.⁹

We have thus far used monthly returns as a proxy for expected returns in equation (1) and daily returns in the construction of the conditional variance estimator. However, using higher frequency returns at, say, weekly or daily intervals may improve the estimate of γ because of the availability of additional data points. Alternatively, it may be argued that quarterly returns increase the efficiency of the estimator of γ because they are less volatile. A general analytical argument is difficult to formulate without making additional assumptions about the data generating process. Similarly, the returns used to forecast volatility can be sampled at different frequencies from intra-daily to weekly or even monthly observations. Fortunately, the MIDAS approach can easily be implemented at different frequencies on the left-hand and on the right-hand side. This can be achieved with the same parametric specification and with the same number of parameters. Hence, we can directly compare the estimates of γ and their statistical significance across different frequencies.

⁸Alternatively, we could use GMM for more flexibility in the relative weighting of the conditional moments in the objective function.

⁹More specifically, using Theorem 2.1 in Bollerslev and Wooldridge (1992), we compute the covariance matrix of the parameter estimates as $A_T^{-1} B_T A_T^{-1} / T$, where A_T^{-1} is an estimate of the Hessian matrix of the likelihood function and B_T is an estimate of the outer product of the gradient vector with itself.

2.2 Empirical Analysis

We estimate the ICAPM with the MIDAS approach using excess returns on the stock market from January of 1928 to December of 2000. We use the CRSP value-weighted portfolio as a proxy for the stock market and the yield of the three-month Treasury bill as the risk-free interest rate. Daily market returns are obtained from CRSP for the period July of 1962 to December of 2000, and from William G. Schwert's website for the period January of 1928 to June of 1962 (see Schwert (1990a) for a description of those data). The daily risk-free rate, obtained from Ibbotson Associates, is constructed by assuming that the Treasury bill rates stay constant within the month and suitably compounding them. Monthly excess returns are obtained by compounding the daily excess returns. In what follows, we refer to excess returns simply as returns.

Table 1 displays summary statistics for the monthly returns and the monthly realized variance of returns computed from within-month daily data (as explained in equation (5) below). We show the summary statistics for the full 1928-2000 sample and, for robustness, we also analyze two subsamples of approximately equal length, 1928 to 1963 and 1964 to 2000.

The monthly market return has a mean of 0.649 percent and a standard deviation of 5.667 percent (variance of 0.321×10^2).¹⁰ Returns are negatively skewed and slightly leptokurtic. The first order autoregressive coefficient of monthly returns is 0.068. The average market return during 1928-1963 is considerably higher than that observed during 1964-2000. The variance of monthly returns is also higher in the first subsample. Both subsamples exhibit negative skewness and high kurtosis. The realized variance has a mean of 0.262 in the overall sample, which closely matches the variance of monthly returns (the small difference between the two numbers is due to Jensen's inequality). The mean of the variance in the first subsample is much higher than in the second, mostly due to the period of the Great Depression. The realized variance process displays considerable persistence, with an autoregressive coefficient of 0.608 in the entire sample. Again, the first subsample shows more persistence in the variance process. As expected, realized variance is highly skewed and leptokurtic. The results from these summary statistics are well-known in the empirical finance literature.

¹⁰This and later tables report variances rather than more customary standard deviations because the risk-return tradeoff postulates a relation between returns and their variance, not their standard deviation.

Table 2 contains the main result of the paper, the estimation of the risk-return tradeoff equation with the MIDAS conditional variance. The estimated ICAPM coefficient γ is 2.606 in the full sample, with a highly significant t -statistic (corrected for heteroscedasticity with the Bollerslev and Wooldridge method) of 6.710. Most importantly, the *magnitude* of γ lines up well with the theory. According to the ICAPM, γ is the coefficient of relative risk aversion of the representative investor and a risk aversion coefficient of 2.606 matches a variety of empirical studies (see Hall (1988) and references therein). The significance of γ is robust in the subsamples, with estimated values of 1.547 and 3.748, and t -statistics of 3.382 and 8.612. These results are consistent with Mayfield (2003) who uses a regime-switching model for conditional volatility and finds that the risk-return tradeoff holds within volatility regimes. The estimated magnitude and significance of the γ coefficient in the ICAPM relation are remarkable in light of the ambiguity of previous results. The intercept μ is always significant, which, in the framework of the ICAPM, may capture compensation for covariance of the market return with other state variables (which we address in section 5) or compensation for jump risk (see Pan (2002) and Chernov, Gallant, Ghysels, and Tauchen (2002)).

Table 2 also reports the estimated parameters of the MIDAS weight function (3). Both coefficients are statistically significant in the full sample and the subsamples. Furthermore, a likelihood ratio test of their joint significance, $\kappa_1 = \kappa_2 = 0$, has a p -value smaller than 0.001. Since the restriction $\kappa_1 = \kappa_2 = 0$ corresponds to placing equal weights on all lagged squared daily returns, we conclude that the estimated weight function is statistically different from a simple equally-weighted scheme. We cannot interpret the magnitudes of the coefficients κ_1 and κ_2 individually but only jointly in the weighting function (3). In Figure 1, we plot the estimated weights, $w_d(\kappa_1, \kappa_2)$, of the conditional variance on the lagged daily squared returns for the full sample and the subsamples. In all cases, we observe that the weights are a slowly declining function of the lag length. For example, only 31 percent of the weight is placed on the first lagged month of daily data (22 days), 56 percent on the first two months, and it takes more than three months for the cumulative weight to reach 75 percent. The weight profiles for the subsamples are very similar. We conclude that it takes a substantial amount of daily return data to accurately forecast the variance of the stock market. This result stands in sharp contrast to the common view that one month of daily returns is sufficient to reliably estimate the variance.

To assess the predictive power of the MIDAS variance for the market return we run

a regression of the realized return in month $t + 1$, R_{t+1} , on the forecasted variance for that month, V_t^{MIDAS} . The coefficient of determination for the regression using the entire sample, R_R^2 , is 1.9 percent, which is a reasonably high value for a predictive regression of returns at monthly frequency. This coefficient increases to 5.0 percent in the second subsample.

We also examine the ability of the MIDAS estimator to forecast realized variance. We estimate realized variance from within-month daily returns as:

$$\sigma_{t+1}^2 = \sum_{d=0}^{22} r_{t+1-d}^2. \quad (5)$$

Table 2 reports the coefficient of determination, $R_{\sigma^2}^2$, from regressing the realized variance, σ_{t+1}^2 , on the MIDAS forecasted variance, V_t^{MIDAS} . MIDAS explains over 40 percent of the fluctuations of the realized variance in the entire sample. Given that σ_{t+1}^2 in (5) is only a noisy proxy for the true variance in the month, the $R_{\sigma^2}^2$ obtained is impressively high.¹¹ The value of $R_{\sigma^2}^2$ in the second subsample is only 0.082, due to the crash of 1987. If we eliminate the 1987 crash from the second subsample, the $R_{\sigma^2}^2$ jumps to 0.283. Figure 2 displays the realized variance together with the MIDAS forecast for the entire sample. We see that the estimator does a remarkable job of forecasting next month's variance.

Thus far we have estimated γ in a MIDAS regression of monthly returns on variance estimated from daily returns. However, this is not the only possible frequency choice. With higher frequency data on the left-hand side, we have more observations, but also more noise in the returns. With lower frequency data, we have a better estimate of expected returns, but fewer observations. We now investigate what return horizon in the left-hand side of the MIDAS regression yields the most precise estimates of the risk-return tradeoff. Table 3 presents estimates of γ in the ICAPM regression of returns at daily, weekly, monthly, bi-monthly, and quarterly horizons on the MIDAS conditional variance, estimated with daily squared returns. We find that the estimates of γ range from 1.964 to 2.880 as we vary the frequency of returns. The t -statistics of γ increase systematically from 1.154 at daily frequency to 6.710 at monthly frequency. The standard error of the estimates does not change much across horizons, so the improvement in the t -statistics is mostly due to the higher point estimate of γ at the monthly horizon.

¹¹ Andersen and Bollerslev (1998) and Andersen, Bollerslev, and Meddahi (2002) show that the maximum R^2 obtainable in a regression of this type is much lower than 100 percent, often on the order of 40 percent. The high standard deviation of the realized variance and the relatively low persistence of the process, shown in Table 1, indicate a high degree of measurement error.

A similar pattern emerges from the goodness-of-fit measures R_R^2 and $R_{\sigma_2}^2$ in Table 3. The use of high-frequency data as a proxy for the conditional mean of returns decreases the ability of the MIDAS estimator to forecast realized variance. The $R_{\sigma_2}^2$ at daily and weekly horizons are only 0.059 and 0.119. At monthly, bi-monthly, and quarterly horizons, they are markedly higher at 0.407, 0.309, and 0.329. Of course, the realized variances at daily or weekly frequency are a very noisy measure of the true variance since they are estimated with only one or five daily returns. The subsamples in Table 3 yield similar results. We conclude that the choice of monthly frequency strikes the best balance between sample size and signal-to-noise ratio. Hence, in the subsequent analysis, we only use monthly returns on the left-hand side of our MIDAS models.

3 Why MIDAS Works: Comparison with Other Tests

To understand why tests based on the MIDAS approach support the ICAPM when the extant literature offers conflicting results, we compare the MIDAS estimator with previously used estimators of conditional variance. We focus our attention on rolling window and GARCH estimators of conditional variance. For conciseness, we report results for the entire sample, but the conclusions also hold in the subsamples.

3.1 Rolling Window Tests

As an example of the rolling window approach, French, Schwert, and Stambaugh (1987) use within-month daily squared returns to forecast next month's variance:

$$V_t^{\text{RW}} = 22 \sum_{d=0}^D \frac{1}{D} r_{t-d}^2 \quad (6)$$

where D is the number of days used in the estimation of variance.¹² Again, daily squared returns are multiplied by 22 to measure the variance in monthly units. French, Schwert, and

¹²French, Schwert, and Stambaugh (1987) include a correction for serial correlation in the returns that we ignore for now. We follow their example and do not adjust the measure of variance by the squared mean return as this is likely to have only a minor impact with daily data. Additionally, French, Schwert, and Stambaugh actually use the fitted value of an ARMA process for the one-month rolling window estimator to model the conditional variance.

Stambaugh choose the window size to be one month, or $D = 22$. Besides its simplicity, this approach has a number of advantages. First, as with the MIDAS approach, the use of daily data increases the precision of the variance estimator. Second, the stock market variance is very persistent (see Officer (1973) and Schwert (1989)), so the realized variance on a given month ought to be a good forecast of next month's variance.

However, it is not clear that we should confine ourselves to using data from the last month only to estimate the conditional variance. We may want to use a larger window size D in equation (6), corresponding to more than one month's worth of daily data. Interestingly, this choice has a large impact on the estimate of γ .

We estimate the parameters μ and γ of the risk-return tradeoff (1) with maximum likelihood using the rolling window estimator (6) for the conditional variance. Table 4 reports the estimates of the risk-return tradeoff for different sizes D of the window used to estimate the conditional variance. The first line corresponds to using daily data from the previous month only so the measure of V_t^{RW} is similar to the one reported in French, Schwert, and Stambaugh (1987). The estimate of γ is 0.546 and statistically insignificant. In their study, French, Schwert, and Stambaugh (1987) estimate a γ of -0.349, also insignificant. The difference between the estimates is due to the difference in sample periods. When we use their sample period from 1928 to 1984, we obtain the same results as French, Schwert, and Stambaugh (1987).

As we increase the window size to two through four months, the magnitude of γ increases and becomes significant, with a higher R_R^2 . When the rolling window includes four months of data, the estimated γ coefficient is 2.149 and statistically significant.¹³ This coefficient is very similar to the estimated γ with the MIDAS approach, only the level of significance is lower. Finally, as the window size increases beyond four months, the magnitude of the estimated γ decreases as does the likelihood value. This suggests that there is an optimal window size to estimate the risk-return tradeoff.

These results are striking. They confirm our MIDAS finding, namely, that there is a positive and significant tradeoff between risk and return. Indeed, the rolling window approach can be thought of as a robust check of the MIDAS regressions since it is such a simple estimator of conditional variance with no parameters to estimate. Moreover, Table 4 helps us reconcile the MIDAS results with the findings of French, Schwert, and Stambaugh

¹³These findings are consistent with Brandt and Kang (2003), and Whitelaw (1994) who report a lagged relation between the conditional variance and the conditional mean.

(1987). That paper missed out on the tradeoff by using too small a window size (one month) to estimate the variance. One month’s worth of daily data simply is not enough to reliably estimate the conditional variance and to measure its impact on expected returns.

The maximum likelihood across window sizes is obtained with a four-month window. This window size implies a constant weight of 0.011 in the lagged daily squared returns of the previous four months. Of the different window lengths we analyze, these weights are closest to the optimal MIDAS weights shown in Figure 1, which puts roughly 80 percent of the weight in those first four months of past daily squared returns.

The rolling window estimator is similar to MIDAS in its use of daily squared returns to forecast monthly variance. But it differs from MIDAS in that it constrains the weights to be constant and inversely proportional to the window length. This constraint on the weights affects the performance of the rolling window estimator compared to MIDAS. For instance, the rolling window estimator does not perform as well as the MIDAS estimator in forecasting realized returns or realized variance. The coefficient of determination for realized returns is 1.2 percent compared to 1.9 percent for MIDAS, and for realized variance it is 38.4 percent which is lower than the 40.7 percent obtained with MIDAS. A more detailed comparison of the forecasts from the rolling window and the MIDAS estimators is provided below.

3.2 GARCH Tests

The most popular approach to study the ICAPM risk-return relation has been with GARCH-in-mean models estimated with *monthly* return data (see Engle, Lilien, and Robins (1987), French, Schwert, and Stambaugh (1987), Campbell and Hentschel (1992), Glosten, Jagannathan, and Runkle (1993), among others). The simplest model in this family can be written as:

$$V_t^{\text{GARCH}} = \omega + \alpha \epsilon_t^2 + \beta V_{t-1}^{\text{GARCH}} \quad (7)$$

where $\epsilon_t = R_t - \mu - \gamma V_{t-1}^{\text{GARCH}}$. The squared innovations ϵ_t^2 in the variance estimator play a role similar to the monthly squared return in the MIDAS or rolling window approaches and, numerically, they are very similar (since the squared average return is an order of magnitude smaller than the average of squared returns). For robustness, we also estimate an absolute GARCH model, ABS-GARCH:

$$(V_t^{\text{ABSGARCH}})^{1/2} = \omega + \alpha |\epsilon_t| + \beta (V_{t-1}^{\text{ABSGARCH}})^{1/2}. \quad (8)$$

Note that the GARCH model (7) can be rewritten as:

$$V_t^{\text{GARCH}} = \frac{\omega}{1 - \beta} + \alpha \sum_{i=0}^{\infty} \beta^i \epsilon_{t-i}^2. \quad (9)$$

The GARCH conditional variance model is thus approximately a weighted average of past monthly squared returns. Compared to MIDAS, the GARCH model uses monthly rather than daily squared returns. Moreover, the functional form of the weights implied by the dynamics of variance in GARCH models exhibits less flexibility than the MIDAS weighting function. Indeed, even though the GARCH process is defined by three parameters, the shape of the weight function depends exclusively on β . This shape is similar to MIDAS when the parameter κ_2 is set to zero.

Table 5 shows the coefficient estimates of the GARCH and the ABS-GARCH models, estimated with quasi-maximum likelihood. Both models yield similar results, so we concentrate on the simple GARCH case. For that model, the estimate of γ is 1.060 and insignificant, with a t -statistic of 1.292 (obtained using Bollerslev and Wooldridge standard errors). French, Schwert, and Stambaugh (1987) obtain a higher estimate for γ of 7.809 in a different sample, but they also find it to be statistically insignificant. Using a symmetric GARCH model, Glosten, Jagannathan, and Runkle (1993) estimate γ to be 5.926 and again insignificant. In similar sample periods, we replicate the findings of these studies. As a further robustness check, we estimate higher order GARCH(p,q) models (not shown for brevity), with $p = 1, \dots, 3$ and $q = 1, \dots, 3$, and obtain estimates of γ that are comparable in magnitude and still insignificant. In sum, although GARCH models find a positive estimate of γ , they lack the power to find statistical significance for the coefficient. Also, the coefficients of determination from predicting returns, R_R^2 , and realized variances, $R_{\sigma^2}^2$, are 0.5 and 35.9 percent for the GARCH model, and appear low when compared with the coefficients of 1.9 and 40.7 percent obtained with MIDAS.

The success of MIDAS relative to GARCH in finding a significant risk-return tradeoff resides in the extra power that mixed-data frequency regressions obtain from the use of daily data in the conditional variance estimator. Put differently, MIDAS has more power than GARCH because it estimates two rather than three parameters and uses a lot more observations to do it. Also, relative to GARCH, MIDAS has a more flexible functional form for the weights on past squared returns. The interplay of mixed-frequency data and flexible weights explain the higher estimates of γ and the higher t -statistics obtained by MIDAS. In

section 3.4 we will come back to this comparison in more detail.

3.3 Comparison of Filtered Variance Processes

To further understand the similarities and differences between MIDAS, GARCH, and rolling window estimators, we turn our attention to the filtered time series of conditional variance produced by each of the three approaches. For the rolling window estimator, we use a window length of one month which is similar to what has been used in the literature. Panel A of Table 6 presents summary statistics of the three conditional variance processes. The GARCH forecast is the most persistent with an AR(1) coefficient of 0.970, has the highest mean (0.325), and the lowest variance (0.187). The rolling window forecast is the least persistent (AR(1) of 0.608), has a much lower mean (0.262), and the highest variance (0.323). The high variance and low persistence is partly due to this estimator's high measurement error. The high mean of the GARCH variance relative to the realized variance (which has the same mean as the rolling windows) indicates that GARCH has some bias. With an AR(1) of 0.872, the persistence of MIDAS conditional variance is between that of the GARCH and the rolling windows approaches. MIDAS variance has a mean of 0.256 which is very similar to the rolling windows mean and is lower than the GARCH mean. Finally, the variance of the MIDAS conditional variance is between that of GARCH and of rolling windows.

The difference between MIDAS, GARCH, and rolling windows is also apparent from a plot of the time series of their (in-sample) forecasted variances displayed against the realized volatility in Figure 2. In the top graph, the MIDAS forecasts (solid line) and the realized variance (thin dotted line) are very similar. In particular, MIDAS is successful at capturing periods of extreme volatility such as during the first twenty years of the sample and around the crash of 1987. GARCH forecasts, shown in the middle graph (again in solid line), are smoother than realized variance. This is not surprising since GARCH uses only data at monthly frequency. More importantly, in periods of relatively low volatility, GARCH forecasts are higher than the realized variance. This translates into higher unconditional means of filtered GARCH variances, as observed in Table 6. Finally, the variances filtered with rolling windows, shown in the bottom graph, are the shifted values of the realized variance. From visual inspection of the time series of the conditional variance processes, MIDAS produces the best forecasts of realized volatility.

As a more systematic way of analyzing the differences between realized variance and

the filtered series, we show in Figure 3 scatterplots of realized variance against forecasted variances. The scatterplots are displayed in log-log scale to facilitate comparison of the series during periods of low and high volatility periods. If a model fits the realized variances well, we expect a tight clustering of points around the 45 degree line. In the top graph, the MIDAS forecasts do plot closely to the realized variance observations. While there are some outliers on both sides of the 45 degree line, there are no discernible asymmetries. In contrast, GARCH forecasts, shown in the middle graph, are systematically higher than realized variance at the low end of the variance scale (between 10^{-4} and 10^{-3}), while the fit at the high end of the scale is no better than MIDAS. This is yet another manifestation of the finding that GARCH forecasts have higher mean and are too smooth when compared to the realizations of the variance process. Finally, the bottom scatterplot displays the realized variance plotted against the rolling window forecasts. There are no systematic biases, but the scatterplot is much more dispersed when compared to the MIDAS and GARCH plots. This is true for all variances, but is especially evident at the high end of the variance scale (between 10^{-2} and 10^{-1}).

We now examine in more detail the dynamics of the three estimators of conditional variance. Previously, we argued that the MIDAS weights implicitly determine the dynamic behavior of the monthly filtered variance. The MIDAS weights in Figure 1 suggest that the estimated volatility process is persistent and the time series plotted in Figure 2 confirms that intuition. It is instructive to analyze the dynamics of V_t^{MIDAS} in the framework of ARMA(p,q) models. A theoretical correspondence between the weight function and the ARMA(p,q) parameters is difficult to derive largely because of the mixed-frequency nature of the problem. Instead, we pursue a data-driven approach. Using the filtered time series of MIDAS conditional variance, we estimate $\Phi(L)V_t^{\text{MIDAS}} = \Psi(L)e_t$, where $\Phi(L) = 1 - \phi_1L - \phi_2L^2 \dots - \phi_pL^p$ and $\Psi(L) = 1 - \psi_1L - \psi_2L^2 \dots - \psi_qL^q$. We study all combinations of $p = 1, \dots, 12$ and $q = 0, 1, \dots, 12$.

In the AR(1) case, we obtain an estimate of $\phi_1 = 0.872$. In general, for the purely autoregressive ARMA($p,0$) models, the persistence of the process is captured by the highest autoregressive root of the corresponding polynomial. In the AR(2), AR(3), and AR(4) cases, the highest autoregressive roots are 0.853, 0.859, and 0.864, respectively, which are comparable to the estimate of ϕ_1 in the AR(1) case. We choose the best-fitting ARMA(p,q) model using the Akaike Information Criterion (AIC) and the Schwartz Criterion (SC) which not only maximize fit but also penalize for the number of estimated parameters. The

AIC and SC select an ARMA(7,5) and an ARMA(7,3), respectively, as the models that best fit V_t^{MIDAS} .¹⁴ It is remarkable that MIDAS can generate such rich dynamics for the conditional variance process from a very parsimonious representation of the weight function. For comparison, the realized variance process is best approximated by an ARMA (5,6) (selected by both the AIC and the SC). The ARMA process that best captures the dynamics of the conditional variance filtered with GARCH is a simple AR(1). We conclude that MIDAS approximates the dynamic structure of realized variance better than GARCH. The rolling window estimator trivially inherits the dynamics of the realized variance process.

In Panel B of Table 6, we investigate whether the filtered conditional variances can adequately capture fluctuations in the realized variances. If a forecasted variance approximates closely the true conditional variance, then the standardized residuals from the risk-return tradeoff should be approximately standard normally distributed (with a mean of zero and variance of unity). We take the demeaned monthly returns and divide them by the square root of the forecasted variance according to each of the methods. We find that the standardized residuals using the MIDAS approach are the closest to standard normality. Their variance, skewness and kurtosis are closer to one, zero, and three, respectively, than with the other two methods. They are still skewed and leptokurtic but much less so than using rolling windows and GARCH.

The above statistics give us a good idea of the statistical properties of the filtered variances. However, since the time series properties of the filtered series are different, it is not clear which one of the three methods provides the most accurate forecasts (in a MSE sense). To judge the forecasting power of the three methods, we compute a goodness-of-fit measure which is defined as one minus the sum of squared forecasting errors (i.e., the sum of squared differences between forecasted variance and realized variance) divided by the total sum of squared realized variance. This goodness-of-fit statistic measures the forecasting power of each method for the realized variance.¹⁵ The goodness-of-fit statistics are shown

¹⁴The estimated autoregressive parameters are 1.007, -0.050, -0.711, 0.918, -0.220, -0.030, 0.014 and the moving average parameters are 0.089, 0.255, -0.577, 0.360, 0.257 in the ARMA(7,5) case. For the ARMA(7,3), the autoregressive parameters are 1.002, -0.045, -0.710, 0.904, -0.221, -0.032, 0.017 and the moving average parameters are 0.112, 0.313, -0.409.

¹⁵This measure is similar to the previously used $R_{\sigma_2}^2$ of a regression of realized variance on the forecasted variance. The only difference is that now the intercept of the regression is constrained to be zero and the slope equal to one. It measures the total forecasting error, rather than the correlation between realized variance and forecasted variance. It is not enough for a forecast to be highly correlated with the realized variance; its level must also be on target. For instance, a forecast that always predicts twice the realized variance would have an R^2 of one in a regression but would have a modest goodness-of-fit value.

in the last column of Table 6, Panel B. MIDAS produces the most accurate forecasts with a goodness-of-fit measure of 0.494. For comparison, the goodness-of-fit of GARCH is 0.440, while that of rolling windows is 0.354.

Panel C of Table 6 presents the correlation matrix of the MIDAS, GARCH, rolling window, and realized variance series. MIDAS correlates highly with GARCH and rolling windows, 0.802 and 0.798, respectively. In contrast, the correlation between GARCH and rolling windows is only 0.660. The correlation of the three forecasts with the realized volatility is shown as a reference point. Not surprisingly, realized variance has the highest correlation of 0.638 with the MIDAS forecasts, as the squared correlation is identical to the $R_{\sigma^2}^2$ in Tables 2, 4, and 5. This evidence, in conjunction with the statistics in Panels A and B, suggests that MIDAS combines the information of GARCH and rolling windows and that each of these individually has less information than MIDAS.

The high volatility of rolling windows compared to the other methods suggests that it is a noisy measure of conditional variance. Similarly, rolling windows displays little persistence, which is also likely due to measurement error. These two related problems hinder the performance of this estimator in the risk-return tradeoff. Indeed, the errors-in-variables problem will bias downward the slope coefficient and lower the corresponding t -statistic in the regression of monthly returns on the rolling windows conditional variance. The GARCH estimator does not suffer from either of these problems. However, it does show a bias as a forecaster of realized variance, especially in periods of low volatility. Additionally, the filtered variance process from GARCH is too smooth when compared to the other estimators and the realized variance. These problems undoubtedly affect the ability of GARCH to explain the conditional mean of returns. The MIDAS estimator has better properties than GARCH and rolling windows: it is unbiased both in high and low volatility regimes, displays little estimation noise, and is highly persistent. These properties make it a good explanatory variable for expected returns.

3.4 Mixed Frequencies and Flexible Weights

Thus far, we have found a positive and significant risk-return tradeoff with the MIDAS estimator that cannot be obtained with either rolling windows or GARCH. The MIDAS tests have two important features: they use mixed-frequency data and the weights of forecasted variance on past squared returns are parameterized with a flexible functional form. This

raises the question of whether one of the two features is predominantly responsible for the power of the MIDAS tests or whether they interact in a particularly favorable fashion. To answer this question, we run two comparisons. First, to isolate the effect of the weight function, we compare MIDAS with GARCH estimated with mixed-frequency data. Second, we study the impact of using mixed-frequency data by comparing monthly GARCH with MIDAS estimated from monthly data alone.

To assess the importance of flexibility in the functional form of the weights, we compare the MIDAS results with GARCH estimated with mixed-frequency data. To estimate the mixed-frequency GARCH, we assume that *daily* variance follows a GARCH(1,1) process as in equation (7). At any point in time, this process implies forecasts for the daily variance multiple days into the future. Summing the forecasted variances over the following 22 days yields a forecast of next month’s variance.¹⁶ We can then jointly estimate the coefficients of the daily GARCH and the parameter γ by quasi-maximum likelihood using monthly returns and the forecast of monthly variance together in the density (4).¹⁷

The first row of Table 7 displays the tests of the risk-return tradeoff using this mixed-frequency GARCH process. For comparison, we reproduce the results of the MIDAS test from Table 2 which is estimated with the same mixed-frequency data. The estimate of γ using the mixed-frequency GARCH estimator is still low at 0.431 and insignificant, with a t -statistic of 0.592, which compare poorly with the MIDAS estimate of 2.606 and t -statistic of 6.710. The estimator has low explanatory power for monthly returns, with an R_R^2 of 0.3 percent (1.9 percent for MIDAS), and low explanatory power for future realized variance, with an $R_{\sigma_2}^2$ of 29.1 percent (40.7 percent for MIDAS). These results point to the importance of having a flexible functional form for the weights on past daily squared returns. Indeed, the only difference between the MIDAS and the mixed-frequency GARCH estimator is the shape of the weight function. Figure 4 plots the weights of the two estimators (plotted as a solid line and labeled “daily MIDAS” and “daily GARCH”) on past daily squared returns.

¹⁶The one-month ahead forecast of the variance in GARCH(1,1) is:

$$\sum_{d=1}^{22} (\alpha + \beta)^d V_t^{\text{GARCH}} + \frac{(1 - (\alpha + \beta)^d)\omega}{1 - \alpha - \beta}.$$

¹⁷We also tried a two-step procedure whereby we first estimate a daily GARCH model (not GARCH in mean) and then run a regression of monthly returns on the forecasted monthly variance from the daily GARCH. The results are similar, albeit slightly less significant, to those from the procedure described above and are not reported.

The decay of the daily GARCH weights is much faster than in the corresponding MIDAS model. In other words, the persistence of the estimated GARCH variance process is lower than that of MIDAS. The first-order serial correlation of the monthly variance estimated from daily GARCH is 0.781, which is considerably less than the 0.872 serial correlation of the MIDAS variance. It is also interesting to note that the daily GARCH estimator performs worse than the previously studied monthly GARCH, with statistics also reported in the table for comparison (reproduced from Table 5).

To analyze the gains from mixing frequencies, we compare the daily MIDAS and daily GARCH results with the same models estimated with monthly (not mixed-frequency) data. We define a MIDAS variance estimator using only monthly data by:

$$V_t^{\text{MIDAS}} = \sum_{m=1}^{\infty} w_m R_{t-m}^2 \quad (10)$$

where the functional form of the weights on lagged monthly squared returns is still given by (3).¹⁸ Although this estimator no longer uses mixed frequency data, we still refer to it as a MIDAS estimator. The second row of Table 7 shows the tests of the risk-return tradeoff with the monthly GARCH and monthly MIDAS estimators. We see that the monthly MIDAS estimator performs rather well, with an estimate of γ of 2.553 and a t -statistic of 2.668. The major difference relative to the daily MIDAS model is the significance of the γ coefficient (the t -statistic drops from 6.710 to 2.668) and the lower explanatory power for monthly returns (R_R^2 drops from 1.9 to 1.1 percent) and future realized variance ($R_{\sigma^2}^2$ drops from 40.7 to 38.2 percent). We conclude that using mixed-frequency data increases the power of the risk-return tradeoff tests. The first panel of Figure 4 compares the weights placed by monthly MIDAS on lagged returns (shown as a step function with the weights constant within each month) with the daily MIDAS weights. There is little difference between the two weight functions which translates into similar persistence of the corresponding variance processes (AR(1) coefficients of 0.893 and 0.872 respectively). Finally, we see that the tests using monthly MIDAS dominate the monthly GARCH tests. The estimate of γ and its t -statistic are more than twice as large. The forecasting power of the monthly MIDAS variance for returns and realized variance is also higher.

We conclude that the power of the MIDAS tests to uncover a tradeoff between risk and return in the stock market comes both from the flexible shape of the weight function

¹⁸Again, for practical purposes we truncate the infinite sum at one year lag.

and the use of mixed-frequency returns in the test.

4 Asymmetries in the Conditional Variance

In this section, we present a simple extension of the MIDAS specification that allows positive and negative returns to have not only an asymmetric impact on the conditional variance, but also to exhibit different persistence. We compare the asymmetric MIDAS model to previously used asymmetric GARCH models in tests of the ICAPM. Our results clarify the puzzling findings in the literature.

4.1 Asymmetric MIDAS Tests

It has long been recognized that volatility is persistent and increases more following negative shocks than positive shocks.¹⁹ Using asymmetric GARCH models, Nelson (1991) and Engle and Ng (1993) confirm that volatility reacts asymmetrically to positive and negative return shocks. Glosten, Jagannathan, and Runkle (1993) use an asymmetric GARCH-in-mean formulation to capture the differential impact of negative and positive lagged returns on the conditional variance and use it to test the relation between the conditional mean and the conditional variance of returns.²⁰ They find that the sign of the tradeoff changes from insignificantly positive to significantly negative when asymmetries are included in GARCH models of the conditional variance. This result is quite puzzling and below we explain its provenance.

To examine whether the risk-return tradeoff is robust to the inclusion of asymmetric effects in the conditional variance, we introduce the asymmetric MIDAS estimator:

$$V_t^{\text{ASYMIDAS}} = 22 \left[\phi \sum_{d=0}^{\infty} w_d(\kappa_1^-, \kappa_2^-) \mathbf{1}_{t-d}^- r_{t-d}^2 + (2 - \phi) \sum_{d=0}^{\infty} w_d(\kappa_1^+, \kappa_2^+) \mathbf{1}_{t-d}^+ r_{t-d}^2 \right] \quad (11)$$

¹⁹This is the so-called “feedback effect,” based on the time-variability of the risk-premium induced by changes in variance. See French, Schwert, and Stambaugh (1987), Pindyck (1984) and Campbell and Hentschel (1992). Alternatively, Black (1976) and Christie (1982) justify the negative correlation between returns and innovations to the variance by the “leverage” effect. Bekaert and Wu (2000) conclude that the feedback effect dominates the leverage effect.

²⁰See also Campbell and Hentschel (1992) for an examination of the risk-return tradeoff with asymmetric variance effects.

where $\mathbf{1}_{t-d}^+$ denotes the indicator function for $\{r_{t-d} \geq 0\}$, $\mathbf{1}_{t-d}^-$ denotes the indicator function for $\{r_{t-d} < 0\}$, and ϕ is in the interval $(0, 2)$. This formulation allows for a differential impact of positive and negative shocks on the conditional variance. The coefficient ϕ controls the total weight of negative shocks on the conditional variance. A coefficient ϕ between zero and two ensures that the total weights sum up to one since the indicator functions are mutually exclusive and each of the positive and negative weight functions add up to one. A value of ϕ equal to one places equal weight on positive and negative shocks. The two sets of parameters $\{\kappa_1^-, \kappa_2^-\}$ and $\{\kappa_1^+, \kappa_2^+\}$ characterize the time profile of the weights from negative and positive shocks, respectively.

Table 8 reports the estimates of the risk-return tradeoff (1) with the conditional variance estimator in equation (11). The estimated coefficient γ is 2.482 and highly significant in the entire sample. In contrast to the findings of Glosten, Jagannathan, and Runkle (1993) with asymmetric GARCH models, in the MIDAS framework, allowing the conditional variance to respond asymmetrically to positive and negative shocks does *not* change the sign of the risk-return tradeoff. Hence, asymmetries in the conditional variance are consistent with a positive coefficient γ in the ICAPM relation.

In agreement with previous studies, we find that asymmetries play an important role in driving the conditional variance. The statistical significance of the asymmetries can easily be tested using a likelihood ratio test. The restricted likelihood function under the null hypothesis of no asymmetries is presented in Table 2, whereas the unrestricted likelihood with asymmetries appears in Table 8. The null of no asymmetries, which is a joint test of $\kappa_1^+ = \kappa_1^-$, $\kappa_2^+ = \kappa_2^-$, and $\phi = 1$, is easily rejected with a p -value of less than 0.001.

The κ coefficients are of interest because they parameterize the weight functions $w_d(\kappa_1^-, \kappa_2^-)$ and $w_d(\kappa_1^+, \kappa_2^+)$. We plot these weight functions in Figure 5. Interestingly, the weight profiles of negative and positive shocks are markedly different. All the weight of negative shocks (dash-dot line) on the conditional variance is concentrated in the first 30 daily lags. In other words, negative shocks have a strong impact on the conditional variance, but that impact is transitory. It disappears after only one month. In contrast, positive returns (dash-dash line) have a much smaller immediate impact, but their effect persists up to a year after the shock. Their decay is much slower than the usual exponential rate of decay obtained in the case of GARCH models.

We find that the estimated value of ϕ is less than one. Since ϕ measures the *total* impact

of negative shocks on the conditional variance, our finding implies that *positive* shocks have overall a greater weight on the conditional variance than do *negative* shocks. This asymmetry is statistically significant. A t -test of the null hypothesis of $\phi = 1$ is rejected with a p -value of 0.009. The combined effect of positive and negative shocks, weighted by ϕ , is plotted as a thick solid line in Figure 5 (the symmetric weight is also plotted for reference as a thin solid line). In the short run, negative returns actually have a higher impact on the conditional variance since their estimated weight in the first month is so much larger than the weight on positive shocks in the same period. For longer lag lengths, the coefficient ϕ determines that positive shocks actually become more important.

We thus find that the asymmetry in the response of the conditional variance to positive and negative returns is more complex than previously documented. Negative shocks have a higher immediate impact but are ultimately dominated by positive shocks. Also, there is a clear asymmetry in the persistence of positive and negative shocks, with positive shocks being responsible for the persistence of the conditional variance process beyond one month.

Our results are consistent with the recent literature on multi-factor variance models (Alizadeh, Brandt, and Diebold (2002), Chacko and Viceira (2003), Chernov, Gallant, Ghysels, and Tauchen (2002), and Engle and Lee (1999), among others) which finds reliable support for the existence of two factors driving the conditional variance. The first factor is found to have high persistence and low volatility, whereas the second factor is transitory and highly volatile. The evidence from estimating jump-diffusions with stochastic volatility points in a similar direction. For example, Chernov, Gallant, Ghysels, and Tauchen (2002) show that the diffusive component is highly persistent and has low variance, whereas the jump component is by definition not persistent and is highly variable.

Using the asymmetric MIDAS specification, we are able to identify the first factor with lagged positive returns and the second factor with lagged negative returns.²¹ Indeed, if we decompose the conditional variance estimated with equation (11) into its two components, $\phi \sum_{d=0}^{\infty} w_d(\kappa_1^-, \kappa_2^-) \mathbf{1}_{t-d}^- r_{t-d}^2$ and $(2 - \phi) \sum_{d=0}^{\infty} w_d(\kappa_1^+, \kappa_2^+) \mathbf{1}_{t-d}^+ r_{t-d}^2$, we verify that their time series properties match the results in the literature on two-factor models of variance. More precisely, the positive shock component is very persistent, with an AR(1) coefficient of 0.989, whereas the negative shock component is temporary, with an AR(1) coefficient of only 0.107. Also, the standard deviation of the negative component is twice the standard deviation of the positive component. These findings are robust in the subsamples.

²¹Engle and Lee (1999) have a similar finding using a two-component asymmetric GARCH model.

4.2 Asymmetric GARCH Tests

For comparison with the asymmetric MIDAS results, we estimate three different asymmetric GARCH-in-mean models: an asymmetric GARCH (ASYGARCH), an exponential GARCH (EGARCH), and a quadratic GARCH (QGARCH). The ASYGARCH and EGARCH formulations are widely used to model asymmetries in the conditional variance and have been used in the risk-return tradeoff literature by Glosten, Jagannathan, and Runkle (1993). The QGARCH model was introduced by Engle (1990) and is used in the risk-return tradeoff context by Campbell and Hentschel (1992). We also estimate a more general GARCH-in-mean class of models, proposed by Hentschel (1995), that nests not only the previous three GARCH specifications, but also the simple GARCH and the ABS-GARCH from the previous section, and several other GARCH models. Following Hentschel (1995), a general class of GARCH models can be written as

$$\frac{V_t^\lambda - 1}{\lambda} = \omega + \alpha V_{t-1}^\lambda (|u_t + b| + c(u_t + b))^\nu + \beta \frac{V_{t-1}^\lambda - 1}{\lambda} \quad (12)$$

where u_t is the residual normalized to have a mean of zero and unit variance. This Box and Cox (1964) transformation of the conditional variance is useful because it nests all the previously discussed models. The simple GARCH model obtains when $\lambda = 1$, $\nu = 2$, and $b = c = 0$ and the ABS-GARCH obtains when $\lambda = 1/2$, $\nu = 1$, and $b = c = 0$.

The asymmetric GARCH models are nested when we allow the parameters b or c to be different from zero. The ASYGARCH model corresponds to the restrictions $\lambda = 1$, $\nu = 2$, and $b = 0$, with the value of c unrestricted. The coefficient c captures the asymmetric reaction of the conditional variance to positive and negative returns. A negative c indicates that negative returns have a stronger impact on the conditional variance. When $c = 0$, the ASYGARCH model reduces to simple GARCH. The EGARCH model obtains when $\lambda \rightarrow 0$, $\nu = 1$, $b = 0$, and c is left unrestricted, because $\lim_{\lambda \rightarrow 0} \frac{V^\lambda - 1}{\lambda} = \ln V$. This model is similar in spirit to ASYGARCH, but imposes an exponential form on the dynamics of the conditional variance as a more convenient way of ensuring positiveness. Again, when c is negative, the variance reacts more to negative return shocks. The QGARCH model corresponds to the restrictions $\lambda = 1$, $\nu = 2$, and $c = 0$, with b left unrestricted.²² When b is negative, the variance reacts more to negative returns and for $b = 0$, the QGARCH model collapses into

²²The formulation of Campbell and Hentschel (1992) has a negative sign in front of the b term. We write the QGARCH model differently to maintain the interpretation of a negative b corresponding to a higher impact of negative shocks on the conditional variance.

the simple GARCH specification. For more details on these models, see Hentschel (1995).

In Table 9, we first estimate (12) by imposing the coefficient restrictions of ASYGARCH, EGARCH, and QGARCH in order to facilitate comparison of the results with the previous literature. We also estimate the unrestricted version of (12) to show that none of the results are driven by the restrictions. The estimated coefficients of the restricted and unrestricted asymmetric GARCH models are shown in Table 9. We confirm the finding in Glosten, Jagannathan, and Runkle (1993) that asymmetries in the ASYGARCH and EGARCH produce a *negative*, albeit statistically insignificant, estimate of the risk-return tradeoff parameter γ . Our estimates of the model are similar to theirs. The QGARCH model also produces a negative and statistically insignificant estimate of γ , which is comparable (although lower in absolute terms) to the negative and statistically insignificant estimates obtained in Campbell and Hentschel (1992).²³ In all three restricted models, the estimates of b or c are negative and statistically different from zero, indicating that the asymmetries are important and that, in asymmetric GARCH models, negative shocks tend to have a higher impact on the conditional variance than positive shocks. The same observations hold true for the unrestricted GARCH model, where the estimate of γ is slightly lower in absolute value, but still negative and insignificant. Our results are in general agreement with Hentschel (1995), who uses daily data and a slightly shorter time period. Finally, comparing the $R_{\sigma^2}^2$ from Tables 5 and 9, we notice that the asymmetric GARCH models produce forecasts of the realized variance that are better than those from the symmetric GARCH models.

The persistence of the conditional variance in the above asymmetric GARCH models is driven by the β parameter. It is important to note that the asymmetric GARCH specifications do not allow for differences in the persistence of positive and negative shocks. In other words, positive and negative shocks decay at the same rate, determined by β . Furthermore, the estimated conditional variance in such asymmetric GARCH processes loads heavily on negative shocks, which we know from the MIDAS results (Figure 5) have a strong immediate impact on volatility. However, we have also seen that the impact of negative shocks on variance is transitory. Hence, it is not surprising that the estimates of the persistence parameter β in the asymmetric GARCH models shown in Table 9 (similar to Glosten, Jagannathan, and Runkle (1993)) are much lower than in the symmetric GARCH models.²⁴ This implicit restriction leads Glosten, Jagannathan, and Runkle to conclude that

²³In addition to this result, Campbell and Hentschel (1992) estimate the risk-return tradeoff imposing a constraint from a dividend-discount model. In that case, they estimate a positive and significant γ .

²⁴This constraint can be relaxed in the GARCH framework. Using a two-component GARCH model,

“the conditional volatility of the monthly excess return is not highly persistent.” In contrast, the asymmetric MIDAS model allows the persistence of positive and negative shocks to be different, resulting in overall higher persistence of the variance process.

To demonstrate the implications of the asymmetric GARCH restriction on the persistence of positive and negative shocks, we compute the AR(1) coefficient of the filtered variance processes. The AR(1) coefficients of the ASYGARCH, EGARCH, QGARCH, and generalized asymmetric GARCH conditional variance processes are only 0.457, 0.414, 0.284, and 0.409, respectively.²⁵ These coefficients are surprisingly low given what we know about the persistence of conditional variance (Officer (1973) and Schwert (1989)). The constraint that asymmetric GARCH models place, that positive and negative shocks be equally persistent, thus imposes a heavy toll on the overall persistence of the forecasted variance process. In contrast, the AR(1) coefficient of the symmetric GARCH and the symmetric MIDAS estimators (reported in Table 6) are 0.970 and 0.872, respectively. It is worth noting that the lack of persistence is not due to the asymmetry in the variance process as specification (12) allows for a very flexible form of asymmetries. Contrary to the asymmetric GARCH models, the AR(1) coefficient of the asymmetric MIDAS estimate is still high at 0.844, showing that the conditional variance process can have both asymmetries and high persistence.

It is thus not surprising that asymmetric GARCH models are incapable of explaining expected returns in the ICAPM relation.²⁶ This explains the puzzling findings of Glosten, Jagannathan, and Runkle (1993) that the risk-return tradeoff turns negative when we take into account asymmetries in the conditional variance. Their results are not driven by asymmetries. Instead, they depend on the lack of persistence in the conditional variance induced by the restriction in the asymmetric GARCH processes. To adequately capture the dynamics of variance, we need both asymmetry in the reaction to negative and positive shocks and a different degree of persistence of those shocks. When we model the conditional variance with the asymmetric MIDAS specification, the ICAPM continues to hold.

Engle and Lee (1999) show that only the persistent component of variance has explanatory power for stock market returns. Also, Hentschel (1995) finds higher estimates of β using daily data.

²⁵In the subsamples, we have observed AR(1) coefficients close to zero or even negative.

²⁶Indeed, Poterba and Summers (1986) show that persistence in the variance process is crucial for it to have any economically meaningful impact on stock prices.

5 The Risk-Return Tradeoff with Additional Predictive Variables

In this section, we extend the ICAPM relation between risk and return to include other predictive variables. Specifically, we modify the ICAPM equation (1) as:

$$E_t[R_{t+1}] = \mu + \gamma \text{Var}_t[R_{t+1}] + \theta^\top Z_t \quad (13)$$

where Z_t is a vector of variables known to predict the return on the market and θ is a conforming vector of coefficients. The variables in Z_t are known at the beginning of the return period.

Campbell (1991), Campbell and Shiller (1988), Chen, Roll, and Ross (1986), Fama (1990), Fama and French (1988, 1989), Ferson and Harvey (1991), and Keim and Stambaugh (1986), among many others, find evidence that the stock market can be predicted by variables related to the business cycle. At the same time, Schwert (1989, 1990b) shows that the variance of the market is highly counter-cyclical. Therefore, our findings about the risk-return tradeoff could simply be due to the market variance proxying for business cycle fluctuations. To test this “proxy” hypothesis, we examine the relation between the expected return on the stock market and the conditional variance using macro variables as controls for business cycle fluctuations.

Alternatively, specification (13) can be understood as a version of the ICAPM with additional state variables. When the investment opportunity set changes through time, Merton shows that:

$$E_t[R_{t+1}] = \mu + \gamma \text{Var}_t[R_{t+1}] + \pi^\top \text{Cov}_t[R_{t+1}, S_{t+1}], \quad (14)$$

where the term $\text{Cov}_t[R_{t+1}, S_{t+1}]$ denotes a vector of covariances of the market return with innovations to the state variables, S , conditional on information known at date t . If the relevant information to compute these conditional covariances consists of the variables in the vector Z_t , we can interpret the term $\theta^\top Z_t$ in (13) as an estimate of the conditional covariance term, $\pi^\top \text{Cov}_t[R_{t+1}, S_{t+1}]$ in (14). Campbell (1987) and Scruggs (1998) emphasize this version of the ICAPM, which predicts only a *partial* relation between the conditional

mean and the conditional variance after controlling for the other covariance terms.²⁷

The predictive variables that we study are the dividend-price ratio, the relative Treasury bill rate, the default spread, and the lagged monthly return (all available at monthly frequency). These variables have been widely used in the predictability literature (Campbell and Shiller (1988), Campbell (1991), Fama and French (1989), Torous, Valkanov, and Yan (2003) and, for a good review, Campbell, Lo, and MacKinlay (1997)). The dividend-price ratio is calculated as the difference between the log of the last twelve month dividends and the log of the current level of the CRSP value-weighted index. The three-month Treasury bill rate is obtained from Ibbotson Associates. The relative Treasury bill stochastically detrends the raw series by taking the difference between the interest rate and its twelve-month moving average. The default spread is calculated as the difference between the yield on BAA- and AAA-rated corporate bonds, obtained from the FRED database. We standardize the control variables (subtracting the mean and dividing by the standard deviation) to ensure comparability of the μ coefficients in equations (1) and (13).

There is an additional reason to include the lagged squared return as a control variable.²⁸ Note that the MIDAS estimator uses lagged squared returns as a measure of conditional variance. This is not strictly speaking a measure of variance but rather a measure of the second (uncentered) moment of returns. In particular, it includes the squared conditional mean of returns. Omitting serial correlation from the return model and including the mean return in the variance filter may induce a spurious relation between conditional mean and conditional variance. To illustrate this point, consider the lagged monthly squared return as a simple estimator of variance. Assume further that returns follow an AR(1) process:

$$\begin{aligned} R_{t+1} &= \phi R_t + \epsilon_{t+1} \\ \sigma_t^2 &= R_t^2. \end{aligned} \tag{15}$$

In this system, the autocorrelation of returns and the inclusion of the mean in the variance

²⁷Scruggs uses the covariance between stock market returns and returns on long bonds as a control and finds a significantly positive risk-return tradeoff.

²⁸We thank the referee for this insight and the following example.

filter imply that:

$$\begin{aligned}\text{Cov}(R_{t+1}, \sigma_t^2) &= \phi \text{Cov}(R_t, R_t^2) \\ &= \phi (ER_t^3 + ER_t^2 ER_t).\end{aligned}\tag{16}$$

Hence, there is a mechanical correlation between returns and conditional variance unless returns are not autocorrelated ($\phi = 0$), or returns have zero skewness and zero mean, or there is some fortuitous cancelation between skewness and mean. Adding lagged returns as a control variable in the risk-return relation addresses this problem.

Once the effect of the control variables in the conditional expected return is removed, γ captures the magnitude of the risk-return tradeoff, while the MIDAS weight coefficients still determine the lag structure of conditional variance. Table 10 presents the results from estimating equation (13) with both the simple MIDAS weights (3) (in Panel A) and the asymmetric MIDAS weights (11) (in Panel B). The results strongly suggest that neither business cycle fluctuations nor serial correlation in returns account for our findings. Indeed, the coefficients of the risk-return relation with controls are remarkably similar to those estimated without controls (shown in Tables 2 and 8). The estimates of μ and γ are almost identical in the two tables across all sample periods. This indicates that the explanatory power of the forecasted variance for returns is largely orthogonal to the additional macro variables. Although lagged market returns are significant in the first subsample, in which returns exhibit stronger serial correlation, as we noted in Table 1, controlling for their effect has no significant effect on the estimates of γ . Moreover, the estimates of κ_1 , and κ_2 are also very similar to the estimates without controls, implying that the weights placed on past squared returns are not changed.

The macro variables and lagged market returns enter significantly in the ICAPM conditional mean either in the sample or in the subsamples. A likelihood ratio test of their joint significance in the entire sample has a p -value of less than 0.001. The coefficient of determination of the regression of realized returns on the conditional variance *and* the control variables, R_R^2 , is 2.8 percent in the full sample. This is significantly higher than the corresponding coefficient without the control variables, which is only 1.9 percent. The adjusted $R_{\sigma^2}^2$ is unchanged by the inclusion of the predetermined monthly variables.

We conclude that the risk-return tradeoff is largely unaffected by including extra predictive variables in the ICAPM equation and the forecasting power of the conditional

variance is not merely proxying for the business cycle. Also, the estimated positive risk-return tradeoff is unlikely to be due to serial correlation in the conditional mean of returns.

6 Conclusion

This paper takes a new look at Merton's ICAPM, focusing on the tradeoff between conditional variance and conditional mean of the stock market return. In support of the ICAPM, we find a *positive* and *significant* relation between risk and return. This relation is robust in subsamples, does not change when the conditional variance is allowed to react asymmetrically to positive and negative returns, and is not affected by the inclusion of other predictive variables.

Our results are more conclusive than those from previous studies due to the added power obtained from the new MIDAS estimator of conditional variance. This estimator is a weighted average of past daily squared returns and the weights are parameterized with a flexible functional form. We find that the MIDAS estimator is a better forecaster of the stock market variance than rolling window or GARCH estimators, which is the reason why our tests can robustly find the ICAPM's risk-return tradeoff.

We obtain new results about the asymmetric reaction of volatility to positive and negative return shocks. We find that, compared to negative shocks, positive shocks: have a bigger impact overall on the conditional mean of returns; are slower to be incorporated into the conditional variance; and are much more persistent and indeed account for the persistent nature of the conditional variance process. Surprisingly, negative shocks have a large initial, but very temporary effect on the variance of returns.

The MIDAS estimator offers a powerful and flexible way of estimating economic models by taking advantage of data sampled at various frequencies. While the advantages of the MIDAS approach have been demonstrated in the estimation of the ICAPM and conditional volatility, the method itself is quite general in nature and can be used to tackle several other important questions.

References

- Abel, Andrew B., 1988, Stock prices under time-varying dividend risk: An exact solution in an infinite-horizon general equilibrium model, *Journal of Monetary Economics* 22, 375–393.
- Alizadeh, Sassan, Michael W. Brandt, and Francis X. Diebold, 2002, Range-based estimation of stochastic volatility models, *Journal of Finance* 57, 1047–1091.
- Andersen, Torben, and Tim Bollerslev, 1998, Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* 39, 885–905.
- , and Nour Meddahi, 2002, Analytic evaluation of volatility forecasts, Working paper, CIRANO.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens, 2001, The distribution of realized stock return volatility, *Journal of Financial Economics* 61, 43–76.
- Andreou, Elena, and Eric Ghysels, 2002, Time variations and covariations in the expectation and volatility of stock market returns, *Journal of Business and Economic Statistics* 20, 363–376.
- Backus, David K., and Allan W. Gregory, 1993, Theoretical relations between risk premiums and conditional variance, *Journal of Business and Economic Statistics* 11, 177–185.
- Baillie, Richard T., and Ramon P. DeGennaro, 1990, Stock returns and volatility, *Journal of Financial and Quantitative Analysis* 25, 203–214.
- Barndorff-Nielsen, Ole E., and Neil Shephard, 2002, Econometric analysis of realised volatility and its use in estimating stochastic volatility models, *Journal of the Royal Statistical Society, Series B* 64, 253–280.
- Bekaert, Geert, and Guojun Wu, 2000, Asymmetric volatility and risk in equity markets, *Review of Financial Studies* 13, 1–42.
- Black, Fischer, 1976, Studies in stock price volatility changes, in *Proceedings of American Statistical Association, Business and Economic Statistics Section* pp. 177–181.

- Bollerslev, Tim, and Jeffrey M. Wooldridge, 1992, Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances, *Econometric Reviews* 11, 143–172.
- Box, G. E., and D. R. Cox, 1964, An analysis of transformations, *Journal of the Royal Statistical Society B* 26, 211–243.
- Brandt, Michael, and Qiang Kang, 2003, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Economics* forthcoming.
- Campbell, John Y., 1987, Stock returns and the term structure, *Journal of Financial Economics* 18, 373–399.
- , 1991, A variance decomposition for stock returns, *Economic Journal* 101, 157–179.
- , 1993, Intertemporal asset pricing without consumption data, *American Economic Review* 83, 487–512.
- , and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281–318.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay, 1997, *The Econometrics of Financial Markets* (Princeton University Press: Princeton).
- Campbell, John Y., and Robert J. Shiller, 1988, Stock prices, earnings, and expected dividends, *Journal of Finance* 43, 661–676.
- Chacko, George, and Luis Viceira, 2003, Spectral GMM estimation of continuous-time processes, *Journal of Econometrics* 116, 259–292.
- Chan, K. C., G. Andrew Karolyi, and René M. Stulz, 1992, Global financial markets and the risk premium on U.S. equity, *Journal of Financial Economics* 32, 137–167.
- Chen, Nai-Fu, Richard Roll, and Stephen A. Ross, 1986, Economic forces and the stock market, *Journal of Business* 59, 383–403.
- Chernov, Mikhael, A. Ronald Gallant, Eric Ghysels, and George Tauchen, 2002, Alternative Models for Stock Price Dynamics, *Journal of Econometrics* 116, 225–257.

- Christie, Andrew A., 1982, The stochastic behavior of common stock variances: Value, leverage and interest rate effects, *Journal of Financial Economics* 10, 407–432.
- Engle, Robert, David Lilien, and Russell Robins, 1987, Estimation of time varying risk premia in the term structure: the arch-m model, *Econometrica* 55, 391–407.
- Engle, Robert F., 1990, Discussion: Stock market volatility and the crash of '87, *Review of Financial Studies* 3, 103–106.
- , and G. Lee, 1999, A long-run and short-run component model of stock return volatility, in R. F. Engle, and H. White, ed.: *Cointegration, Causality and Forecasting - A Festschrift in Honour of Clive W. J. Granger*. Oxford University Press.
- Engle, Robert F., and Victor Ng, 1993, Measuring and testing the impact of news on volatility, *Journal of Finance* 48, 1749–1778.
- Fama, Eugene F., 1990, Stock returns, expected returns, and real activity, *Journal of Finance* 45, 1089–1108.
- , and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.
- , 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- Ferson, Wayne E., and Campbell R. Harvey, 1991, The variation of economic risk premiums, *Journal of Political Economy* 99, 385–415.
- Foster, Dean P., and Daniel B. Nelson, 1996, Continuous record asymptotics for rolling sample variance estimators, *Econometrica* 64, 139–174.
- French, Kenneth R., William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.
- Genotte, Gerard, and Terry A. Marsh, 1993, Variations in economic uncertainty and risk premiums on capital assets, *European Economic Review* 37, 1021–1041.
- Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2003, Predicting volatility: Getting the most out of return data sampled at different frequencies, Working paper, UCLA and UNC.

- 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.
- Goyal, Amit, and Pedro Santa-Clara, 2002, Idiosyncratic risk matters!, *Journal of Finance* forthcoming.
- Hall, Robert E., 1988, Intertemporal substitution in consumption, *Journal of Political Economy* 96, 339–357.
- Harvey, Campbell R., 2001, The specification of conditional expectations, *Journal of Empirical Finance* 8, 573–638.
- Hentschel, Ludger, 1995, All in the family: Nesting symmetric and asymmetric GARCH models, *Journal of Financial Economics* 39, 71–104.
- Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357–390.
- Lettau, Martin, and Sydney C. Ludvigson, 2002, Measuring and modeling variation in the risk-return tradeoff, in Yacine Ait-Sahalia, and Lars P. Hansen, ed.: *Handbook of Financial Econometrics*.
- Mayfield, E. Scott, 2003, Estimating the market risk premium, *Journal of Financial Economics* forthcoming.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.
- , 1980, On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323–361.
- Nelson, Daniel B., 1991, Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* 59, 347–370.
- Officer, Robert R., 1973, The variability of the market factor of the New York Stock Exchange, *Journal of Business* 46, 434–53.
- Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

- Pindyck, Robert S., 1984, Risk, inflation, and the stock market, *American Economic Review* 74, 335–351.
- Poterba, James M., and Lawrence H. Summers, 1986, The persistence of volatility and stock market fluctuations, *American Economic Review* 76, 1142–1151.
- Santa-Clara, Pedro, and Shu Yan, 2001, Risk and return in the stock market: Lessons from option data, Working paper, University of Arizona.
- Schwert, G. William, 1990a, Indexes of united states stock prices from 1802 to 1987, *Journal of Business* 63, 399–426.
- Schwert, William G., 1989, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115–1153.
- , 1990b, Stock volatility and the crash of 87, *Review of Financial Studies* 3, 77–102.
- Scruggs, John T., 1998, Resolving the puzzling intertemporal relation between the market risk premium and conditional market variance: A two-factor approach, *Journal of Finance* 52, 575–603.
- Taylor, Stephen J., and Xuxinzhong Xu, 1997, The incremental volatility information in one million foreign exchange quotations, *Journal of Empirical Finance* 4, 317–340.
- Torous, Walter, Rossen Valkanov, and Shu Yan, 2003, On Predicting Stock Returns with Nearly Integrated Explanatory Variables, *Journal of Business* forthcoming.
- Turner, Christopher M., Richard Startz, and Charles R. Nelson, 1989, A markov model of heteroskedasticity, risk, and learning in the stock market, *Journal of Financial Economics* 25, 3–22.
- Whitelaw, Robert F., 1994, Time variations and covariations in the expectation and volatility of stock market returns, *Journal of Finance* 49, 515–541.

Appendix

To better understand the MIDAS estimator, consider a continuous-time model of the instantaneous return dp_t (where p_t is the log price) with stochastic volatility:

$$\begin{aligned} dp_t &= \mu(\sigma_t)dt + \sigma_t dW_{1t} \\ d\sigma_t^2 &= \zeta(\sigma_t)dt + \delta(\sigma_t)dW_{2t} \end{aligned} \tag{17}$$

where W_1 and W_2 are standard Brownian motions (possibly correlated) and the functions $\mu(\cdot)$, $\zeta(\cdot)$ are continuous and $\delta(\cdot)$ is strictly positive. Merton (1980) considered models where σ_t is constant over non-overlapping time intervals and $\mu(\sigma_t)$ is a linear function of variance. Appendix A of Merton's paper shows that sampling data at very high frequency yields arbitrarily accurate volatility estimates. This insight prompted Merton to consider estimating volatility with equally weighted block-sampled data, which is a simple rolling window estimator. This approach has been used extensively by Merton (1980), French, Schwert, and Stambaugh (1987), and Schwert (1989), who typically used a month's worth of equally-weighted daily data in the rolling window estimator.

Foster and Nelson (1996) extended this line of work to processes with stochastic volatility, i.e., where the diffusion governing volatility dynamics in (17) is genuinely taken into account. Foster and Nelson use continuous-record asymptotic theory (which assumes that a fixed span of data is sampled at ever finer intervals) and propose volatility estimators based on sampling returns at a frequency $1/m$ that can generically be written as:

$$\sum_{\tau} \omega_{t-\tau} (r_{t-\tau}^{(m)})^2 \tag{18}$$

where $\omega_{t-\tau}$ is some weighting scheme and $r_t^{(m)}$ denotes returns sampled at frequency $1/m$.²⁹ Given the temporal dependence of volatility, one would expect that recent squared returns get more weight than distant ones. This intuition is indeed correct. Theorem 5 of Foster and Nelson (1996, p. 154) shows that the optimal weights for a class of stochastic volatility diffusions are of the form $\omega_{t-\tau} = \alpha \exp \alpha \tau$. Hence, the weights are exponentially declining

²⁹This estimator assumes that the drift over short intervals is negligible, which is justified by the analysis in Merton (1980).

at rate α .³⁰ Unfortunately, estimating α is rather involved. Indeed, to the best of our knowledge, apart from the small empirical application in Foster and Nelson (1996), there are no implementations of this estimator in the literature.

The MIDAS regression approach allows us to bypass the estimation of α in the Foster and Nelson (1996) optimal weighting scheme. Instead, the weighting function is chosen to maximize the likelihood of the data. The Foster and Nelson scheme is “optimal” in a minimum MSE sense, yet this optimality is only established for a restricted class of diffusions. In particular, optimal weighting schemes have not been explicitly derived for more general data generating processes such as diffusions involving asymmetric volatility. The MIDAS approach relies on a different optimality principle, namely that of maximum likelihood. It is not directly comparable with the optimality criterion of Foster and Nelson, but has the advantage of being easy to implement and widely applicable.

³⁰Foster and Nelson show that $\alpha = \sqrt{\Lambda/\theta}$, where, intuitively, θ is closely related to the local martingale component of the Doob-Meyer decomposition associated with (17), and Λ is the variance of the conditional variance process (in the example above that would be δ^2). Formal expressions for θ and Λ require definitions and concepts that are not of direct interest here. Details can be found in Foster and Nelson (1996, p 142-143).

Table 1: Summary Statistics of Returns and Realized Variance

The table shows summary statistics of monthly excess returns, R_t , of the stock market, and realized monthly variance computed from within-month daily data, σ_t^2 . The proxy for the stock market is the CRSP value-weighted portfolio and the risk-free rate is the yield on the three-month Treasury bill. The table shows the mean, variance, skewness, kurtosis, first-order serial correlation, and the sum of the first 12 auto-correlations, for each of the variables. The statistics are shown for the full sample and for two subsamples of approximately equal length.

Panel A: Monthly Excess Returns (R_t)							
Sample	Mean ($\times 10^2$)	Variance ($\times 10^2$)	Skewness	Kurtosis	AR(1)	AR(1-12)	T
1928:01-2000:12	0.649	0.321	-0.189	10.989	0.068	0.126	876
1928:01-1963:12	0.782	0.461	-0.095	10.105	0.077	0.199	432
1964:01-2000:12	0.499	0.198	-0.566	5.261	0.045	-0.031	444

Panel B: Monthly Realized Variance (σ_t^2)							
Sample	Mean ($\times 10^2$)	Variance ($\times 10^4$)	Skewness	Kurtosis	AR(1)	AR(1-12)	T
1928:01-2000:12	0.262	0.323	7.046	71.651	0.608	0.840	876
1928:01-1963:12	0.372	0.551	5.275	42.006	0.648	0.860	432
1964:01-2000:12	0.162	0.087	13.210	226.977	0.265	0.482	444

Table 2: MIDAS Tests of the Risk-Return Tradeoff

The table shows estimates of the risk-return tradeoff (1) with the MIDAS estimator of conditional variance in equations (2) and (3). Daily returns are used in the construction of the conditional variance estimator. Monthly returns are used in the estimation of the risk-return tradeoff parameter γ . The coefficients and corresponding t -statistics (in brackets) are shown for the full sample and the two subsamples. The t -statistics are computed using Bollerslev-Wooldridge standard errors. R_R^2 and $R_{\sigma^2}^2$ quantify the explanatory power of the MIDAS variance estimator in predictive regressions for realized returns and variances, respectively. LLF is the log-likelihood value.

Sample	μ ($\times 10^3$)	γ	κ_1 ($\times 10^3$)	κ_2 ($\times 10^5$)	R_R^2	$R_{\sigma^2}^2$	LLF
1928:01-2000:12	6.430 [11.709]	2.606 [6.710]	-5.141 [-4.528]	-10.580 [-5.241]	0.019	0.407	1421.989
1928:01-1963:12	11.676 [5.887]	1.547 [3.382]	-0.909 [-3.770]	-10.807 [-2.106]	0.011	0.444	681.237
1964:01-2000:12	3.793 [5.673]	3.748 [8.612]	-6.336 [-7.862]	-18.586 [-7.710]	0.050	0.082	807.193

Table 3: MIDAS Tests of the Risk-Return Tradeoff at Different Frequencies

The table shows estimates of the risk-return tradeoff (1) with the MIDAS estimator of conditional variance in equations (2) and (3) at different horizons. Daily returns are used in the construction of the conditional variance estimator. Daily, weekly, monthly, bi-monthly, and quarterly returns are used in the estimation of the risk-return tradeoff parameter γ . The coefficients and corresponding t -statistics (in brackets) are shown for the full sample and the two subsamples. The t -statistics are computed using Bollerslev-Wooldridge standard errors. R_R^2 and $R_{\sigma^2}^2$ quantify the explanatory power of the MIDAS variance estimator in predictive regressions for realized returns and variances, respectively. LLF is the log-likelihood value.

Horizon	μ ($\times 10^3$)	γ	R_R^2	$R_{\sigma^2}^2$	LLF
Sample: 1928:01-2000:12					
Daily	0.275 [13.422]	2.684 [1.154]	0.004	0.059	57098.422
Weekly	1.320 [13.156]	2.880 [3.127]	0.009	0.119	8441.573
Monthly	6.430 [11.709]	2.606 [6.710]	0.019	0.407	1421.989
Bi-Monthly	14.218 [12.007]	1.964 [4.158]	0.018	0.309	583.383
Quarterly	24.992 [12.029]	2.199 [4.544]	0.016	0.329	377.901
Sample: 1928:01-1963:12					
Daily	0.319 [2.762]	2.120 [1.167]	0.002	0.096	24627.123
Weekly	2.463 [3.441]	1.870 [0.054]	0.008	0.181	3529.511
Monthly	11.676 [5.887]	1.547 [3.382]	0.011	0.444	681.237
Bi-Monthly	23.547 [6.087]	1.627 [3.123]	0.012	0.468	305.823
Quarterly	36.741 [6.565]	1.682 [3.270]	0.010	0.421	217.287
Sample: 1964:01-2000:12					
Daily	0.214 [2.210]	3.377 [1.906]	0.004	0.043	31437.438
Weekly	0.846 [3.303]	3.804 [3.060]	0.009	0.080	4851.063
Monthly	3.793 [5.673]	3.748 [8.612]	0.050	0.082	807.193
Bi-Monthly	7.223 [6.917]	3.660 [6.246]	0.040	0.079	369.865
Quarterly	7.812 [8.395]	3.476 [5.028]	0.021	0.081	264.068

Table 4: Rolling Window Tests of the Risk-Return Tradeoff

The table shows estimates of the risk-return tradeoff (1) with the rolling window estimators of conditional variance (6). The coefficients and corresponding t -statistics (in brackets) are shown for the entire sample, 1928:01–2000:12. The t -statistics are computed using Bollerslev-Wooldridge standard errors. R_R^2 and $R_{\sigma^2}^2$ quantify the explanatory power of the MIDAS variance estimator in predictive regressions for realized returns and variances, respectively. LLF is the log-likelihood value.

Horizon (Months)	μ ($\times 10^3$)	γ	R_R^2	$R_{\sigma^2}^2$	LLF
1	9.523 [4.155]	0.546 [0.441]	0.007	0.370	1292.454
2	7.958 [4.071]	1.494 [1.532]	0.009	0.379	1325.528
3	6.823 [3.240]	2.171 [1.945]	0.011	0.382	1308.923
4	6.830 [3.344]	2.149 [2.212]	0.012	0.384	1346.685
5	7.972 [3.506]	1.458 [1.325]	0.011	0.383	1335.114
6	7.924 [3.409]	1.483 [1.316]	0.011	0.382	1334.067

Table 5: GARCH Tests of the Risk-Return Tradeoff

The table shows estimates of the risk-return tradeoff (1) with the GARCH estimators of conditional variance (7) and (8). The coefficients and corresponding t -statistics (in brackets) are shown for the entire sample, 1928:01–2000:12. The t -statistics are computed using Bollerslev-Wooldridge standard errors. R_R^2 and $R_{\sigma^2}^2$ quantify the explanatory power of the MIDAS variance estimator in predictive regressions for realized returns and variances, respectively. LLF is the log-likelihood value.

Model	μ ($\times 10^3$)	γ	ω ($\times 10^3$)	α	β	R_R^2	$R_{\sigma^2}^2$	LLF
GARCH-M	8.310 [3.899]	1.060 [1.292]	0.069 [0.675]	0.116 [7.550]	0.864 [31.840]	0.005	0.359	1400.086
ABSGARCH-M	7.439 [3.082]	1.480 [1.415]	0.174 [0.218]	0.091 [5.005]	0.900 [75.410]	0.004	0.332	1399.280

Table 6: Comparison of Conditional Variance Models

Panel A of the table displays means, variances, first-order serial correlations, and sums of the first 12 auto-correlations of the MIDAS, GARCH, and rolling window (RW) filtered conditional variances. Panel B shows the variance, skewness, and kurtosis of the standardized residuals, defined as the demeaned monthly returns divided by square root of the respective forecasted variance. The goodness-of-fit measure in the last column of Panel B is computed as one minus the sum of squared forecasting errors divided by the total sum of squared realized variance. The goodness-of-fit measures the forecasting power of each approach for the realized variance. Panel C displays the correlation matrix between the MIDAS, GARCH, rolling window, and realized conditional variances. The statistics are shown for the entire sample, 1928:01–2000:12.

Panel A: Summary Statistics				
Estimator	Mean ($\times 10^2$)	Variance ($\times 10^4$)	AR(1)	AR(1-12)
MIDAS	0.256	0.198	0.872	0.914
GARCH	0.325	0.187	0.970	0.964
RW	0.262	0.323	0.608	0.840

Panel B: Performance of Conditional Variance Models				
Estimator	Variance Std. Resids	Skewness Std. Resids	Kurtosis Std. Resids	Goodness of Fit
MIDAS	0.994	2.176	17.193	0.494
GARCH	0.992	8.562	91.101	0.440
RW	1.103	9.423	23.424	0.354

Panel C: Correlations				
	MIDAS	GARCH	RW	Realized
MIDAS	1.000	—	—	—
GARCH	0.802	1.000	—	—
RW	0.798	0.660	1.000	—
Realized	0.638	0.599	0.609	1.000

Table 7: Comparison of MIDAS and GARCH using Daily and Monthly Returns

The table shows estimates of the risk-return tradeoff (1). We use MIDAS and GARCH estimators of conditional variance with daily and monthly data in forecasting future variance. The daily MIDAS estimator is reproduced from Table (2) (entire sample). The monthly MIDAS estimator is defined in equations (10) and (3), where we use 12 lagged monthly returns instead of 252 lagged daily returns. The daily GARCH parameters are estimated with daily market returns and forecasts of monthly variances obtained by summing 22 daily variance forecasts. The monthly GARCH estimator is reproduced from Table (5). The coefficients and corresponding t -statistics (in brackets) are shown for the full sample in all specifications. The t -statistics are computed using Bollerslev-Wooldridge standard errors. R_R^2 and $R_{\sigma^2}^2$ quantify the explanatory power of the variance estimators in predictive regressions for realized returns and variances, respectively. Panel B displays the correlation matrix of the forecasted variance using the four models and the realized variance.

Panel A: Estimates and Model Fit													
Frequency of RHS Variable	MIDAS						GARCH-M						
	μ ($\times 10^3$)	γ	κ_1 ($\times 10^3$)	κ_2 ($\times 10^5$)	R_R^2	$R_{\sigma^2}^2$	μ ($\times 10^3$)	γ	ω ($\times 10^6$)	α	β	R_R^2	$R_{\sigma^2}^2$
Daily	6.430 [11.709]	2.606 [6.710]	-5.141 [-4.528]	-10.580 [-5.241]	0.019	0.407	9.038 [3.843]	0.431 [0.592]	1.395 [2.062]	0.032 [9.943]	0.957 [41.037]	0.003	0.291
Monthly	5.815 [8.213]	2.553 [2.668]	($\times 10^1$) -3.672 [-4.822]	($\times 10^2$) -2.821 [-3.084]	0.011	0.382	8.310 [3.899]	1.060 [1.292]	($\times 10^3$) 0.069 [0.675]	0.116 [7.550]	0.864 [31.840]	0.005	0.359

Panel B: Correlations					
	MIDAS Daily	MIDAS Monthly	GARCH-M Daily	GARCH-M Monthly	Realized
MIDAS Daily	1				
MIDAS Monthly	0.885	1			
GARCH-M Daily	0.561	0.557	1		
GARCH-M Monthly	0.802	0.752	0.516	1	
Realized Monthly	0.638	0.618	0.539	0.599	1

Table 8: Asymmetric MIDAS Tests of the Risk-Return Tradeoff

The table shows estimates of the risk-return tradeoff (1) with the Asymmetric MIDAS estimator of conditional variance (11). Daily returns are used in the construction of the conditional variance estimator. Monthly returns are used in the estimation of the risk-return tradeoff parameter γ . The coefficients and corresponding t -statistics (in brackets) are shown for the full sample and the two subsamples. The t -statistics are computed using Bollerslev-Wooldridge standard errors. R_R^2 and $R_{\sigma^2}^2$ quantify the explanatory power of the MIDAS variance estimator in predictive regressions for realized returns and variances, respectively. LLF is the log-likelihood value.

Sample	μ ($\times 10^3$)	γ	κ_1^- ($\times 10^2$)	κ_2^- ($\times 10^3$)	κ_1^+ ($\times 10^2$)	κ_2^+ ($\times 10^5$)	ϕ	R_R^2	$R_{\sigma^2}^2$	LLF
1928:01-2000:12	7.912 [12.133]	2.482 [3.449]	18.838 [3.655]	-12.694 [-4.407]	0.188 [7.352]	-2.230 [-1.252]	0.572 [7.817]	0.041	0.429	1482.667
1928:01-1963:12	10.114 [14.242]	2.168 [2.493]	13.866 [5.242]	-10.924 [-3.743]	0.176 [5.496]	-3.241 [-1.435]	0.537 [7.054]	0.023	0.461	698.835
1964:01-2000:12	5.521 [6.470]	2.603 [4.544]	27.616 [3.790]	-15.767 [-3.920]	-0.392 [-3.259]	-0.050 [-0.052]	0.697 [13.948]	0.092	0.088	819.237

Table 9: Asymmetric GARCH Tests of the Risk-Return Tradeoff

The table shows estimates of the risk-return tradeoff in equation (1) where the conditional variance follows Hentschel (1995)'s generalized GARCH model (12). The process (12) is estimated under the ASYGARCH restrictions ($\lambda = 1$, $\nu = 2$, and $b = 0$), EGARCH restrictions ($\lambda = 0$, $\nu = 1$, and $b = 0$), and QGARCH restrictions ($\lambda = 1$, $\nu = 2$, and $c = 0$). It is also estimated with no restrictions on the λ , ν , b , and c parameters. The coefficients and corresponding t -statistics (in brackets) are shown for the entire sample, 1928:01-2000:12. The t -statistics are computed using Bollerslev-Wooldridge standard errors. R_R^2 and $R_{\sigma^2}^2$ quantify the explanatory power of the filtered variance estimator in predictive regressions for realized returns and variances, respectively. LLF is the log-likelihood value.

Model	μ ($\times 10^3$)	γ	ω ($\times 10^2$)	α	β	λ	ν	b ($\times 10^2$)	c ($\times 10^2$)	R_R^2	$R_{\sigma^2}^2$	LLF
ASYGARCH(1,1)-M	3.419 [4.008]	-1.354 [-0.995]	0.045 [0.854]	0.043 [3.245]	0.623 [14.823]	1	2	0	-29.596 [-6.682]	0.007	0.369	1409.803
EGARCH(1,1)-M	11.645 [9.354]	-1.668 [-1.345]	-484.622 [-2.867]	-0.112 [-2.847]	0.593 [12.847]	0	1	0	-3.985 [-3.294]	0.007	0.371	1409.954
QGARCH(1,1)-M	16.936 [11.056]	-1.098 [-0.747]	0.051 [0.921]	0.094 [4.662]	0.327 [7.328]	1	2	-15.235 [-5.095]	0	0.007	0.372	1410.635
Generalized GARCH(1,1)-M	5.678 [7.112]	-0.713 [-0.496]	0.039 [0.746]	0.069 [3.905]	0.589 [11.654]	0.862 [8.934]	1.764 [11.056]	-11.947 [-4.132]	-9.967 [-3.881]	0.009	0.389	1417.966

Table 10: MIDAS Tests of the Risk-Return Tradeoff Controlling for Other Predictive Variables

The table shows estimates of the risk-return tradeoff in equation (13) with the MIDAS estimator of conditional variance (2) and other predictive variables: the default spread (θ_1), the stochastically detrended risk-free interest rate (θ_2), the market's dividend yield (θ_3), and lagged market return (θ_4). To facilitate comparison of the MIDAS coefficients with previous tables, the four control variables are normalized to have mean zero and unit variance. Panels A and B present the results without and with asymmetries, respectively. The coefficients and corresponding t -statistics (in brackets) are shown for the full sample and the two subsamples. The t -statistics are computed using the Bollerslev-Wooldridge standard errors. R_R^2 and $R_{\sigma^2}^2$ quantify the explanatory power of the MIDAS variance estimator in predictive regressions for realized returns and variances, respectively. LLF is the log-likelihood value.

Panel A: No Asymmetries

Sample	μ ($\times 10^3$)	γ	κ_1 ($\times 10^3$)	κ_2 ($\times 10^5$)	θ_1 ($\times 10^3$)	θ_2 ($\times 10^3$)	θ_3 ($\times 10^3$)	θ_4 ($\times 10^3$)	R_R^2	$R_{\sigma^2}^2$	LLF
1928:01-2000:12	8.557 [0.414]	2.473 [7.866]	-5.985 [-10.905]	-10.531 [-5.518]	6.494 [2.542]	-5.077 [-2.619]	5.349 [3.319]	9.733 [2.964]	0.028	0.406	1429.630
1928:01-1963:12	2.417 [0.035]	1.694 [3.157]	-0.767 [-3.299]	-7.436 [-2.417]	15.883 [2.998]	-10.410 [-4.469]	7.760 [2.397]	11.012 [3.852]	0.015	0.418	710.239
1964:01-2000:12	9.050 [0.882]	3.459 [5.014]	-6.144 [-6.899]	-8.904 [-5.501]	8.597 [2.124]	-3.050 [-2.954]	10.123 [3.642]	0.112 [0.208]	0.059	0.082	828.476

Table continued on next page ...

Panel B: With Asymmetries

Sample	μ ($\times 10^3$)	γ	κ_1^- ($\times 10^2$)	κ_2^- ($\times 10^3$)	κ_1^+ ($\times 10^2$)	κ_2^+ ($\times 10^5$)	ϕ	θ_1 ($\times 10^3$)	θ_2 ($\times 10^3$)	θ_3 ($\times 10^3$)	θ_4 ($\times 10^3$)	R_R^2	$R_{\sigma^2}^2$	LLF
1928:01-2000:12	11.512 [3.684]	2.511 [3.579]	15.941 [3.800]	-12.926 [-5.031]	0.170 [6.307]	-1.950 [-1.512]	0.559 [7.381]	7.181 [3.109]	-5.871 [-2.546]	6.021 [3.804]	9.843 [2.107]	0.047	0.436	1489.057
1928:01-1963:12	11.973 [5.106]	2.186 [3.014]	14.857 [5.746]	-11.036 [-3.882]	0.181 [5.107]	-2.935 [-1.907]	0.534 [6.869]	9.548 [2.453]	-11.455 [-5.673]	8.016 [2.400]	-0.291 [-0.238]	0.026	0.471	732.067
1964:01-2000:12	11.546 [7.405]	2.791 [4.289]	21.397 [5.392]	-16.001 [-4.207]	-0.377 [-3.450]	-0.081 [-0.336]	0.670 [9.304]	9.567 [8.862]	-3.499 [-0.601]	10.991 [11.519]	0.123 [0.229]	0.101	0.096	829.927

Table continued from previous page.

Figure 1: MIDAS Weights

The figure plots the weights that the MIDAS estimator (2) and (3) places on lagged daily squared returns. The weights are calculated by substituting the estimated values of κ_1 , and κ_2 into the weight function (3). The estimates of κ_1 , and κ_2 are shown in Table 2. The figure displays the weights for the entire sample and for the two subsamples.




Figure 2: Filtered Conditional Variances and Realized Variance

The figure plots the filtered MIDAS, GARCH, and rolling window conditional variances, plotted in thick solid lines, and compares them with the realized variance (5), which is displayed in thin dotted line. The parameter values used to compute the filtered MIDAS, GARCH, and rolling window variances are in Tables 2, 5, and 4, respectively. For clarity of presentation, the conditional variances have been truncated to 0.04.

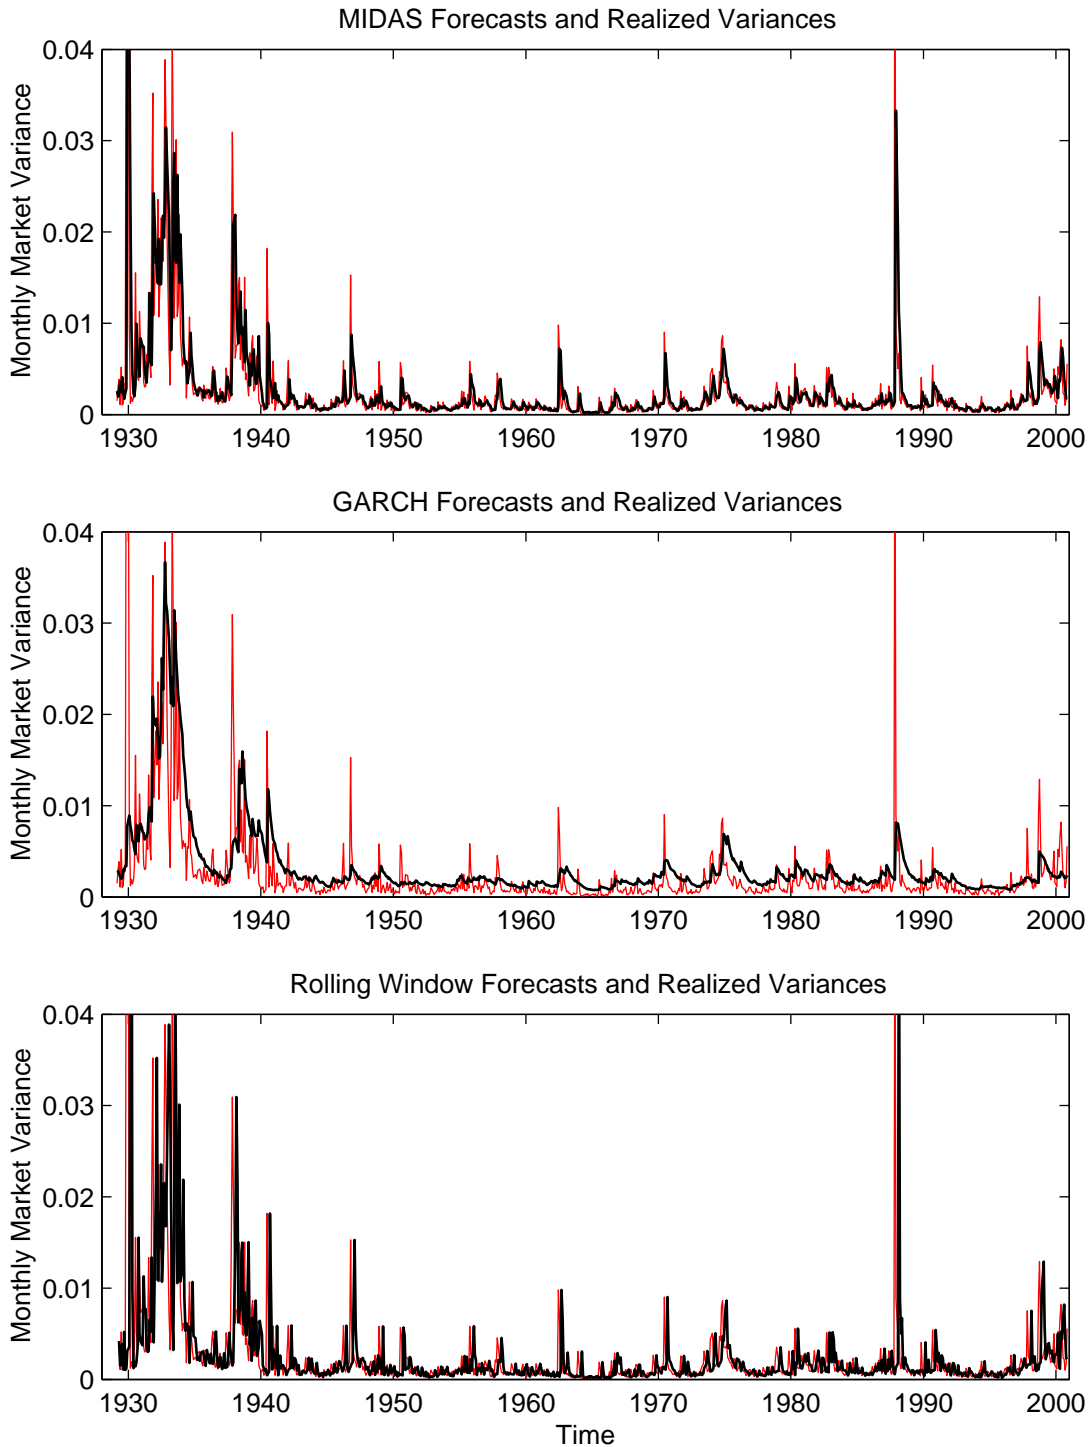


Figure 3: Scatterplot of Forecasted Variances versus Realized Variance

The figure displays scatterplots of the realized variance against the conditional MIDAS, GARCH, and rolling window conditional variances for each month in the 1928:01-2000:12 sample. The plots are shown in a log-log scale to facilitate comparison of the series during periods of low and extremely high volatility. A 45 degree solid line offers a reference point (indicating perfect fit) between realized variance and the filtered series.




Figure 4: MIDAS and GARCH Weights using Daily and Monthly Data

The first panel plots the weights that the daily and monthly MIDAS estimators place on lagged squared returns. The weights are calculated by substituting the estimated values of κ_1 and κ_2 from daily and monthly MIDAS into the weight function (3). The estimates of κ_1 and κ_2 are shown in Table 7. The second figure displays the weights that the daily and monthly GARCH estimators place on lagged squared returns. The weights are calculated by substituting the estimated values of α and β from daily and monthly GARCH into the weight function (9). The estimates of α and β are shown in Table 7. The figure displays the weights estimated from the entire sample.

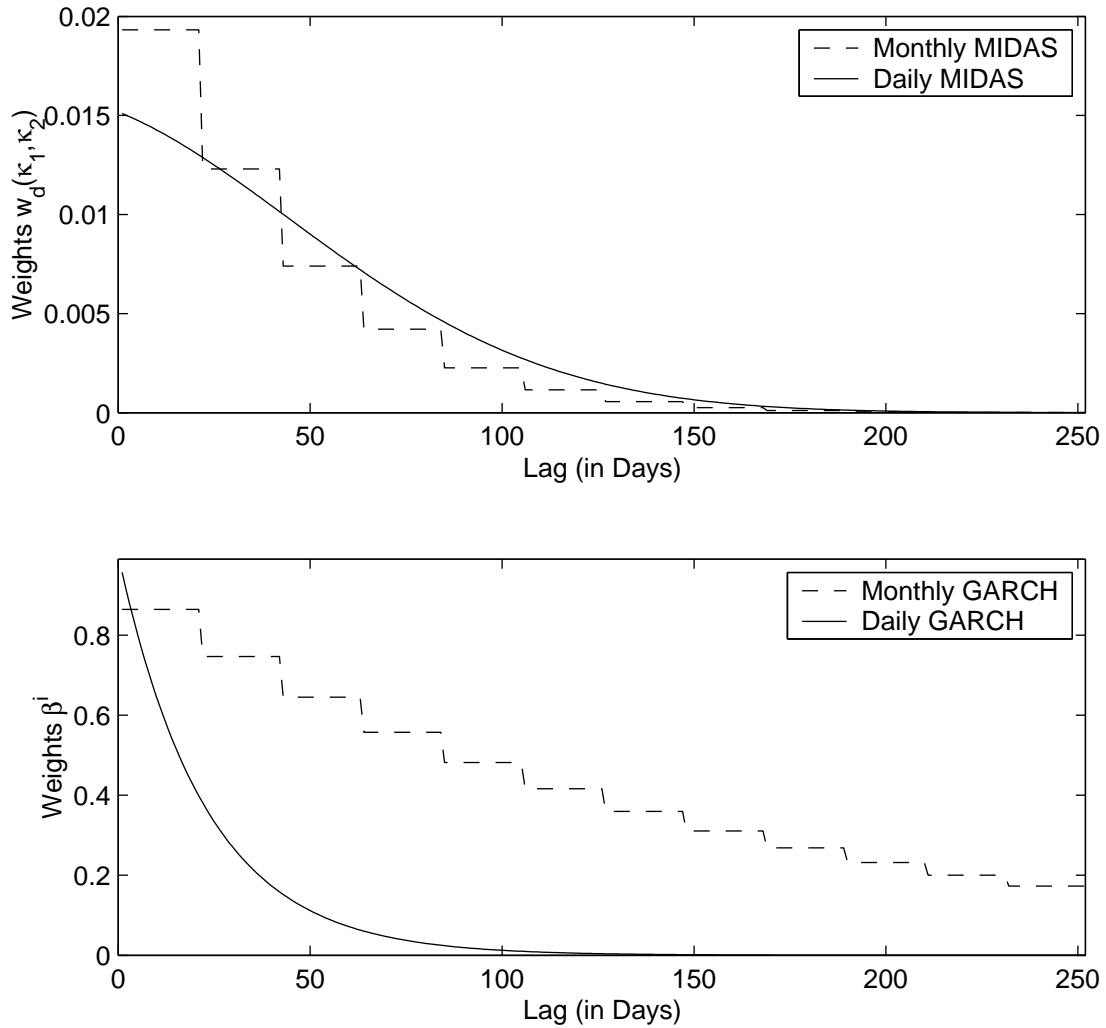


Figure 5: Asymmetric MIDAS Weights

The figure plots the weights that the asymmetric MIDAS estimator (11) and (3) places on lagged daily squared returns, conditional on the sign of the returns. The data sample is 1928:01-2000:12. The weights on the negative shocks ($r < 0$) are calculated by substituting the estimated values of κ_1^- , and κ_2^- into (3). Similarly, the weights on the positive shocks ($r \geq 0$) are calculated by substituting the estimated values of κ_1^+ , and κ_2^+ into (3). The total asymmetric weights, plotted using equation (11), take into account the overall impact of asymmetries on the conditional variance through the parameter ϕ . The estimates of all parameters are shown in Table 8. The symmetric weights from Figure 1 are also plotted for comparison.

