Business Conditions, Market Volatility and Option Prices

Christian Dorion*

Department of Finance, HEC Montreal, Montreal, Canada (christian.dorion@hec.ca)

January 13, 2011

First Draft: February 13, 2009

Abstract We introduce a dynamic volatility model in which stock market volatility varies around a time-varying fundamental level. This fundamental level is determined by macroeconomic risk, quantified using a mixed data sampling (MIDAS) structure to account for changes in the recently introduced Aruoba-Diebold-Scotti (ADS) Business Conditions Index. The new model outperforms the benchmark in fitting asset returns and in pricing options, especially around the 1990-1991 and 2001 recessions. The benchmark model exhibits a counter-cyclical option-valuation bias across all maturities and moneyness levels, and the newly introduced model removes this cyclicality by allowing the conditional expected level of volatility to evolve with business conditions. We extract the volatility premium implied by the model and find that an economically significant 13% of its variation through time can be explained by the impact of macroeconomic risk.

Keywords Business conditions; Macroeconomic risk; Generalized autoregressive conditional heteroscedasticity; Mixed data sampling; Option valuation; Volatility.

EFM Classification Codes: 410, 450

*I am especially indebted to Peter F. Christoffersen and Kris Jacobs for their invaluable suggestions. Thanks to Jan Ericsson, Chayawat Ornthanalai, Redouane Elkamhi, Vadim di Pietro, and seminar participants at HEC Montreal, McGill, HEC Paris, HEC Lausanne, Copenhagen Business School, FCEE-Católica, the U. of Exeter and Rouen Business Schools, Université Laval, and at the Bank of Canada for their enriching comments, and to the *Institut de Finance Mathématique de Montréal* for financial support. All errors are mine.

1 Introduction

Volatility is one of the main determinants of option prices, and much emphasis is placed on improving the dynamic volatility models used to value options. Interestingly, most of these models do not include observables and let volatility mean-revert to a constant level regardless of the current business conditions. However, stock market volatility is robustly found to be highly counter-cyclical, and so the data suggest that a model's volatility process should mean-revert to different levels depending on macroeconomic conditions. Recent work by Engle, Ghysels, and Sohn (2008) builds on this insight and lets volatility vary around a time-varying mean-reversion level that evolves along with the economic fundamentals.[1] These authors find that this fundamental volatility process is significantly related to such factors as inflation and industrial production growth.[2] This study extends their results and investigates the extent to which the impact of business conditions on stock market volatility is reflected in option prices.

Central to this analysis is the new business conditions index recently introduced by Aruoba, Diebold, and Scotti (ADS 2009). Using this index within a mixed data sampling (MIDAS) model,[3] we suggest a model in which volatility varies around a fundamental volatility process that accounts for recent volatility levels and for changes in business conditions. We refer to this model as the MacroHV-MIDAS model, where HV stands for historical volatility. Our model nests Duan's (1995) GARCH model and significantly outperforms it in fitting asset returns and stock market volatility. These results are consistent with the growing consensus that two-factor volatility processes better capture the time-series properties of volatility by accounting separately for transient and high-persistence volatility shocks.[4]

However, two-factor models mostly rely on two latent, autoregressive volatility factors. Whereas the actual drivers of these processes are usually left unidentified, our fundamental volatility process acknowledges that macroeconomic determinants do impact

[1]Regime-switching models, *à la* Hamilton and Susmel (1994), would provide another approach to allowing for different mean-reversion levels. However, using economic fundamentals has the advantage of identifying the determinants of gradual changes in conditional expectations.

[2]Engle, Ghysels, and Sohn (2008) refer to this fundamental volatility process as the secular volatility process.

[3]On MIDAS models, see, for instance, Ghysels, Santa-Clara, and Valkanov (2005); Forsberg and Ghysels (2007); Ghysels, Sinko, and Valkanov (2007); and Engle, Ghysels, and Sohn (2008).

[4]On two-factor models, see, amongst others, Engle and Lee (1999); Andersen, Bollerslev, Diebold, and Ebens (2001); Alizadeh, Brandt, and Diebold (2002); and Engle and Rangel (2008).

conditional volatility expectations. Our results demonstrate that changes in business conditions are an important determinant of the fundamental volatility process. Considering a restricted version of the model in which the business conditions are constrained not to contribute, we find that the constrained model still offers a significantly better fit to asset returns than that of Duan's benchmark model, but offers a significantly worse fit than that of the MacroHV-MIDAS model. Thus, changes in business conditions have an impact on conditional volatility expectations that extends beyond that of recent volatility levels.

These results strengthen those of Engle and Rangel (2008) and Engle, Ghysels, and Sohn (2008) who extensively study physical volatility processes and find them to be counter-cyclical. However, neither study directly discusses the implications for financial derivatives. With the MacroHV-MIDAS model, we propose a risk neutralization that accounts for the correlation between financial returns and changes in business conditions. This risk-neutralized form of the model warrants an analysis of option-pricing errors on twenty years of weekly option data, spanning from June 1988 to December 2007, for a total of 1020 weeks of observations. This is one of the most extensive data sets analyzed in the option pricing literature. We find that our MacroHV-MIDAS model consistently outperforms Duan's (1995) benchmark model in pricing options.[5]

By explicitly accounting for changes in business conditions, our model furthers understanding of the impact of business conditions on option prices. Notably, the analysis of option-pricing errors from a time-series perspective reveals that much of the MacroHV-MIDAS model's improvement over its benchmark arises from its ability to better capture the spot volatility and its dynamics around the 1990-1991 and 2001 recessions. Duan's benchmark model exhibits counter-cyclical biases on options of all maturities and all moneyness levels. By allowing the conditional expected level of volatility to evolve with business conditions, our model is able to remove this cyclicality in the bias, across all maturities and moneyness levels.

The MacroHV-MIDAS model also allows us to measure the contribution of macroeconomic risk to the model-implied volatility premium, defined as the difference between the volatility processes under the risk-adjusted and physical measures. We estimate the price of risk parameter of the MacroHV-MIDAS model using VIX data and then extract the volatility premium implied by the model. The contribution of macroeconomic risk to

[5]Our results are thus consistent with those of, for instance, Christoffersen, Jacobs, Ornthanalai, and Wang (2008), and Christoffersen, Dorion, Jacobs, and Wang (2010). In these articles, however, the long-run volatility component is driven only by innovations to the return process, and thus the model offers no insight regarding the fundamental drivers of expected stock market volatility levels.

the premium is economically significant. While short-term volatility is found to be the main driver of the volatility premium, explaining up to 79% of its variation through time, the volatility impact of changes in business conditions is found to account for a sizeable 13% of the premium's fluctuations.[6]

In summary, this paper shows that a simple dynamic volatility model is able to draw on the informational content of the ADS Business Conditions Index to better capture and understand properties of stock market volatility and of option prices. This ability could prove highly relevant in better understanding the risk inherent to option portfolios throughout the business cycle. In cross sections of option returns, Aramonte (2009) finds macroeconomic uncertainty to be a priced factor, and his results are robust to controlling for a variety of relevant factors such as market and liquidity factors, higher moments of intra-daily returns, and the SMB and HML factors. Our paper suggests that these results are a consequence of an intuitive reality: the expected level of stock market volatility varies along with business conditions.

This article is organized as follows. Section 2 presents the MacroHV-MIDAS model. Section 3 briefly discusses the ADS Business Conditions Index, and compares it to other macroeconomic series of interest before discussing the estimation of the MacroHV-MIDAS model using maximum likelihood. Section 4 uses the maximum likelihood estimates to price twenty years of option data and analyzes the impact of business conditions on the model's implied volatilities. Using option data and nonlinear least-squares estimation, Section 5 refines the results of Section 4 and studies the time-series dynamics of the volatility premium implied by the MacroHV-MIDAS model. Finally, Section 6 concludes.

2 The MacroHV-MIDAS Model

2.1 The Model's Foundations

Dynamic volatility models can be divided into two broad categories: GARCH models and stochastic volatility models. In a stochastic volatility model, the volatility process is driven by unobservable shocks that are imperfectly correlated with shocks to the return process.

[6]In related work, Corradi, Distaso, and Mele (2009) also analyze the volatility risk premium and obtain similar results in a no-arbitrage framework in which the asset price process endogenously determines volatility dynamics that are linked with macroeconomic factors. These authors focus on time-varying risk premia, while the focus of this paper is time-varying conditional expectations; both papers are thus somewhat complementary.

In a GARCH model, the shocks to the volatility process are assumed to result from a deterministic transformation of the return innovations.

On the market, one observes returns and can estimate volatility, but never observes it. In this way, stochastic volatility models are somewhat more realistic. However, by assuming a single source of randomness, GARCH models offer a framework where, given the observable return process $R_t = \log(S_t/S_{t-1})$, where S_t is the stock price, the filtration of the return shocks is trivial. In stochastic volatility models, inferring two unobservable shocks using a single observable is a more difficult task. For this reason, we choose to cast our study in a GARCH framework. The return process is given by

$$R_{t+1} = \mu_{t+1} + \sqrt{h_{t+1}}\varepsilon_{t+1}\,, \qquad \varepsilon_{t+1} \overset{\mathbb{P}}{\sim} \mathcal{N}(0,1)\,, \tag{2.1}$$

$$h_{t+1} = f(\,\cdot \mid \Theta, \mathcal{F}_t\,)\,, \tag{2.2}$$

where μ_{t+1}, the conditional expected return, and h_{t+1}, the conditional variance of returns, are $\mathcal{F}_t-$measurable, and where Θ is a (deterministic) vector of parameters.

Most dynamic volatility models, ergo most GARCH models, eventually mean revert to a constant volatility level, a somewhat undesirable property. The newly introduced GARCH-MIDAS model of Engle, Ghysels, and Sohn (EGS 2008) was introduced to capture a simple intuition: the stock market volatility process should mean revert to different levels depending on macroeconomic conditions. Consider the following multiplicative variance specification, suggested by Engle and Rangel (2008) and EGS:

$$h_{t+1} = g_{t+1}\tau_{t+1}\,, \tag{2.3}$$

$$g_{t+1} = (1-\alpha-\beta) + \alpha g_t \varepsilon_t^2 + \beta g_t\,, \tag{2.4}$$

where τ_t, which refers to the fundamental volatility process, can be interpreted as a time-varying conditional expectation for the level of stock market volatility. The g_t process, which has an unconditional mean of one, accounts for transient shocks to the volatility process by allowing short-run volatility to diverge from fundamental volatility.

Historical volatility, defined as the sum of squared returns over a given horizon,

$$HV_t = \sum_{n=0}^{N-1} R_{t-n}^2\,, \tag{2.5}$$

provides a statistically consistent estimate of stock market volatility.[7] EGS thus suggest the following specification for the fundamental volatility process:

$$\log\left(\tau_{t+1}\right) = m + \theta_{hv} \sum_{k=0}^{K-1} \phi_k\left(w_{hv}\right) HV_{t-k} \,, \tag{2.6}$$

where historical volatilities are computed on a daily basis using the N last daily returns observed on the market.[8] Rather than focusing solely on the last historical volatility measure, this specification smoothly loads on recent observations in the MIDAS spirit. Here, ϕ_k is a Beta weighting scheme,

$$\phi_k(w) = \frac{(1-k/K)^{w-1}}{\sum_{j=0}^{K-1} (1-j/K)^{w-1}} \,,$$

which discards past observations at a rate controlled by w; the larger the w, the faster past historical volatility levels are discarded.[9] While w can be estimated through maximum likelihood, the number of observations used, N in Equation (2.5), just as the number of lags considered, K in Equation (2.6), are selected using the Bayesian information criterion (hereafter referred to as the BIC, Schwarz 1978). In their analysis, EGS find quarterly historical volatilities ($N = 63$ trading days) computed on each day of the past four years ($K = 1008$) to be the best BIC-performing time spans. As our data set largely overlaps theirs, we use these time spans in our analysis.

[7]EGS refer to the estimate of Equation (2.5) as a realized volatility (RV) estimate. Strictly speaking, the estimator is indeed a RV estimate, but some readers may associate RV with the intraday, high-frequency version of the estimator in Equation (2.5). We use the historical volatility (HV) terminology to highlight the low-frequency nature of the RV estimator used here. For more on realized volatility, see, amongst many others, Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Christoffersen, and Diebold (2006), Liu and Maheu (2008), and Andersen and Benzoni (2008).

[8]Equation (2.6) is based on distributed lags of historical volatility measures that are positive by construction. Hence, there is no need to model the *logarithm* of the fundamental variance. However, EGS show that there is little impact from doing so, and this specification has the advantage of allowing for negative values to enter the smoothing function, which proves handy when it comes to using macroeconomic series.

[9]Beta weights are usually parameterized by two parameters; we omitted the one allowing for hump-shaped weightings for the sake of parsimony as preliminary experiments showed it came at little cost. Engle, Ghysels, and Sohn (2008) do the same in their analysis of historical volatility.

2.2 MacroHV-MIDAS: Recent Volatility Levels and Business Conditions

This paper studies the following model:

$$R_{t+1} = r + \lambda \sqrt{h_{t+1}} - \frac{1}{2} h_{t+1} + \sqrt{h_{t+1}} \varepsilon_{t+1} , \tag{2.7}$$

$$h_{t+1} = g_{t+1} \tau_{t+1} , \tag{2.8}$$

$$g_{t+1} = (1 - \alpha(1+\gamma^2) - \beta) + \alpha g_t (\varepsilon_t - \gamma)^2 + \beta g_t , \tag{2.9}$$

$$\log(\tau_{t+1}) = m + \theta_{hv} \sum_{k=0}^{K-1} \phi_k(w_{hv}) HV_{t-k} + \theta_m \sum_{k=0}^{K-1} \phi_k(w_m) \Delta x_{t-k} , \tag{2.10}$$

where x_t is a daily indicator of the quality of business conditions, and where $\Delta x_t = x_t - x_{t-N}$ denotes a measure of the improvement (or deterioration) of business conditions over the last N business days. As we will see in Section 3, the ADS Business Conditions Index provides an appropriate measure of x_t. While the MIDAS framework is, first and foremost, useful for dealing with data sampled at mixed frequencies, it still proves relevant here even though HV_t and Δx_t are both available on a daily basis. Indeed, the functional form of Equation (2.10) allows for a rich lag structure that enables the model to combine past observations of historical volatilities and business conditions in a non-trivial way.

EGS estimate the GARCH-MIDAS model of Equations (2.3)–(2.6) replacing historical volatilities by measures of inflation or of industrial production growth. In the 1953–2004 period, they find the level of these variables to explain 35% and 17% of the expected volatility, respectively. We suggest that changes in business conditions constitute a source of risk that contributes to expected volatility levels beyond what is measured by recent historical volatility levels. In this way, we are essentially pairing the informational content of historical volatilities and business conditions. In order to evaluate the relevance of accounting jointly for both observables, we also consider two restricted versions of the MacroHV-MIDAS model: (i) the HV-MIDAS model, in which we constrain θ_m to be zero, and (ii) the Macro-MIDAS model, in which we constrain θ_{hv} to be zero.

Besides, note that we introduce a γ parameter in the short-run variance specification of Equation (2.9) to allow for the well documented leverage effect (Black 1976), which is particularly important when considering the option-valuation properties of a model for index options.[10] Now, by fixing τ_t to the constant value $e^m = \omega/(1 - \alpha(1+\gamma^2) - \beta)$, one retrieves the nested non-affine GARCH model of Duan (1995) in which h_t simply varies

[10]See, for instance, Nandi (1998), Heston and Nandi (2000), Chernov and Ghysels (2000), Christoffersen and Jacobs (2004), and Christoffersen, Heston, and Jacobs (2006).

around a constant expected variance level parameterized by ω,

$$h_{t+1} = \omega + \alpha h_t(\varepsilon_t - \gamma)^2 + \beta h_t \,. \tag{2.11}$$

In fact, our model could have been designed so as to nest the affine GARCH model of Heston and Nandi (2000). This latter model offers the advantage of admitting a quasi-closed form solution for the value of European calls, and thus relieves the computational burden inherent to a GARCH option-pricing exercise involving Monte Carlo simulations. However, Hsieh and Ritchken (2005) find that the non-affine GARCH specification (hereafter, NGARCH) is superior at removing biases from pricing residuals for all moneyness and maturity categories. These results are supported by Christoffersen, Dorion, Jacobs, and Wang (CDJW 2010) and extended to models allowing for two (additive) variance components and for non-normal innovations. CDJW also show that the NGARCH specification outperforms its affine counterpart from an asset-returns perspective. In results not reported here, we confirm that the superiority of the NGARCH model over its affine counterpart holds in our data sets, both from the asset returns and option valuation standpoints. We, thus, choose the non-affine specification as a more stringent benchmark.

3 Estimating the Model using the ADS Business Conditions Index

The fundamental variance process of the MacroHV-MIDAS model, defined in Equation (2.10), requires a daily measure of business conditions. Before discussing the estimation of the model, this section presents the ADS Business Conditions Index and argues that it is well suited to fulfill the role implied by our characterization of the fundamental variance process.

3.1 The ADS Business Conditions Index

On January 9, 2009, the Federal Reserve Bank of Philadelphia introduced the *ADS Business Conditions Index*, an index that is built on the work of Aruoba, Diebold, and Scotti (ADS 2009). In their paper, ADS develop a sophisticated model that infers latent business conditions from daily term spread observations, weekly initial jobless claims, monthly (non-agricultural) payroll employment, and quarterly real GDP.[11]

[11]The term spread is defined here as the difference between ten-year and three-month Treasury yields. The index published by the Federal Reserve Bank of Philadelphia is based on the ADS paper, but includes some

The ADS procedure cleverly handles missing data, temporal aggregation, complex lag structures and time trends so that they ultimately obtain a linear state-space representation. The authors are thus able to filter out a daily autoregressive process

$$x_t = \varphi x_{t-1} + v_t\,, \qquad v_t \overset{\mathbb{P}}{\sim} \mathcal{N}(0,1)\,, \tag{3.1}$$

that is referred to as the business conditions index. The average value of the ADS index, $\mathbb{E}[x_t]$, is zero, and progressively larger positive values indicate progressively better-than-average conditions. The converse is true for negative values. The v_t innovations are assumed to have unit variance. The first column of Table 1 reports summary statistics on the index, and the lower right panel of Figure 1 plots its value through time, with shaded regions highlighting the NBER recessions. All deep troughs of the index coincide with NBER recessions. In that sense, the index clearly seems to adequately captures the business conditions' relative quality level through time.Note that while the NBER typically announces that the economy reached a peak or a trough several months after it actually occurred, the ADS index value can be updated each time one of its input series is updated. For instance, the ADS Index captured the U.S. economy's December 2007 downturn in real time, while the NBER officially announced it 12 months later, on December 1, 2008.

ADS have to rely on the very simple dynamics of Equation (3.1) for the x_t business conditions index; in particular, the homoskedasticity assumption is necessary for identification. The second column of Table 1, however, makes clear that v_t innovations are all but standard normal. The third and fourth columns of the same table are obtained by fitting a simple GARCH(1,1) to the v_t innovations of Equation (3.1),

$$v_t = \sqrt{h_t^x}\,u_t\,, \qquad u_t \overset{\mathbb{P}}{\sim} \mathcal{N}(0,1)\,, \tag{3.2}$$

$$h_t^x = \omega_x + \alpha_x v_{t-1}^2 + \beta_x h_{t-1}^x\,. \tag{3.3}$$

That is, given the filtered index values, we relax the homoskedasticity assumption and allow innovations to the business conditions index to have a time-varying variance h_t^x. While the u_t are still far from normally distributed, as indicated by the value of the Jarque-Bera statistic, their likelihood and moments are nonetheless more reasonable. The model of Equations (3.1)–(3.3) is, thus, preferable for forecasting purposes.

---

modifications. However, we use the data as kindly provided by Aruoba, Diebold, and Scotti; the series was computed on April 7^{th}, 2008.

3.2 Business Conditions and Other Macroeconomic Series

To justify the use of the ADS Business Conditions Index in our model, this subsection demonstrates that this index conveys some of the informational content usually attributed to indicators such as inflation and industrial production growth. Figure 1 reports some key macroeconomic series throughout the time period considered in this paper.

The lower left panel of Figure 1 plots Ang and Piazzesi's (2003) Inflation factor. This factor is computed from the principal component of three inflation measures based on the consumer price index (CPI), the production price index for finished goods (PPI), and spot market commodity prices as given by the CRB Spot Index (PCOM). For these three indices, we follow Ang and Piazzesi in computing a growth measure, $\log\left(\frac{P_t}{P_{t-12}}\right)$, where P_t is the index level. The resulting series, which are used to compute the principal component, are displayed above the Inflation factor in Figure 1. Analogously, the central panels plot the series that are used to compute Ang and Piazzesi's Real Activity factor. These series are the growth rate—$\log\left(\frac{I_t}{I_{t-12}}\right)$, where I_t is the level—of employment (EMPLOY) and of industrial production (IP), and the unemployment rate (UE). We refer the reader to Ang and Piazzesi (2003) for more details on how the series are processed.[12]

Table 2 reports the correlations between the business conditions index, the rate on the three-month treasury bill, the term spread, the growth rate of employment, the unemployment rate, the (detrended) real GDP, [13] and Ang and Piazzesi's factors. Daily series are sampled monthly for monthly correlations; daily and monthly series are sampled quarterly for quarterly correlations.

The index is strongly and positively correlated with the growth rate of employment (72.4%), with the Real Activity factor (79.4%), and, to a lesser extent, with real GDP (14.6%). As the term spread is the sole daily driver of the index, it is not surprising that it has a relatively strong correlation with the index at the daily level (−27%); interestingly, this correlation remains mainly unchanged by sampling the series monthly (−27.6%). As expected, the index has a strong negative correlation with the unemployment rate (−47.5%). Finally,

[12]The series were obtained from Federal Reserve and the Commodity Research Bureau websites. This paper does not account for the Index of Help Wanted Advertising in Newspapers in its replication of Ang and Piazzesi's Real Activity factor. At first glance, this omission does not yield any notable qualitative difference.

[13]We are interested in the index's correlation with these four series (TS, EMPLOY, UE and GDP) since they are closely related to the index's inputs. Yet, note that the EMPLOY, UE, and GDP series are, here, processed as in Ang and Piazzesi (2003) while Aruoba, Diebold, and Scotti (2009) use levels directly in a much more sophisticated approach.

the index's correlation with the Inflation factor is negative (−7.3%) but, surprisingly, insignificant.

In sum, the business conditions index seems to covary intuitively with many macroeconomic series of interest, while offering the great advantage of accounting for daily innovations.

3.3 Model Estimation: Asset Returns and Stock Market Volatility

Equipped with the ADS Index as a measure of business conditions, we can now estimate the MacroHV-MIDAS model given by Equations (2.7)–(2.10). Table 3 reports the maximum likelihood estimates obtained using S&P 500 returns between January 1968 and December 2007 for the MacroHV-MIDAS model, for the nested HV- and Macro-MIDAS constrained versions, as well as for the NGARCH benchmark from Duan (1995).[14]

All MIDAS models significantly outperform the nested NGARCH model. That the variance is decomposed into two components allows each component to take on one of the two fundamentally different roles that must unduly be assumed by the single component in the NGARCH model. The fundamental variance process captures the long-memory-like properties of stock market volatility. The MacroHV-MIDAS model's fundamental variance process, for instance, has a persistence of only 0.70 basis points (0.70×10^{-4}) below unity. The θ_{hv} and θ_m loadings on historical volatilities and business conditions are positive and negative, respectively. The former captures the persistence of variance following financial turmoils; fundamental variance strongly and positively loads on recent historical variance levels. Fundamental variance loads negatively on recent changes in business conditions thus capturing the counter-cyclical nature of volatility; when business conditions deteriorate, the expected variance level rises.

Note that the θ parameters are smaller (in absolute terms) in the MacroHV-MIDAS model than in the HV- and Macro-MIDAS models. This highlights that the historical volatility levels are not independent from changes in business conditions; when the latter deteriorate, the historical volatility levels tend to rise. Along the same line, both w parameters rise when recent volatilities and changes in business conditions are paired. That is, the MacroHV-MIDAS model weights the recent values of both signals more than the nested

[14] This study uses S&P 500 data (SPX) because of its availability over a long horizon and because options on the SPX have been actively traded for a long time. Returns on major indexes and their volatility are usually highly correlated, and volatility tends to be higher in recessions regardless of the index being considered. We are thus confident that the results obtained in this paper would also obtain using data for other stock indexes.

models that discard the older values slower. Yet, while not orthogonal, the informational content of both signals is clearly not the same; tests based on the likelihoods reported in Table 3 strongly reject the HV- and Macro-MIDAS nested models in favor of the MacroHV-MIDAS model—the likelihood ratio statistics and their p-values are not reported, but the latter are below 1e-6. The MacroHV-MIDAS model would also be selected according to the Bayesian information criterion.

Figure 2 illustrates how the MacroHV-MIDAS model blends both fundamental volatility processes implied by the nested HV- and Macro-MIDAS models. As a reference, we plot a horizontal line at 16.31%, the NGARCH model's expected variance level implied by $\sqrt{252\,\mathbb{E}\left[h_t\right]} = \left(252\omega/(1-\alpha(1+\gamma^2)+\beta)\right)^{\frac{1}{2}}$. The MacroHV-MIDAS model's fundamental volatility process ranges from 11.3% to 25%, at times driven by the contribution of historical volatilities, at times by the contribution of changes in business conditions. The contribution of the ADS Index is remarkably dominant around recessions. The contribution of historical volatilities is most important around the October 1987 crash. Besides, historical volatilities have a surprisingly modest impact in the late 90s, given the relatively high level of volatility observed during the Russian/LTCM crisis.

Including a fundamental variance component gives the short-run variance component the flexibility to allow for greater volatility of variance and to better capture the leverage effect. Indeed, the value of β, the autoregressive variance coefficient, is lower for MIDAS models than for the NGARCH, and lower for the MacroHV-MIDAS than for the two nested ones. In the same line, values of α and γ are higher for the three MIDAS models than for the GARCH(1,1) benchmark, and even more so in the MacroHV-MIDAS case. Altogether, our MacroHV-MIDAS model allows for an 18% higher volatility of variance than that of the NGARCH model (1.862 vs. 1.576) and yields a correlation of -74.3% between the returns and variance processes, about 4.6% greater in magnitude than that of the NGARCH model. By way of comparison, between January 1990 and December 2007, the correlation between excess returns on the S&P 500 and changes in the VIX is -74.1%; considering changes in variance terms, i.e., ΔVIX^2, the correlation is -73.0%.

Table 3 also reports, for both models accounting for business conditions, the correlation between total market innovations and innovations to the business conditions index. This correlation, $\text{Corr}_t(\varepsilon_{t+1}, u_{t+1})$, is about 5% under both the Macro-MIDAS and the MacroHV-MIDAS models. That the observed correlation is positive is somewhat consistent with the preliminary analysis of Section 3.2, which shows that the business conditions index is negatively correlated with Ang and Piazzesi's (2003) Inflation factor, but positively with their Real Activity factor.[15] Five percent may seem low but is consistent with the fact that the business condition index evolves relatively smoothly through time and does not distinguish between expected and unexpected movements of the underlying business conditions. As Cenesizoglu (2005) highlights, the literature agrees that returns mainly react to the surprise content of news and tend to react negatively to positive unanticipated news. Obtaining a low, positive correlation here suggests that increases in the business conditions index reflect heightened expectations about the state of the economy rather than the arrival of unexpected positive news.

4 Option-Valuation Empirics: An Assessment of the Model's Forecasting Abilities

Accounting for business conditions in modeling the physical volatility process does improve a model's capacity to capture the distribution of the volatility of observed returns. Now, we address the extent to which business conditions impact option prices, of which spot volatility is a major determinant. The first step entails analyzing the option-valuation properties of our MacroHV-MIDAS model. We consider twenty years of call option prices from 1988 to 2007, one of the most extensive data sets in the option pricing literature. Even so, our data set covers only two recessions, the early 1990 and 2001 ones.

4.1 Risk Neutralization

In order to analyze the option-pricing properties of the MacroHV-MIDAS model, a risk-neutral form of the model is needed. Typically, GARCH volatility models include a single source of randomness, the ε_{t+1} innovation of Equation (2.1). However, accounting for time-varying business conditions introduces a second source of randomness in the MacroHV-MIDAS model, that is the macroeconomic u_{t+1} innovation of Equation (3.2). Moreover, as observed in Section 3.3, the correlation between the two innovation pro-

[15]Bodie (1976) finds that stock returns covary negatively with both anticipated and unanticipated inflation. Fama (1981) suggests that this negative relationship is driven by real variables covarying positively with stock returns, but negatively with inflation. Yet, the impact of real macro variables on equity returns has found mitigated support for many years. Flannery and Protopapadakis (2002) note that Chen, Roll, and Ross (1986) express their *"embarrassment"* with the situation and that Chan, Karceski, and Lakonishok (1998) *"are at a loss to explain"* the poor performance of macroeconomic factors in explaining stock returns. However, in their own work, Flannery and Protopapadakis, estimating a GARCH model of equity returns, find that these returns are affected by announcements of nominal and real macroeconomic factors.

cesses, $\mathrm{Corr}_t(\varepsilon_{t+1}, u_{t+1})$, is non-zero: market returns are correlated with macroeconomic news. This correlation is however imperfect, i.e., there is a "pure-market" innovation process $z_{t+1} \overset{\mathbb{P}}{\sim} \mathcal{N}(0,1)$, independent from u_{t+1}, such that

$$\varepsilon_{t+1} = \rho u_{t+1} + \sqrt{1-\rho^2} z_{t+1} \,, \tag{4.1}$$

where $\rho = \mathrm{Corr}_t(\varepsilon_{t+1}, u_{t+1})$ by construction.[16]

Our risk neutralization, building on Christoffersen, Elkamhi, Feunou, and Jacobs (2009), relies on the assumption that the equity risk premium on the macroeconomic source of risk is subsumed by the premium on volatility risk and by the contribution of macroeconomic conditions to the volatility process. This assumption was carried over to the model by maintaining the $\mathcal{F}_t$-measurable λh_{t+1} as sole determinant of the equity risk premium in the expected return specification of Equation (2.7). We show in the appendix that, under this assumption, the correlation structure of Equation (4.1) leads to the following risk neutralization of the macroeconomic and pure-market innovation processes:

$$u_t^* = u_t + \rho\lambda \tag{4.2}$$

$$z_t^* = z_t + \sqrt{1-\rho^2}\lambda \,. \tag{4.3}$$

So, the mean shift on each process is proportional to the conditional correlation of that process with total market innovations. Interestingly, if ρ=0, that is, if market shocks and macroeconomic shocks were uncorrelated, the latter would be unaffected by the risk neutralization.

Given Equations (4.2) and (4.3), the risk-adjusted returns of the MacroHV-MIDAS model are given by

$$R_{t+1} = r - \frac{1}{2}\sqrt{\tau_{t+1} g_{t+1}} + \sqrt{\tau_{t+1} g_{t+1}}\left(\rho u_{t+1}^* + \sqrt{1-\rho^2} z_{t+1}^*\right) \tag{4.4}$$

$$g_{t+1} = (1 - \alpha(1+\gamma^2) - \beta) + \alpha g_t \left(\rho u_t^* + \sqrt{1-\rho^2} z_t^* - \gamma - \lambda\right)^2 + \beta g_t \tag{4.5}$$

$$\log(\tau_{t+1}) = m + \theta \sum_{k=0}^{K-1} \phi_k(w_{hv})\, HV_{t-k} + \theta_m \sum_{k=0}^{K-1} \phi_k(w_m)\, x_{t-k} \tag{4.6}$$

$$x_t = \varphi x_{t-1} + \sqrt{h_t^x}\left(u_t^* - \rho\lambda\right) \tag{4.7}$$

$$h_t^x = \omega_x + \alpha_x h_{t-1}^x \left(u_{t-1}^* - \rho\lambda\right)^2 + \beta_x h_{t-1}^x \,, \tag{4.8}$$

[16] This correlation structure implies that market movements do not feed back into the real economy; this implicit assumption is most likely violated in practice (see, for instance, Bernanke, Gertler, and Gilchrist (1999) on the financial accelerator hypothesis), but is necessary here for the sake of simplicity. Corradi, Distaso, and Mele (2009) rely on a similar assumption in a continuous-time setting.

where u_t^* and z_t^* are independent and serially uncorrelated standard normal innovations under the risk-adjusted measure $\mathbb{Q}$.

4.2 Option Valuation Results

Using parameter estimates of Table 3 and the above risk neutralization, it is easy to evaluate call option prices through simulations. We consider cross sections of call options on the S&P 500 index from June 1988 to December 2007. This data set is assembled from three different segments: (i) from June 1988 to December 1989, we use the data from Bakshi, Cao, and Chen (1997); (ii) from January 1990 to December 1995, we use data from Christoffersen, Dorion, Jacobs, and Wang (2010); and (iii) from January 1996 to December 2007, we use OptionMetrics data. For OptionMetrics data, the midpoint between bid and ask prices is used as the option price, and the dividend yield provided by OptionMetrics is used to infer an ex-dividend index level to be used in the option pricing. We also filter zero-volume quotes, and we apply the filtering rules suggested in Bakshi, Cao, and Chen (1997).

Then, for each model, on each Wednesday t_w, we perform Monte Carlo simulations using 2000 paths of $\{z^*_{t_w+\tau,k}\}$ and, when needed, of $\{u^*_{t_w+\tau,k}\}$ in order to price options quoted on week t_w. The shocks are generated using Sobol sequences and we perform Duan and Simonato's (1998) empirical martingale adjustment.[17] Simulations are performed using only the information set up to time t_w, $\mathcal{F}_{t_w}$, with the notable exception that we use parameter estimates from Section 3.3. As these parameters were estimated using *physical* data spanning from 1968 to 2007, and as they are used to price options within that time frame, this is not strictly an out-of-sample exercise. Yet, as the models were estimated without using option data, this exercise is still a stringent exercise in terms of analyzing a model's capacity to properly describe the likely future behavior of volatility.

Aggregate option-valuation metrics are reported in Table 4. The MacroHV-MIDAS model performs, overall, better than all other models. On short- and medium-term options, MIDAS models better capture the volatility smirk than does the benchmark NGARCH model. On long-term options, however, the MacroHV-MIDAS model is outperformed by its benchmark; we will return to this result shortly. Figure 3 sheds some light on these results by casting them in a time-series perspective. The upper panel reports yearly IVRMSE

[17] Christoffersen, Dorion, Jacobs, and Wang (2010) illustrate the accuracy of these simulation settings by comparing the quasi-Monte Carlo results with the exact results computed using the quasi-analytical solutions for the affine model of Heston and Nandi (2000).

values. First, the MacroHV-MIDAS model outperforms the NGARCH in 1988 and slightly less so in 1989. Looking at the performance of the nested HV-MIDAS and Macro-MIDAS models in these two years, we see that the performance of the MacroHV-MIDAS model is driven by the persistent contribution of past historical volatilities; in the single component model, the volatility impact of the Black Monday wears off too quickly.

In 1990 and 1991, the MacroHV-MIDAS model again offers a better fit to option prices, but this time draws on the informational content of the business conditions index. The same observation holds around the second recession in our sample, while the MacroHV-MIDAS model experiences bad performances during the "irrational exuberance" period and throughout the Russian/LTCM crisis. Figures 4 and 5 further illustrate how these results unfold through time and break up the results along the maturity and moneyness dimensions. Figure 4 reports, throughout the sample, 13-week moving averages of the forecast improvement in IVRMSE terms of the MacroHV-MIDAS model over the NGARCH model; Figure 5 similarly reports the MacroHV-MIDAS model's bias at forecasting implied volatilities.[18] In short, while the benchmark NGARCH model exhibits strongly counter-cyclical biases, the MacroHV-MIDAS model removes this cyclicality, especially (i) for longer maturity options and (ii) as we go from ITM to OTM calls.

As previously noted, in IVRMSE terms, it is only on long-term options that the MacroHV-MIDAS model shows a worse fit to implied volatilities than its benchmark does. However, the MacroHV-MIDAS model actually fits the implied volatilities of long-term options dramatically better than its benchmark around recessions. On the other hand, as evidenced in the lower-left panel of Figure 4, the model shows a significant bias on long-term options in the late 90s and, for this subset of options, this cancels the improvements realized around recessions. In fact, the counter-performance of the MacroHV-MIDAS model over the late 90s, evidenced in Figure 3, is shown in Figure 5 to be due to large negative bias at all maturities and over all moneyness levels during this period.

A second look at Figure 2 can provide us with an intuition why this is so. Indeed, while the VIX reaches all-time highs during the Russian/LTCM crisis, the time-varying volatility expectations captured by the fundamental volatility process are rising at a very slow pace. It is likely that the fundamental volatility process, as specified, is too smooth to account for drastic changes in the market's expectations about future volatility. Besides,

[18]We plot 13-week moving averages solely for the sake of clarity; the weekly measures are very noisy, especially for short-term options and for ITM options. The averages reported above each subplot are, however, based on these weekly measures.

while stock market volatility is relatively high during this period, business conditions are better than average. As defined here, the fundamental volatility process simply sums, in the log-volatility domain, the impact of both historical volatilities and recent changes in business conditions. It is possible that, during the late 90s, this naive blend puts too much emphasis on the better-than-average business conditions and too little on recent volatility levels.

Nonetheless, on average, the MacroHV-MIDAS model exhibits an implied-volatility bias of smaller magnitude than that of its benchmark, for any maturity or moneyness, as evidenced by the averages reported along with Figure 5. However, the MIDAS model still underprices all options on average, and this is even more obvious on short-term and in-the-money calls. This underpricing is likely to be a consequence of the conditional normality assumption, but could also be due to the choice of a linear pricing kernel; see Christoffersen, Elkamhi, Feunou, and Jacobs (2009) for further discussion on these issues.

4.3 The Impact of Business Conditions

Along with Figure 4, we report average improvements of the MacroHV-MIDAS model over its benchmark, conditioned on whether a given week falls within an expansion or a recession period. Average improvements over recessions, reproduced in the Panel A of Table 5, are highly statistically significant. However, of the 1020 weeks in our data set of options, only 72 (approximately 7.1%) fall within a recession. To further study the extent of the MacroHV-MIDAS model's improvements over the benchmark and detail the role of business conditions in these improvements, Panel A of Table 5 also reports statistics conditional on the contemporary level of the ADS Index. That is, instead of relying on a NBER recession dummy to determine that a week falls within a period of bad business conditions, we compute centered quarterly moving averages of the ADS on each Wednesday t,

$$x_t^{(63)} = \frac{1}{63} \sum_{s=t-31}^{t+31} x_s \,. \tag{4.9}$$

As the index average is theoretically zero, we will say that a week is in the middle of a quarter with "bad" business conditions when $x_t^{(63)} < 0$, "severe" business conditions when $x_t^{(63)} < -1$, "extreme" business conditions when $x_t^{(63)} < -1.5$. In our sample, 380 weeks (37.3% of the 1020 weeks) are exposed to bad business conditions, 178 (17.5%) to severe ones, and 45 (4.4%) to extreme ones. Across all maturities and moneyness levels, Panel A of Table 5 reports that the improvement brought by the MacroHV-MIDAS model increases as business conditions deteriorate.

Besides, even if we don't have option data prior to 1988, we can use the models to price a synthetic at-the-money option through time. On each day from January 1968 to December 2007, we use the NGARCH and MacroHV-MIDAS models to price a 30-day option with its strike equal the the index value on that day. The time series of NGARCH implied volatilities for such an option is reported in the upper-left panel of Figure 6, and statistics for this time series are reported in the first row of Table 5's Panel B. Statistics on the difference between the MacroHV-MIDAS model's implied volatilities and those of the NGARCH are reported on the second row of Panel B and are broken down in the mid- and lower-left panels of Figure 6. The right-column panels of Figure 6 and the remainder of Table 5's Panel B report the same results but in the price domain.[19]

As we have seen in Figure 5, both models consistently underprice options. Thus, to improve on the NGARCH model, the MacroHV-MIDAS model should predict higher implied volatilities than those of NGARCH, and Table 5's Panel B reports that it does so on average. Interestingly, this implied-volatility difference is increasing as business conditions deteriorate. For example, under extreme business conditions, the 30-day, at-the-money implied volatility of the MacroHV-MIDAS model is 1.6% higher than that implied by the NGARCH model. In terms of option prices, this translate in a 9.1% higher option price on average, a difference of sizable economic importance. In sum, Table 5 shows that business conditions indeed play a key role the ability of MacroHV-MIDAS model to outperform its benchmark.

Table 6 sheds further light on the role played by business conditions in the MacroHV-MIDAS model's option-valuation performance. In this table, the MacroHV-MIDAS model is compared to the HV-MIDAS, rather than the NGARCH, in order to better grasp the marginal impact of accounting for changes in business conditions. Moreover, the expansion results are unfolded in three parts: "good", "very good" and "exceptional" business conditions when $x_t^{(63)}$ respectively is above 0, 1, and 1.5. Overall, the MacroHV-MIDAS model improves on the HV-MIDAS when business conditions are bad; in relative terms, the former model cuts the benchmark's IVRMSE by 7.1% more than the latter model. When business conditions are good, however, the HV model outperforms the MacroHV by 1.2%.

[19]The average NGARCH IV reported in Panel B is higher in recessions than in expansion, as expected. Interestingly, the average call price is higher in expansion. While this might seem contradictory, it is actually due to the fact that the underlying, the S&P 500 index is, on average, higher in expansion than in recessions.

Interestingly, when business conditions are very good or exceptional, this figure is once more reversed and the MacroHV does better than the HV by 3.7% and 7.2% respectively. Panel A of Table 6 illustrates that this pattern is rather robust to maturities and moneyness levels.

As highlighted by Panel B of Table 6, the MacroHV model predicts higher implied volatilities than the nested HV model when business conditions are bad, and lower ones when business conditions are good. This difference between the two models evolves monotonically as business conditions improve from extreme to exceptional, which is consistent with the smooth, log-linear fashion in which the MIDAS specification accounts for changes in the index in the MacroHV model. The results in Panel A, in particular the under-performance of the MacroHV model when business conditions are good but not very good, are thus probably highlighting, once more, that the log-linear mix of recent historical volatility levels and changes in business conditions is suboptimal. A possible explanation of the above pattern could be that fundamental volatility does not decrease when business conditions are only mildly good since economic agents might still perceive some macroeconomic risk, a phenomenon that the current fundamental volatility specification cannot accommodate.

5 Model-Implied Volatility Premium

Under certain assumptions, the volatility premium, defined as the difference between expected future volatility under the risk-neutral and the objective measure, can be directly related to the risk aversion of a representative agent.[20] As the risk neutralization of the MacroHV-MIDAS model is set under Christoffersen, Elkamhi, Feunou, and Jacobs' (2009) framework, no assumption is made with respect to the representative agent or its utility function. It is nonetheless important to assess whether the volatility premium generated by the MacroHV-MIDAS model has coherent properties. For instance, this premium was consistently found to increase when the stock market volatility rises, and some found it to be even more counter-cyclical than volatility itself.[21] Besides, the MacroHV-MIDAS model splits volatility between its time-varying mean-reversion level and the short-run excess volatility and, moreover, allows to easily isolate the contribution of macroeconomic risk to

[20]See, amongst others, Heston (1993), Eraker (2007), and Bollerslev, Gibson, and Zhou (2010).

[21]See, for instance, Bollerslev, Gibson, and Zhou (2010), Corradi, Distaso, and Mele (2009), and Bollerslev, Tauchen, and Zhou (2009).

the volatility level. These abilities prove interesting when it comes to better understanding the drivers of the premium.

Since the residual implied-volatility bias of the MacroHV-MIDAS model reported in Figure 4 is likely to be partly due to an underestimation of the volatility premium, we refine the model's estimation before analysing the premium. To do so, we perform nonlinear least squares to minimize the tracking error between the model's volatility forecasts under the risk-adjusted measure and the VIX. See Appendix B for details. The resulting daily volatility premium series is displayed in the lower panel of Figure 7.

A simple glance at the figure confirms that the MacroHV-MIDAS model-implied volatility premium is very strongly correlated with the current volatility level and that it is counter-cyclical. The average value of the extracted volatility premium is 3.10%, with a standard deviation of 1.19%. On actual data, over the June 1988 to December 2007 period, the average value of the VIX was 19.70% and the standard deviation of excess returns of 15.61%, for an average premium of 4.09%; over the 1990-2007 period, the average value of the VIX was 18.97% and the standard deviation of excess returns of 15.80%, for an average premium of 3.17%. The model seems to slightly underestimate the premium on average. Much of this underestimation seems to be due to the negative bias observed during the late 90s, in line with our analysis in Section 4.2.

Panel C of Table 7 reports the results of nine linear regressions with, as regressand, our model-implied volatility premium. Regressors are demeaned and, in order to account for the likely strong autocorrelation of the residuals, t-stats are computed using Newey-West standard errors with a lag of 63, corresponding to one quarter of trading days.[22] First, we control for the annualized and demeaned fundamental volatility level under the objective measure. By itself, the current fundamental volatility level accounts for 31.8% of the variation in the premium through time. All else equal, a one-percent increase in the annualized fundamental volatility level causes a statistically significant 24.5 bps increase in the volatility premium. Relative to the 3.10% mean of the extracted volatility premium process, this is an 8% increase (24.5 / 310) and is thus economically significant.

In our framework, macroeconomic risk impacts stock market volatility through the contribution of changes in business conditions to the time-varying volatility mean-reversion level. This contribution is easily extracted by setting θ_{hv} to zero in our model, which effectively nullifies the contribution of recent historical volatilities to the fundamental volatil-

[22]The choice of a quarter-long lag is arbitrary and intended to be very conservative; in our case, $T = 4936$, so that $\lfloor 4(T/100)^{0.25} \rfloor = 10$ would be the lag suggested by Newey and West (1994).

ity level. A one-percent increase in this measure of macroeconomic risk translates into a 45.2 bps increase in the volatility premium, which is a 14.6% change relative to the 3.10% mean, and macroeconomic risk explains 12.9% of the time variation in the volatility premium process. On the other hand, when we focus our attention on how recent historical volatilities contribute to the fundamental volatility level, that same one-percent increase translates into a 26.3 bps increase in the premium, a 8.5% change relatively to the 3.10% mean. That the volatility premium is more sensitive to each of the restricted signals than to the overall fundamental volatility level suggests, once again, that there might be more efficient ways to combine the informational content of historical volatilities and that of business conditions than simply summing them in the log-volatility domain.

In addition, we regress the premium on the standardized value of $\sqrt{g_t}$. Remember that when this control is above its mean, the current level of physical volatility is above its fundamental, expected level. By itself, this excess volatility explains 79.3% percent of the variation in the premium. A one-standard deviation increase in $\sqrt{g_t}$ causes a dramatic 105.9 bps increase in the volatility premium, a 33.9% move relative to the mean of the premium process. Note though that $\sqrt{g_t}$ is within one-standard deviation of its mean on close to 87% of the trading days under consideration. Nonetheless, short-run volatility is undeniably the main driver of our model's volatility premium.

Finally, to assess the cyclicality of the premium process, we consider a dummy variable that has value one when a day falls within an NBER recession and zero otherwise. As can be seen from the first column of Panel C, the average volatility premium on a typical NBER recession day is 44.7% greater than the overall average value, an increase of 138.5 bps over the 310 bps mean. The MacroHV-MIDAS model-implied premium is thus strongly counter-cyclical. In a regression of the premium on the fundamental volatility level and on the NBER dummy, the coefficient on the dummy implies that even when controlling for time-varying volatility expectations, the premium is higher on a recession day by a sizeable 54.3 bps (17.5% relative to the mean). Although the magnitude of the effect could be seen as economically significant, the NBER coefficient is scarcely statistically significant at the 10% level. On the other hand, when controlling for both fundamental and time-varying volatility levels, the coefficient on the NBER dummy is statistically significant at the 5% level, but economically more modest at 14 bps, 4.5% of the 310 bps mean. Nonetheless, the loading on the NBER dummy tells us that, all else equal, the volatility premium is still higher on a typical recession day than what can be explained by the physical volatility components.

In sum, the MacroHV-MIDAS model-implied volatility premium (i) is mainly driven by short-run volatility effects; (ii) is strongly counter-cyclical and a sizeable portion of this counter-cyclicality is driven by changes in expectations with respect to the long-run volatility level (as modeled by the fundamental variance process); and (iii) the premium is slightly more counter-cyclical than what is explained by short-run and long-run volatility effects.

6 Conclusion

This paper introduces the MacroHV-MIDAS model, a dynamic volatility model accounting for both financial and macroeconomic sources of fundamental volatility. This model is shown to outperform the NGARCH benchmark in fitting asset returns and pricing options, especially around the 1990-1991 and 2001 recessions. In particular, the MacroHV-MIDAS model improves on the benchmark's option-valuation abilities by mitigating the counter-cyclicality of its implied-volatility bias, across all maturity and moneyness levels. The MacroHV-MIDAS model also allows us to isolate the contribution of macroeconomic risk to the volatility premium, and this contribution is found to account for a sizeable 13% of the variation in the premium through time.

This work offers several avenues for further research. For instance, conducting our analysis in a stochastic volatility framework would allow us to assess the extent of the relationship between the macroeconomic shocks entering our fundamental volatility process and the unobservable volatility shocks inherent to stochastic volatility models. Apart from that, incorporating analyst forecasts or survey results in the business conditions' forecasting model could further improve the abilities of the MacroHV-MIDAS model to explain observed option prices. Buraschi, Trojani, and Vedolin (2009) also suggest that dispersion in analyst forecasts is strongly related to implied volatility levels. Otherwise, once it is established that business conditions impact option prices, option data could eventually be used to infer market expectations of future business conditions.

Another line of investigation would be to refine how the informational content of historical volatilities and business conditions are combined to model the fundamental volatility process. Our work uses historical volatilities based on daily returns; when intraday data are available, intraday realized volatilities could prove more reactive to current market conditions. Moreover, the MacroHV-MIDAS model simply sums the impact of historical volatilities and business conditions in the log-volatility domain. However, given that responses to macroeconomic news differ depending on the current state of the economy, and as our results suggest that worsening business conditions increase option prices more than improving business conditions lower them, it is likely that an approach allowing for further nonlinearities would prove fruitful. Finally, in our opinion, a study of how higher moments of the stock returns' distribution evolve with changing business conditions could further our understanding of the volatility premium and of its time-series properties. The Lévy GARCH framework of Ornthanalai (2009) or the mixed normal heteroskedasticity framework of Rombouts and Stentoft (2009) could prove to be usuful in that regards.

References

Alizadeh, S., M. Brandt, and F. Diebold (2002). Range-Based Estimation of Stochastic Volatility Models. *Journal of Finance 57*, 1047–1091.

Andersen, T. G. and L. Benzoni (2008, July). Realized Volatility. FRB of Chicago Working Paper No. 2008-14.

Andersen, T. G., T. Bollerslev, P. F. Christoffersen, and F. X. Diebold (2006). Volatility and Correlation Forecasting. In G. ELLIOTT, C. W. J. GRANGER, and A. TIMMERMANN (Eds.), *Handbook of Economic Forecasting*, Volume 1, pp. 777–878. North-Holland.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens (2001). The Distribution of Realized Stock Return Volatility. *Journal of Financial Economics 61*, 43–76.

Ang, A. and M. Piazzesi (2003). A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables. *Journal of Monetary Economics 50*, 745–787.

Aramonte, S. (2009). Macroeconomic Uncertainty and Option Returns. Working Paper, London Business School.

Aruoba, S., F. Diebold, and C. Scotti (2009). Real-Time Measurement of Business Conditions. *Journal of Business and Economic Statistics 27*, 417–427.

Bakshi, G., C. Cao, and Z. Chen (1997). Empirical Performance of Alternative Option Pricing Models. *Journal of Finance 52*, 2003–2049.

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999). Chapter 21 The Financial Accelerator in a Quantitative Business Cycle Framework. Volume 1, Part 3 of *Handbook of Macroeconomics*, pp. 1341 – 1393. Elsevier.

Black, F. (1976). Studies of Stock Price Volatility Changes. In *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association*, pp. 177–181.

Bodie, Z. (1976). Common Stocks as a Hedge Against Inflation. *Journal of Finance 31*, 459–470.

Bollerslev, T., M. S. Gibson, and H. Zhou (2010). Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities. *Journal of Econometrics 160*, 235–245.

Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected Stock Returns and Variance Risk Premia. *The Review of Financial Studies*, Forthcoming.

Bollerslev, T. and J. Wooldridge (1992). Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances. *Econometric Reviews 11*, 143–172.

Brown, S. J., W. N. Goetzmann, and S. A. Ross (1995). Survival. *Journal of Finance 50*(3), 853–873.

Buraschi, A., F. Trojani, and A. Vedolin (2009). Economic Uncertainty, Disagreement, and Credit Markets. Working Paper.

Cenesizoglu, T. (2005). Risk and Return Reaction of the Stock Market to Public Announcements about Fundamentals: Theory and Evidence. Working Paper.

Chan, L. K. C., J. Karceski, and J. Lakonishok (1998). The Risk and Returns from Factors. *Journal of Financial and Quantitative Analysis 33*, 159–188.

Chen, N., R. Roll, and S. Ross (1986). Economic Forces and the Stock Market. *Journal of Business 59*, 383–403.

Chernov, M. and E. Ghysels (2000). A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Option Valuation. *Journal of Financial Economics 56*, 407–458.

Christoffersen, P. F., C. Dorion, K. Jacobs, and Y. Wang (2010). Volatility Components, Affine Restrictions and Non-Normal Innovations. *Journal of Business and Economic Statistics 28*, 483–502.

Christoffersen, P. F., R. Elkamhi, B. Feunou, and K. Jacobs (2009). Option Valuation with Conditional Heteroskedasticity and Non-Normality. *Review of Financial Studies*, Forthcoming.

Christoffersen, P. F., S. Heston, and K. Jacobs (2006). Option Valuation with Conditional Skewness. *Journal of Econometrics 131*, 253–284.

Christoffersen, P. F. and K. Jacobs (2004). Which Volatility Model for Option Valuation? *Management Science 50*, 1204–1221.

Christoffersen, P. F., K. Jacobs, C. Ornthanalai, and Y. Wang (2008). Option Valuation with Long-run and Short-run Volatility Components. *Journal of Financial Economics 90*, 272–297.

Corradi, V., W. Distaso, and A. Mele (2009). Macroeconomic Determinants of Stock Market Returns, Volatility and Volatility Risk-Premia. Working Paper.

DeLong, J. B. and K. Magin (2009). The U.S. Equity Return Premium: Past, Present, and Future. *Journal of Economic Perspectives 23*, 193–208.

Duan, J.-C. (1995). The GARCH Option Pricing Model. *Mathematical Finance 5*, 13–32.

Duan, J.-C. and J.-G. Simonato (1998). Empirical Martingale Simulation for Asset Prices. *Management Science 44*, 1218–1233.

Engle, R., E. Ghysels, and B. Sohn (2008). On the Economic Sources of Stock Market Volatility. Working Paper.

Engle, R. F. and G. Lee (1999). A Permanent and Transitory Component Model of Stock Return Volatility. In R. ENGLE and H. WHITE (Eds.), *Cointegration, Causality, and Forecasting: a Festschrift in Honor of Clive W.J. Granger*, New York, pp. 475–497. Oxford University Press.

Engle, R. F. and J. G. Rangel (2008). The Spline-GARCH Model for Low-Frequency Volatility and Its Global Macroeconomic Causes. *Review of Financial Studies 21*, 1187–1222.

Eraker, B. (2007). The Volatility Premium. Working Paper, Duke University, Department of Economics.

Fama, E. (1981). Stock Returns, Real Activity, Inflation and Money. *The American Economic Review 71*, 545–565.

Flannery, M. J. and A. A. Protopapadakis (2002). Macroeconomic Factors *Do* Influence Aggregate Stock Returns. *Review of Financial Studies 15*, 751–782.

Forsberg, L. and E. Ghysels (2007). Why Do Absolute Returns Predict Volatility So Well? *Journal of Financial Econometrics 5*, 31–67.

Ghysels, E., P. Santa-Clara, and R. Valkanov (2005). There is a Risk-Return Trade-Off After All. *Journal of Financial Economics 76*, 509–548.

Ghysels, E., A. Sinko, and R. Valkanov (2007). MIDAS Regressions: Further Results and New Directions. *Econometric Reviews 26*, 53–90.

Gibson, R. and E. S. Schwartz (1990). Stochastic Convenience Yield and the Pricing of Oil Contingent Claims. *Journal of Finance 45*, 959–976.

Hamilton, J. and R. Susmel (1994). Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics 64*(1-2), 307–333.

Heston, S. L. (1993). A Closed Form Solution for Options with Stochastic Volatility with Applications to Bonds and Currency Options. *Review of Financial Studies 6*, 327–343.

Heston, S. L. and S. Nandi (2000). A Closed-Form GARCH Option Valuation Model. *Review of Financial Studies 13*, 585–626.

Hsieh, K. C. and P. Ritchken (2005). An Empirical Comparison of GARCH Option Pricing Models. *Review of Derivatives Research 8*(3), 129–150.

Li, H. and Y. Xu (2002). Survival Bias and the Equity Premium Puzzle. *Journal of Finance 57*(5), 1981–1995.

Liu, C. and J. M. Maheu (2008). Are There Structural Breaks in Realized Volatility? *Journal of Financial Econometrics 6*(3), 326–360.

Mehra, R. and E. Prescott (1985). The Equity Puzzle. *Journal of Monetary Economics 15*, 145–61.

Nandi, S. (1998). How Important is the Correlation Between Returns and Volatility in a Stochastic Volatility Model? Empirical Evidence from Pricing and Hedging in the S&P 500 Index Options Market. *Journal of Banking and Finance 22*, 589–610.

Newey, W. K. and K. D. West (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica 55*, 703–708.

Newey, W. K. and K. D. West (1994). Automatic Lag Selection in Covariance Matrix Estimation. *The Review of Economic Studies 61*, 631–653.

Ornthanalai, C. (2009). A New Class of Asset Pricing Models with Lévy Processes: Theory and Applications. Working Paper.

Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics 63*, 3–50.

Pollard, M. C. (2009). Luck and the US Equity Premium. Working Paper.

Rombouts, J. V. and L. Stentoft (2009, May). Bayesian Option Pricing Using Mixed Normal Heteroskedasticity Models. Working Paper.

Rosenberg, J. V. and R. F. Engle (2002). Empirical Pricing Kernels. *Journal of Financial Economics 64*, 341–372.

Schwarz, G. (1978). Estimating the Dimension of a Model. *Annals of Statistics 6*, 461–464.

Appendix

A Risk Neutralization

We consider a GARCH model of the form

$$R_{t+1} = r + \lambda_{t+1}\sqrt{h_{t+1}} - \tfrac{1}{2}h_{t+1} + \sqrt{h_{t+1}}\Big(\rho u_{t+1} + \sqrt{1-\rho^2}z_{t+1}\Big) \tag{A.1}$$

$$h_{t+1} = f\left(\,\cdot \mid \Theta, \mathcal{F}_t\right)\,, \tag{A.2}$$

where λ_{t+1} and h_{t+1} are $\mathcal{F}_t$-measurable, and where $\{\,u_t\,\}$ and $\{\,z_t\,\}$ are independent and serially uncorrelated innovation processes. To formally demonstrate the risk neutralization of u_t and z_t as introduced in Equations (4.2) and (4.3), we here draw on Christoffersen, Elkamhi, Feunou, and Jacobs (CEFJ 2009) treatment of two-shocks stochastic volatility models (see CEFJ's Section 7). Note that our model is, however, fundamentally different from a stochastic volatility model in that, here, the "second" shock, u_{t+1}, does not contemporaneously impact the variance but the mean of the return process. We will return to the implications of this fundamental difference shortly.

First, we write the risk neutralization of our return process in terms of the risk neutralization of the bivariate, uncorrelated normal innovations $\{\,u_t, z_t\,\}$ using the following Radon-Nikodym derivative:

$$\xi_\tau \equiv \frac{d\mathbb{Q}}{d\mathbb{P}}\bigg|\mathcal{F}_\tau = \exp\left\{-\sum_{t=1}^{\tau}\left(\eta_{u,t}u_t + \eta_{z,t}z_t + \Psi_t^{u,z}\left(\eta_{u,t},\eta_{z,t}\right)\right)\right\}\,, \tag{A.3}$$

where $\Psi_t^{u,z}$ is natural logarithm of the moment-generating function of the $\{\,u_t, z_t\,\}$ pairs, that is,

$$\Psi_t^{u,z}\left(\eta_u,\eta_z\right) = \tfrac{1}{2}\left(\eta_u^2 + \eta_z^2\right)\,. \tag{A.4}$$

For the probability measure $\mathbb{Q}$ defined by Radon-Nikodym derivative of Equation (A.3) to be an equivalent martingale measure (EMM), it must be the case that

$$\begin{aligned}1 &= \mathbb{E}^{\mathbb{Q}}_{t\text{-}1}\left[\frac{S_t}{S_{t-1}}\Big/\frac{B_t}{B_{t-1}}\right] = \mathbb{E}^{\mathbb{P}}_{t\text{-}1}\left[\frac{\xi_t}{\xi_{t-1}}\frac{S_t}{S_{t-1}}\Big/\frac{B_t}{B_{t-1}}\right]\\ &= \mathbb{E}^{\mathbb{P}}_{t\text{-}1}\left[\exp\left\{-\eta_{u,t}u_t - \eta_{z,t}z_t - \Psi_t^{u,z}\left(\eta_{u,t},\eta_{z,t}\right)\right\}\exp\left\{\lambda_t\sqrt{h_t} - \tfrac{1}{2}h_t + \sqrt{h_t}\Big(\rho u_t + \sqrt{1-\rho^2}z_t\Big)\right\}\right]\end{aligned} \tag{A.5}$$

or, equivalently,

$$0 = \Psi_t^{u,z}\left(\eta_{u,t} - \rho\sqrt{h_t}\,, \eta_{z,t} - \sqrt{(1-\rho^2)h_t}\right) - \Psi_t^{u,z}\left(\eta_{u,t},\eta_{z,t}\right) + \lambda_t\sqrt{h_t} - \tfrac{1}{2}h_t\,, \tag{A.6}$$

which boils down to

$$\rho\eta_{u,t} + \sqrt{1-\rho^2}\eta_{z,t} = \lambda_t\,. \tag{A.7}$$

Equation (A.7) admits an infinity of solutions. Yet, as highlighted above, our model has the specificity that both shocks affect the mean of the return process. Thus, the bivariate normal shocks can be seen as blending into a single stream of standard normal innovations $\{ \varepsilon_t \}$ and Equation (A.1) is equivalent to

$$R_{t+1} = r + \lambda_{t+1}\sqrt{h_{t+1}} - \tfrac{1}{2}h_{t+1} + \sqrt{h_{t+1}}\varepsilon_{t+1}\,. \tag{A.8}$$

This last equation is that of Duan's (1995), which is a special case of CEFJ for which the (linear) Radon-Nikodym derivative can be written as

$$\xi_\tau \equiv \frac{d\mathbb{Q}}{d\mathbb{P}}\bigg|\mathcal{F}_\tau = \exp\left\{-\sum_{t=1}^{\tau}\left(\eta_t\varepsilon_t + \Psi_t^{\varepsilon}(\eta_t)\right)\right\}, \tag{A.9}$$

where Ψ_t^{ε} is natural logarithm of the moment-generating function of the $\{ \varepsilon_t \}$ innovations, that is, $\Psi_t^{\varepsilon}(\eta) = \frac{1}{2}\eta^2$. Again, for the $\mathbb{Q}$ measure defined by Equation (A.9) to be an EMM, it must be that

$$\begin{aligned} 1 &= \mathbb{E}_{\text{t-1}}^{\mathbb{Q}}\left[\frac{S_t}{S_{t-1}}\Big/\frac{B_t}{B_{t-1}}\right] = \mathbb{E}_{\text{t-1}}^{\mathbb{P}}\left[\frac{\xi_t}{\xi_{t-1}}\frac{S_t}{S_{t-1}}\Big/\frac{B_t}{B_{t-1}}\right] \\ &= \mathbb{E}_{\text{t-1}}^{\mathbb{P}}\left[\exp\left\{-\eta_t\varepsilon_t - \Psi_t^{\varepsilon}(\eta_t)\right\}\exp\left\{\lambda_t\sqrt{h_t} - \tfrac{1}{2}h_t + \sqrt{h_t}\varepsilon_t\right\}\right] \qquad \text{(A.10)} \\ \Leftrightarrow 0 &= \Psi_t^{\varepsilon}\left(\eta_t - \sqrt{h_t}\right) - \Psi_t^{\varepsilon}(\eta_t) + \lambda_t\sqrt{h_t} - \tfrac{1}{2}h_t\,, \qquad \text{(A.11)} \end{aligned}$$

which implies that $\eta_t = \lambda_t, \forall t$. Now, for Equations (A.3) and (A.9) to describe the same Radon-Nikodym derivative, it must be that

$$\begin{aligned} \xi_\tau &= \exp\left\{-\sum_{t=1}^{\tau}\left(\lambda_t\varepsilon_t + \tfrac{1}{2}\lambda_t^2\right)\right\} \\ &= \exp\left\{-\sum_{t=1}^{\tau}\left(\lambda_t\rho u_t + \lambda_t\sqrt{1-\rho^2}z_t + \tfrac{1}{2}\overbrace{\left(\rho^2\eta_{u,t}^2 + 2\rho\sqrt{1-\rho^2}\eta_{u,t}\eta_{z,t} + (1-\rho^2)\eta_{z,t}^2\right)}^{\text{Using Equation (A.7)}}\right)\right\} \\ &= \exp\left\{-\sum_{t=1}^{\tau}\left(\eta_{u,t}u_t + \eta_{z,t}z_t + \Psi_t^{u,z}(\eta_{u,t},\eta_{z,t})\right)\right\} \end{aligned}$$

where the last equality holds if, and only if, for all t,

$$\eta_{u,t} = \rho\lambda_t \quad \text{and} \quad \eta_{z,t} = \sqrt{1-\rho^2}\lambda_t\,. \tag{A.12}$$

We thus have that $u_t^* = u_t + \rho\lambda_t$, $z_t^* = z_t + \sqrt{1-\rho^2}\lambda_t$, and

$$\varepsilon_t^* = \varepsilon_t + \lambda_t = \rho u_t^* + \sqrt{1-\rho^2}z_t^*\,. \tag{A.13}$$

B Refining Model Estimation Using the VIX

The VIX levels reflect the market's one-month-ahead expectation of the (risk-neutral) implied volatility process. For the MacroHV-MIDAS model, this implied-volatility expectation is $\mathbb{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right]$ and can be easily computed using Monte-Carlo integration. In order to study the properties of the volatility premium, $\mathbb{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right] - \mathbb{E}_t^{\mathbb{P}}\left[\sqrt{h_{t+21}}\right]$, implied by our model, we must first ensure that the bias of its implied volatility process is minimized; that is, $\mathbb{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right]$ should track the VIX level as closely as possible. However, as can be observed in Figure 4, using the ML parameter estimates the model systematically underprices short-maturity options.

The only difference between the objective and risk-neutral volatility processes of conditionally normal GARCH model lies in the presence of the price of risk parameter, λ, in the risk-neutral specification. The ML estimate of λ is identified by the model's equity risk premium

$$\log \mathbb{E}_t^{\mathbb{P}}\left[e^{-r}\,\frac{S_{t+1}}{S_t}\right] = \log \mathbb{E}_t^{\mathbb{P}}\left[\exp\left\{\lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1} + \sqrt{h_{t+1}}\varepsilon_{t+1}\right\}\right] = \lambda\sqrt{h_{t+1}}. \tag{B.1}$$

Unfortunately, it is notoriously difficult to pin down the magnitude of the equity risk premium, and economists have not even reached a consensus about its very existence.[23] Hence, an interesting alternative that has been pursued by many authors is to estimate a forward-looking price of risk parameter λ from option data.[24] Following this path, we opt for a simple estimation criterion: minimizing, with respect to λ, the tracking error between the model's implied volatility and the VIX level, that is,

$$\min_{\lambda} \sum_{t\in\mathcal{T}} \left(\mathbb{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right] - \mathrm{VIX}_t\right)^2 . \tag{B.2}$$

A nonlinear least squares (NLS) estimation of λ is performed using the implied-volatility sum of squared errors criterion in Equation (B.2). This estimation procedure is most similar to that used by Gibson and Schwartz (1990) who use a mean squared error criterion on futures prices to estimate the market price of convenience yield risk in the NYMEX crude oil futures market. Similarly, Rosenberg and Engle (2002) extract empirical risk aversion levels on a monthly basis from options written on the S&P 500 between 1991 and 1995.

[23]The literature on the subject is overwhelmingly vast, starting with Mehra and Prescott (1985). Some authors, notably Brown, Goetzmann, and Ross (1995), argue that the equity risk premium could be solely due to a survival bias, a hypothesis undermined by others like, for instance, Li and Xu (2002). Pollard (2009) even attributes the premium to luck. DeLong and Magin (2009) offer a long review on the subject, concluding that the equity premium is still a puzzle.

[24]See, for instance, Chernov and Ghysels (2000), Pan (2002), Eraker (2007), and Ornthanalai (2009).

Our estimation procedure involves, for each candidate value of λ suggested by the optimizer, the computation of a Monte-Carlo integral for each date t entering the sum in Equation (B.2). To ease the computational burden inherent to this estimation procedure, only the VIX values that are observed on Wednesdays from June 1988 to December 2007 are used in the estimation.[25] Given the so-obtained NLS-optimal value of λ, we compare the model's implied volatility values with non-Wednesday observations of the VIX as an "out-of-sample" validation of the estimated value of λ. Finally, the same Monte-Carlo integration procedure is performed under the physical measure to obtain the MacroHV-MIDAS model's one-month-ahead, objective expectation $\mathbb{E}_t^{\mathbb{P}}\left[\sqrt{h_{t+21}}\right]$; the model-implied volatility premium, then, obtains by subtracting this objective expectation from the foregoing risk-neutral one. To provide a benchmark, the whole procedure is also applied to the NGARCH model.

In Table 7, Panel A summarizes VIX observations retained for our NLS estimation exercise, while Panel B reports the results obtained by minimizing Equation (B.2). Note that these latter results are reported in implied volatility root mean squared error (IVRMSE) terms, that is, using

$$\text{IVRMSE} = \sqrt{\frac{1}{N_{\mathcal{T}}}\sum_{t\in\mathcal{T}}\left(\mathbb{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right] - \text{VIX}_t\right)^2}. \tag{B.3}$$

The IVRMSE is strictly monotone in the objective function of Equation (B.2) but easier to interpret because it is on the same scale as the VIX. Using the λ values obtained from maximum likelihood on asset returns (Table 3), we compute benchmark IVRMSE values and report them along with the relative improvement achieved by using the NLS estimate of λ. The NLS estimates of λ for both the NGARCH (0.191) and MacroHV-MIDAS (0.199) model are more than twelve times higher than the estimates obtained under ML. Using the value of λ obtained under ML estimation, the IVRMSE of the NGARCH model's is 5.24% and the NLS estimate reduces this error to 2.92%, a 44.3% improvement. For the MacroHV-

[25]When no data are available on a given Wednesday, we use the next trading day's data. Note that this is the same twenty-year period that is analyzed in Section 4.2. However, the (new) VIX values are only available from January 1990. From June 1988 to December 1989, we therefore use VXO values (often referred to as the "old" VIX) that are based on OEX options rather than SPX ones and that are *not* model free. While the VXO is most likely a biased proxy of the value that the VIX would have taken over this early period, we are confident that this bias has little impact on the estimation of a single parameter using twenty years of weekly data.

MIDAS model, the benchmark IVRMSE is lower at 5.09% and yet the improvement is greater at 47.6%.[26]

As seen in Figure 4, using the values of λ obtained under ML on asset returns, both models' implied volatility errors are consistently negative through time. Once the price of risk parameter is estimated using VIX data, the magnitude of the NGARCH bias falls by 94.5% to -23 basis points (bps), while that of the MacroHV-MIDAS model falls even further at -6 bps, a 98.5% improvement. It thus appears that the time-series properties of the MacroHV-MIDAS model's volatility process are closer to those of the VIX than are those of the NGARCH model.

As we use only Wednesday observations of the VIX in our NLS estimation of λ, a legitimate concern is whether the time-series properties of the MacroHV-MIDAS model's implied volatility process are consistent with those of the VIX on a daily basis.By way of validation, we also compute the IVRMSEs and implied volatility biases of both models on the non-Wednesday observations left aside for the estimation. Out-of-sample improvements, both in terms of IVRMSE and IV biases, are very close to the ones obtained in sample. So, we can be confident that the MacroHV-MIDAS model's implied volatility process, $\mathbb{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right]$, is close to bias-free throughout the twenty years of data we consider.

[26]Just like Gibson and Schwartz (1990) in their footnote 17, we acknowledge that NLS relies on the assumption that our implied volatility errors have a normal, independent, and identical distribution. This assumption is most likely violated, as can be seen from the residuals in the middle panel of Figure 7. However, just as in Gibson and Schwartz (1990), our focus is not on the statistical significance of the parameter, and this is why we forgo developing a more involved procedure. In our subsequent analysis of the time-series properties of the extracted volatility premium process, in order to account for the strong autocorrelation of the residuals, we use Newey-West standard errors with a conservative lag of 63, corresponding to one quarter of trading days.

Figure 1: *Macroeconomic Series*

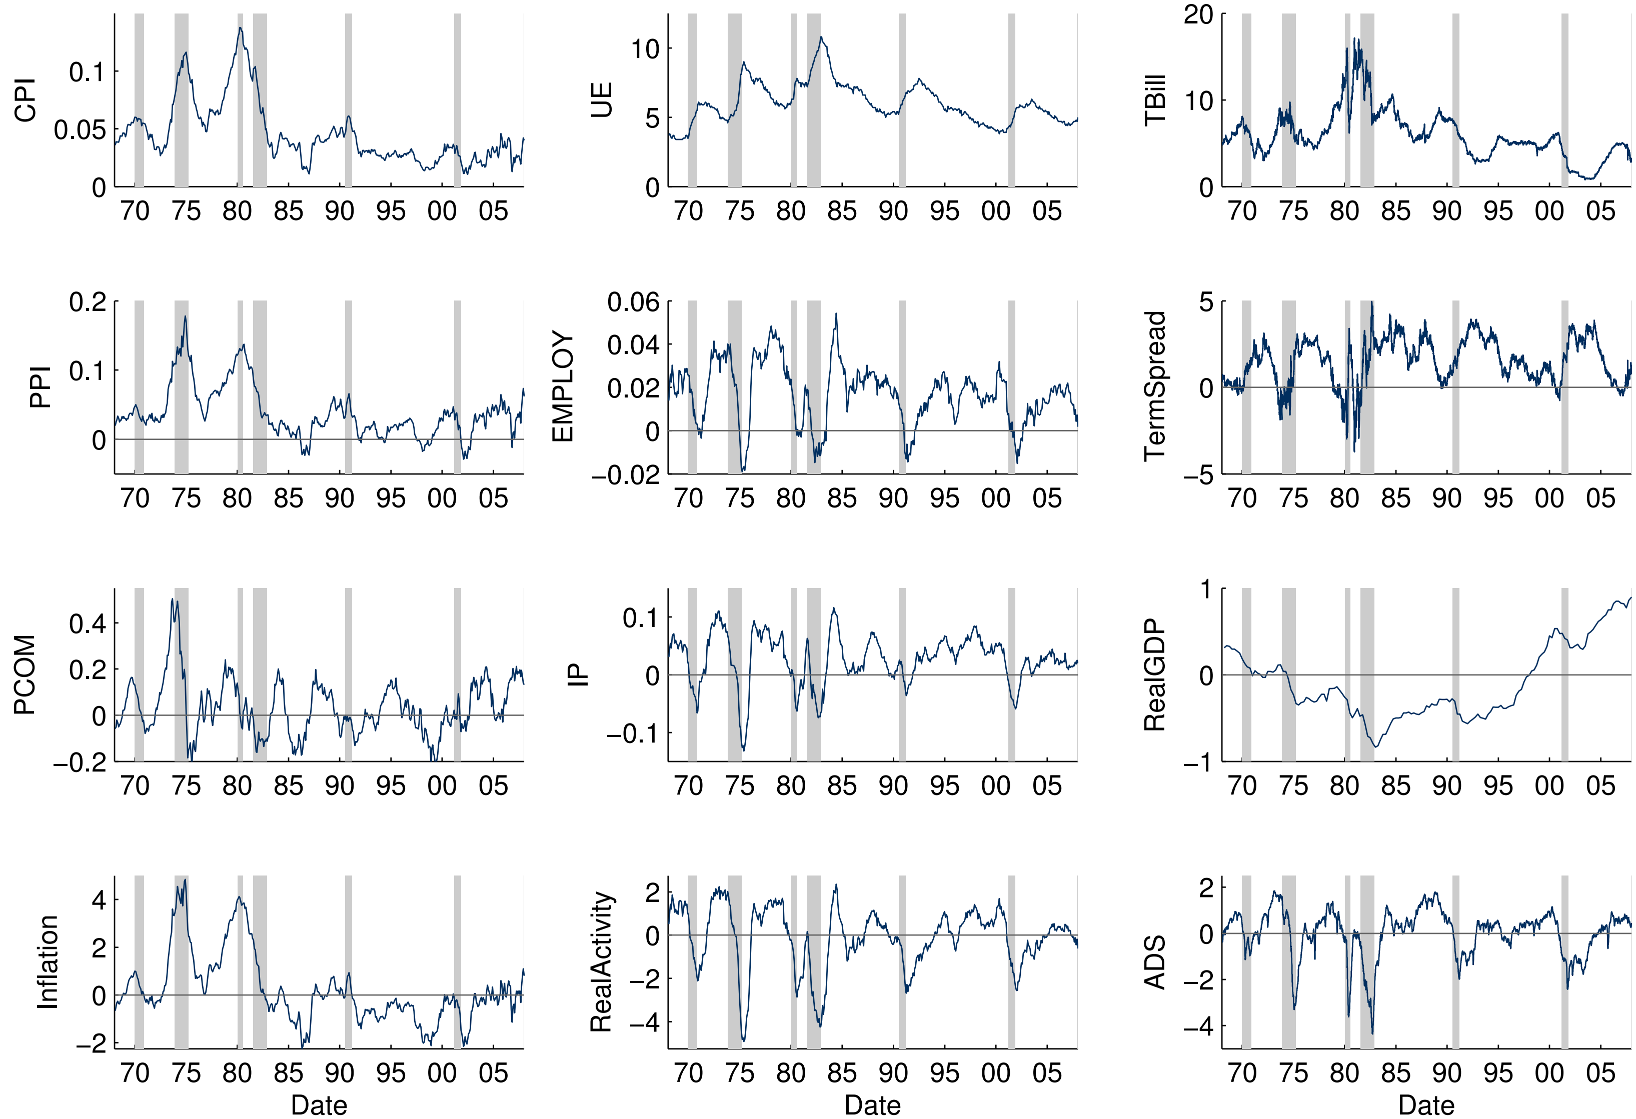

In the bottom row of the left and central columns are Ang & Piazzesi's monthly inflation and real activity indicators; their respective input series are reported in the first three rows. The right column displays the T-Bill rates, the term spread, the detrended real GDP, and the business conditions index.

Figure 2: *Fundamental Variance Processes*

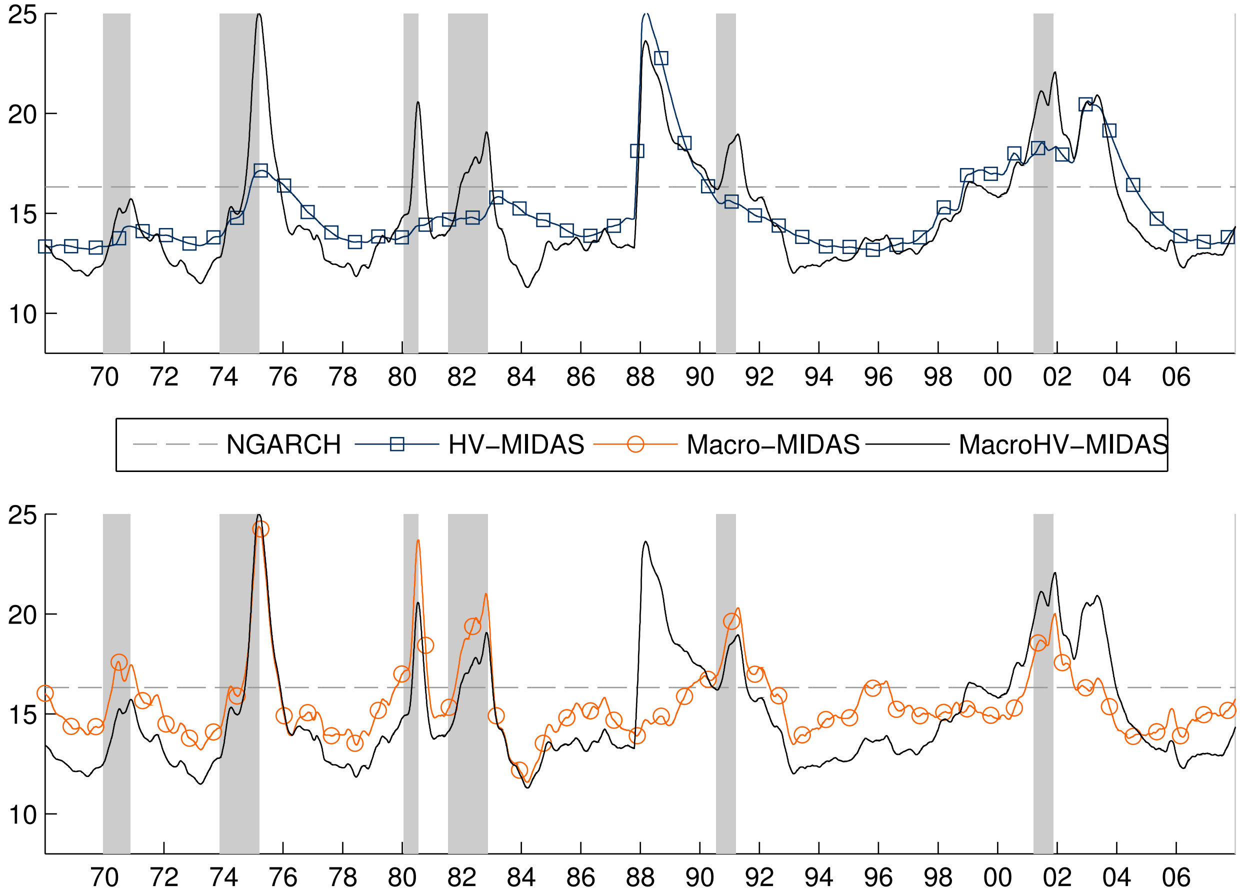

In both panels, we plot, as a solid black line, the annualized fundamental volatility level of the MacroHV-MIDAS model along with a dashed horizontal line at 16.31%, which corresponds to the NGARCH model's expected volatility, $\sqrt{252\omega/(1-\alpha(1+\gamma^2)-\beta)}$. In the upper panel, we superimpose the fundamental volatility level obtained for the HV-MIDAS model; in the lower panel, we superimpose the volatility level obtained for the Macro-MIDAS model.

Figure 3: *Option Valuation: A Time-Series Perspective*

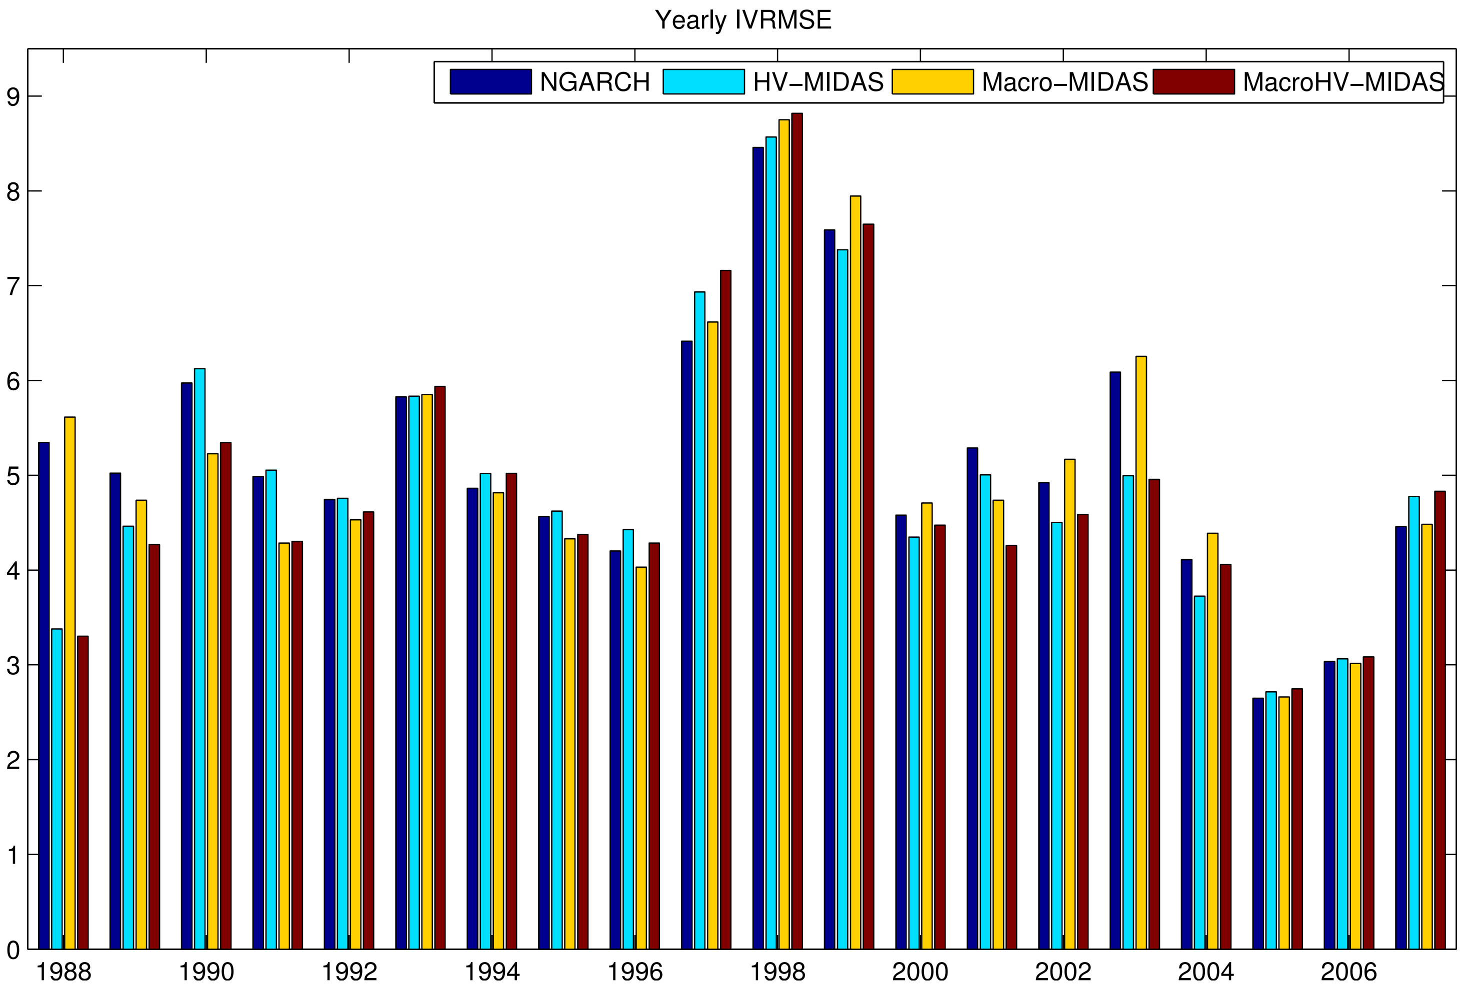

This figure reports yearly IVRMSE measures for the NGARCH model along with those for the MacroHV-MIDAS model and the nested HV-MIDAS and Macro-MIDAS models.

Figure 4: *MacroHV-MIDAS Model's IVRMSE Improvement Over the NGARCH Model*

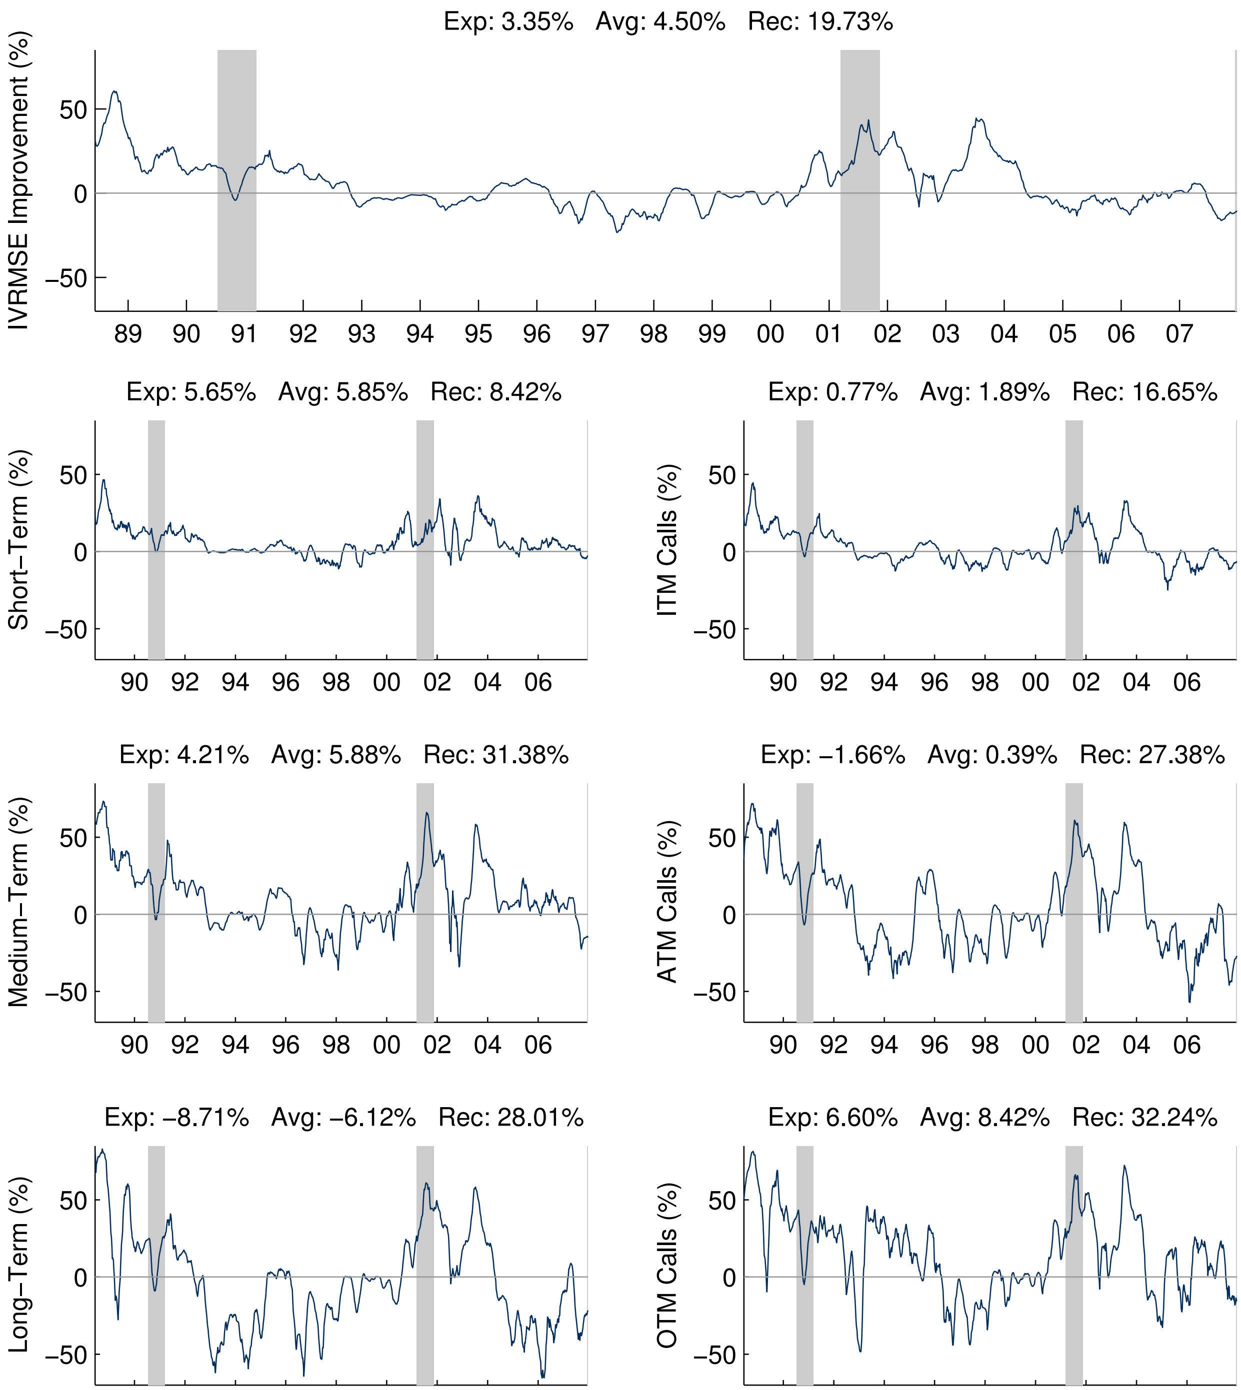

Using parameter estimates in Table 3, we price options and compute weekly IVRMSE measures for the NGARCH model, $IVRMSE_w^{\text{NG}}$, and for the MacroHV-MIDAS model, $IVRMSE_w^{\text{HV},\Delta x}$. This figure reports the relative improvement of using the latter model over the former, that is, $(IVRMSE_w^{\text{NG}} - IVRMSE_w^{\text{HV},\Delta x})/IVRMSE_w^{\text{NG}}$. On the left-hand side, results are divided along the options' maturity: 45 days to maturity (DTM) or less, between 46 and 90 DTM, or more than 90 DTM. On the right-hand side, results are divided along options' moneyness: $K/S \leq 0.975$, $0.975 < K/S < 1.025$, and $K/S \geq 1.025$. Above each subplot, we report the overall average improvement (Avg), as well as the average through expansion and recession periods (Exp/Rec).

Figure 5: *Implied-Volatility-Forecasting Bias of the MacroHV-MIDAS Model*

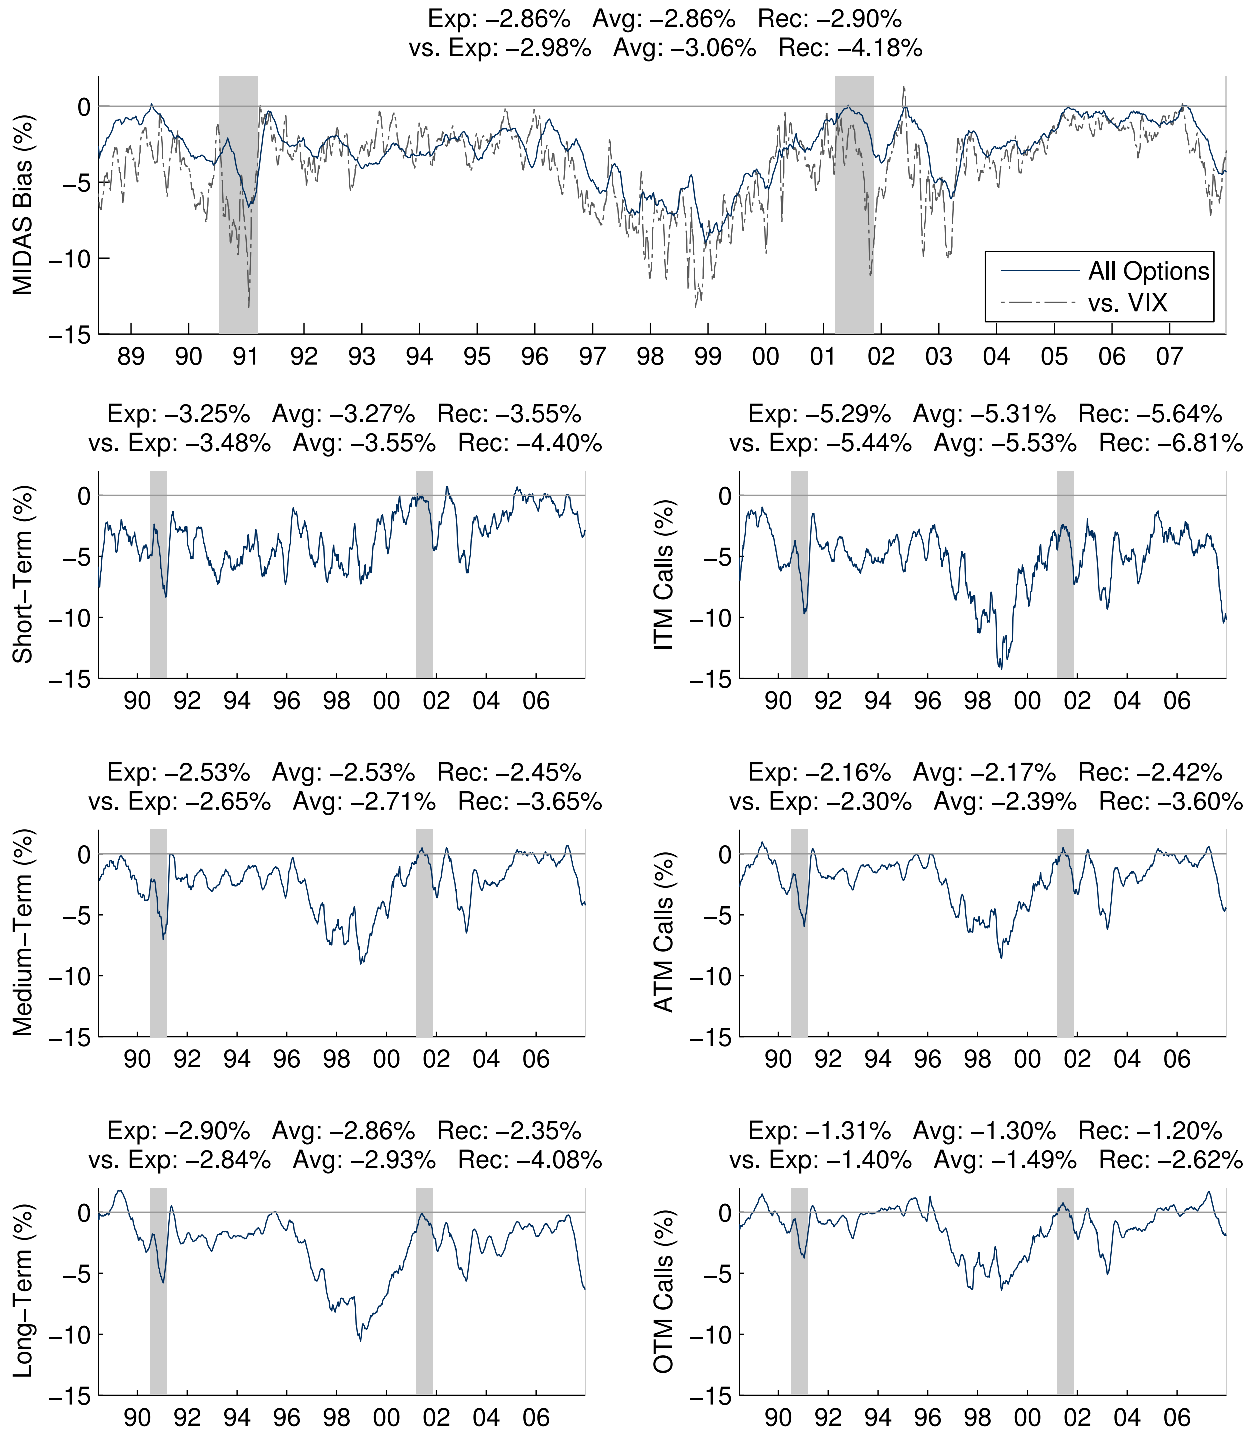

This figure reports the time series of implied-volatility biases for the MacroHV-MIDAS model. On the left-hand side, results are divided along the options' maturity: 45 days to maturity (DTM) or less, between 46 and 90 DTM, or more than 90 DTM. On the right-hand side, results are divided along options' moneyness: $K/S \leq 0.975$, $0.975 < K/S < 1.025$, and $K/S \geq 1.025$. Above each subplot, we report the overall average bias (Avg), as well as the average through expansion and recession periods (Exp/Rec); by way of comparison, the same averages are reported for the NGARCH model (vs).

Figure 6: *Synthetic Options*

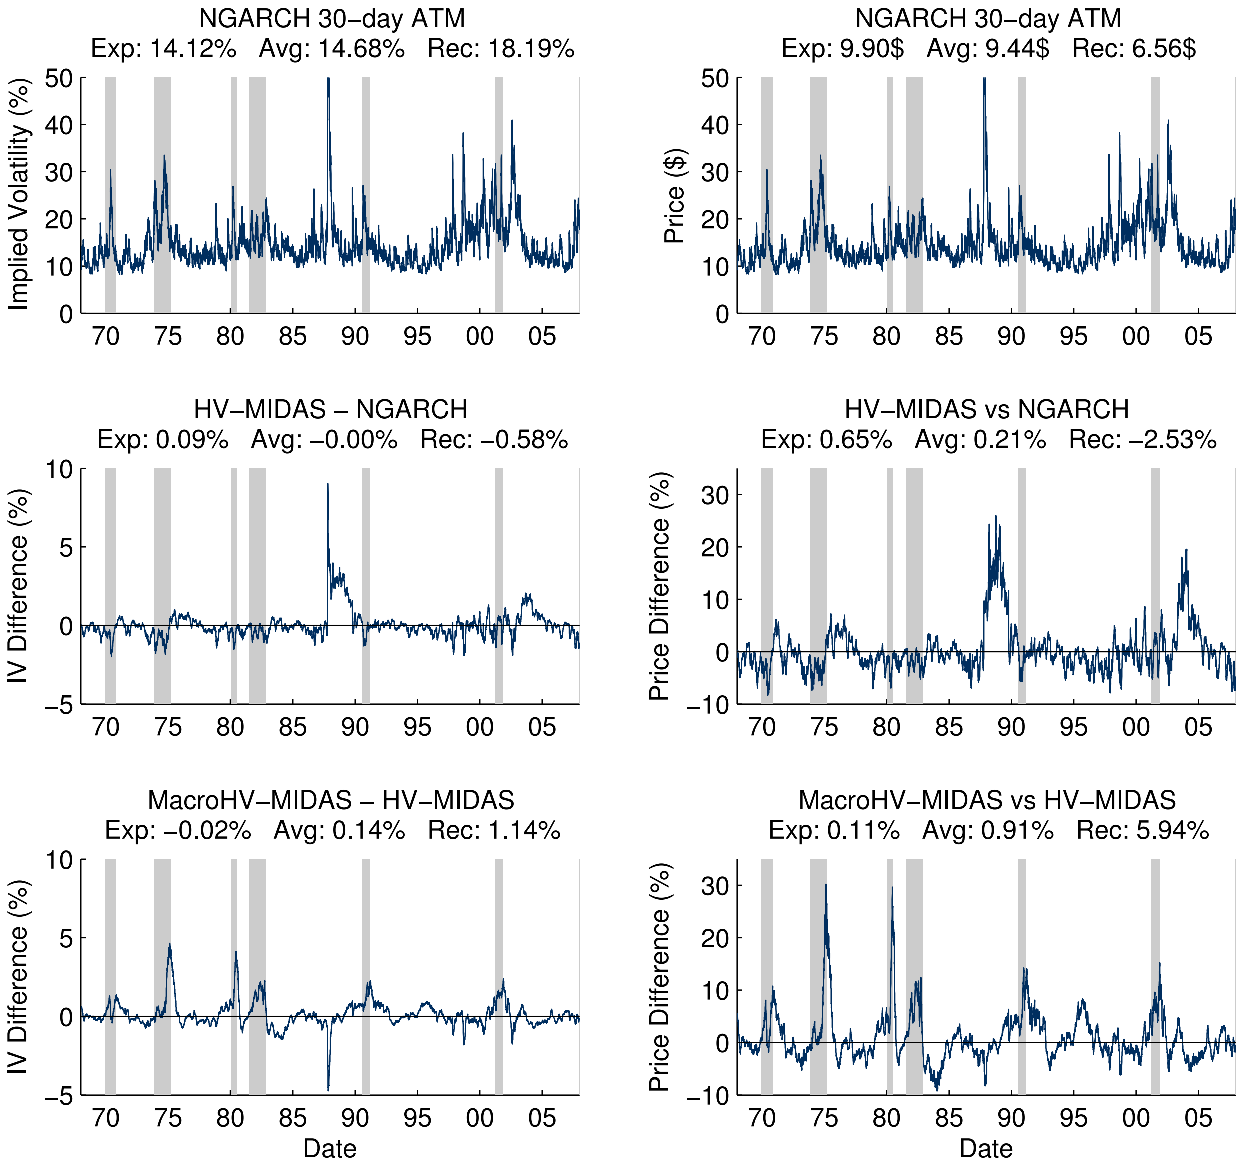

The upper-left panel of this figure reports the time series of implied volatilities obtained using the NGARCH model to price synthetic, at-the-money options with 30 days to maturity, daily from January 1968 to December 2007. These options are created assuming that their strike is equal to the index value on each given day. The mid-left panel reports the difference between the HV-MIDAS model's implied volatilities for these synthetic options and the NGARCH implied volatilities. Similarly, the lower-left panel reports difference in implied volatilities when comparing the MacroHV-MIDAS and HV-MIDAS models. The right-column panels report the data of the left column but in the price domain. Note that price differences are in relative terms, e.g. $\left(C_t^{\mathrm{HV}} - C_t^{\mathrm{NGARCH}}\right) \big/ C_t^{\mathrm{NGARCH}}$.

Figure 7: *The Volatility Premium*

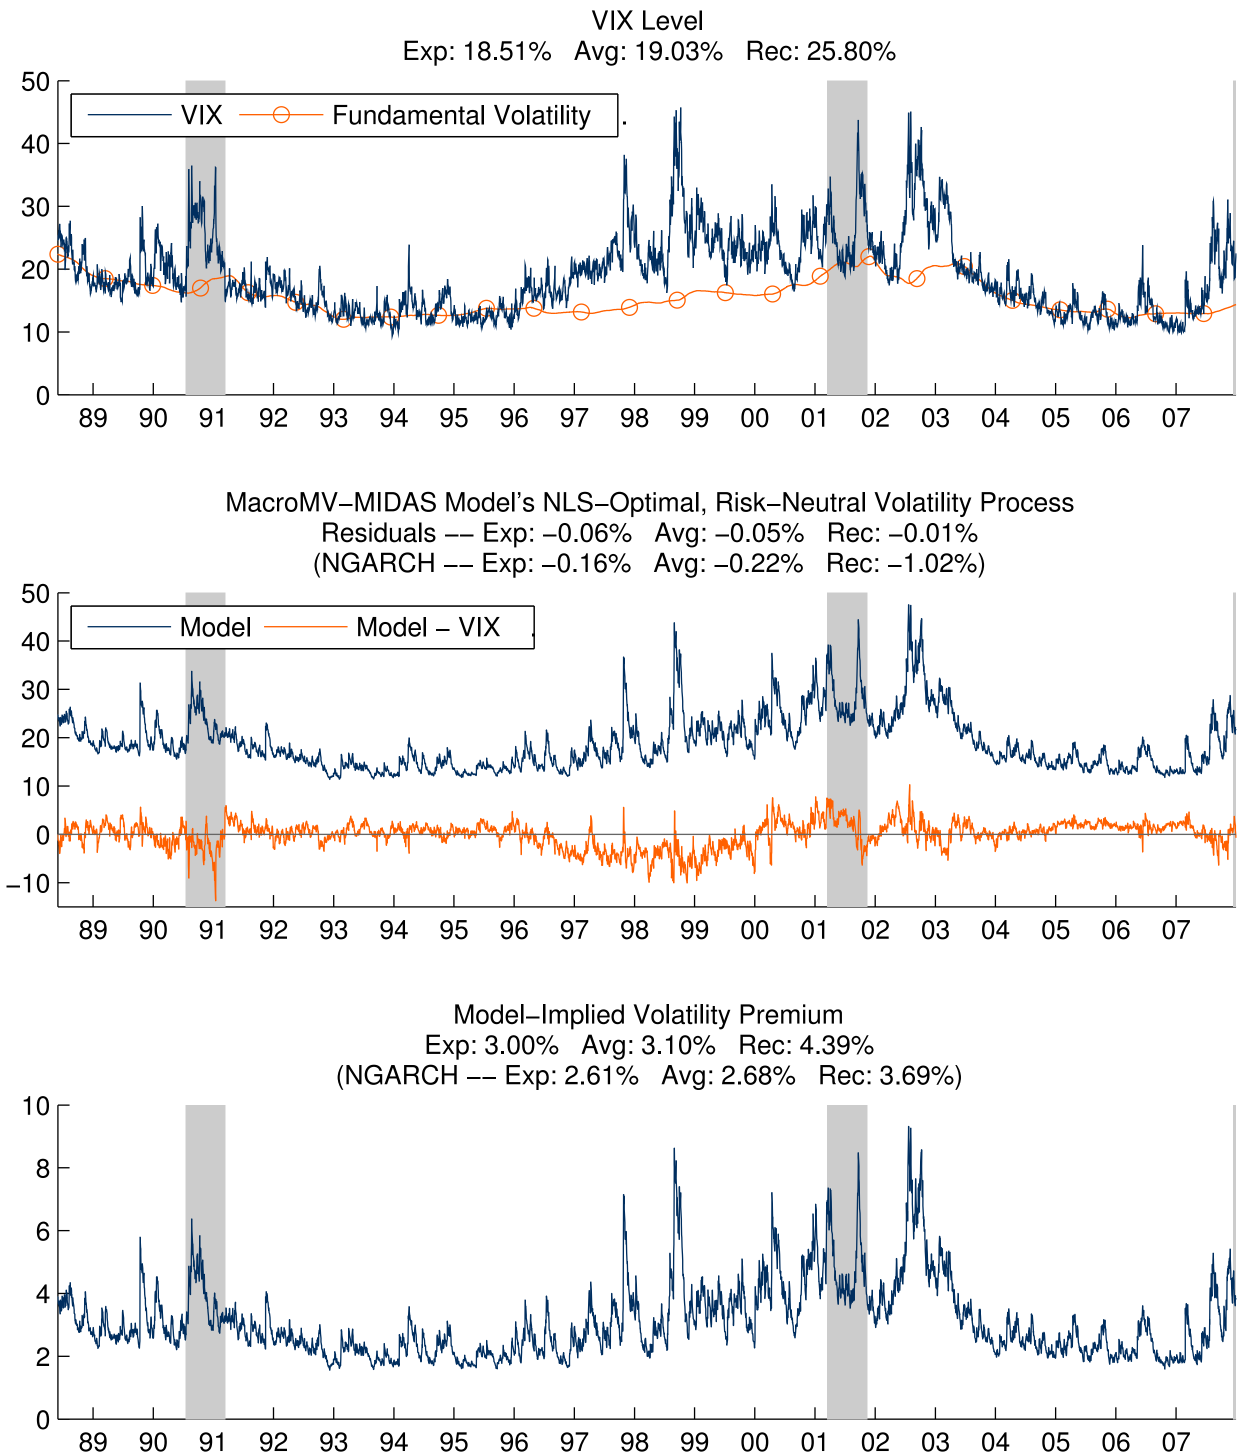

In the upper panel, we plot the VIX through time and report its overall average level as well as that through expansion and recession periods (Exp/Avg/Rec). For comparison sakes, we also plot the MacroHV-MIDAS model's fundamental volatility process; note that the latter is the long-run mean-reversion level for the physical volatility of the model, while the VIX is the expectation of one-month-ahead volatility under the risk-adjusted measure. The middle panel plots the model's NLS-optimal, risk-neutral volatility process and its difference with the VIX. The lower panel reports the volatility premium obtained by substracting the model's expected, one-month-ahead volatility process under the physical measures from the foregoing risk-neutral volatility process.

Table 1: *Descriptive Statistics on the Business Conditions Index*

			GARCH	
	Index: x_t	**Residuals: v_t**	**Residuals: u_t**	**Variances: $h_t \times 1e4$**
Min	-4.36	-0.36	-9.16	0.13
Mean	0.02	0.00	0.02	4.91
Max	1.83	0.21	6.52	171.50
Std. Dev.	1.04	0.02	1.06	9.31
Skewness	-1.21	-0.92	-0.26	6.95
Kurtosis	4.84	23.51	7.88	76.91
Log-Likelihood		-9254.3	27298.5	
Jarque-Bera		177850.0	10099.2	

This table reports summary statistics on the processes of Equation (3.1), (3.2) *and* (3.3) *:*

$$x_t = \varphi x_{t-1} + v_t \,, \qquad v_t = \sqrt{h_t^x} u_t \,, \qquad h_t^x = \omega_x + \alpha_x v_{t-1}^2 + \beta_x h_{t-1}^x \,.$$

Maximum likelihood parameter estimates are $\varphi = 0.999999$, $\omega_x = 1.4117e-06$, $\alpha_x = 0.1183$ *and* $\beta_x = 0.8817$, *implying a variance persistence of* 0.9999.

Table 2: *Correlations between Macroeconomic Series*

		ADS	**TBill**	**TS**	**UE**	**EMPLOY**	**Inflation**	**RA**
Monthly	**TBill**	-0.026*						
	Term Spread	-0.276	-0.436					
	UE	-0.475	0.304	0.532				
	EMPLOY	0.724	0.142	-0.204	-0.282			
	Inflation	-0.073*	0.556	-0.456	$0.090^{4.8\%}$	0.141		
	Real Activity	0.794	-0.059*	-0.286	-0.585	0.903	-0.023	
Quarterly	**Real GDP**	$0.170^{3.1\%}$	-0.502	-0.350	-0.733	0.012*	-0.074*	0.264

This table reports the correlations between eight of the series displayed in Figure 1. For the first six rows, the three daily series are sampled monthly; daily correlations between these are similar to the ones reported here. The last row reports the correlations of the first seven series, sampled quarterly, with the detrended real GDP; again, unreported quarterly correlations are very similar to the monthly ones. Correlations marked with an asterisk are not *statistically significant at the 90% level; the exponent, whenever there is one, is the correlation's p-value; all other correlations are significant at least to the 99% level.*

Table 3: *Maximum Likelihood Estimates*

		MIDAS		
	NGARCH	**HV**	**Macro**	**MacroHV**
$\boldsymbol{\lambda}$	0.0151	0.0156	0.0149	0.0155
	(2.31E-08)	(2.48E-08)	(2.23E-09)	(1.24E-08)
$\boldsymbol{\omega}$, **m**	1.03E-06	-9.714	-9.255	-9.730
	(1.30E-12)	(9.78E-07)	(3.26E-07)	(2.17E-07)
$\boldsymbol{\alpha}$	0.0567	0.0624	0.0589	0.0629
	(1.29E-09)	(4.09E-08)	(3.66E-09)	(8.30E-09)
$\boldsymbol{\beta}$	0.9067	0.8846	0.8951	0.8725
	(4.77E-08)	(1.72E-08)	(2.16E-08)	(5.23E-08)
$\boldsymbol{\gamma}$	0.6873	0.7283	0.7278	0.7850
	(4.33E-07)	(1.10E-06)	(2.29E-07)	(4.44E-07)
$\boldsymbol{\theta}_{\mathbf{hv}}$		68.234		63.825
		(5.27E-05)		(3.21E-06)
$\mathbf{w}_{\mathbf{hv}}$		2.931		3.203
		(2.68E-06)		(1.38E-06)
$\boldsymbol{\theta}_{\mathbf{m}}$			-1.238	-1.009
			(1.54E-06)	(5.53E-07)
$\mathbf{w}_{\mathbf{m}}$			3.370	3.722
			(7.31E-07)	(1.82E-06)
SR Persistence	0.9902	0.9801	0.9852	0.9741
1.0 − LR ($\times 10^4$)		0.6616	0.7376	0.7036
SR VoV ($\times 10^4$)	1.576	1.786	1.677	1.862
LR VoV ($\times 10^4$)		0.0338	0.0483	0.0533
Corr($\mathbf{R}_{t+1}, \mathbf{h}_{t+2}$)	-69.70%	-71.75%	-71.72%	-74.30%
Corr($\varepsilon_{t+1}, \mathbf{u}_{t+1}$)			5.04%	5.03%
Log-Likelihood	33791.3	33806.7	33803.0	33822.0
BIC	-6.7080	-6.7093	-6.7085	-6.7105

This table reports maximum likelihood estimates for the NGARCH model as well as those of MIDAS models for three different specifications: (i) HV-MIDAS: quarterly historical volatilities computed from daily returns; (ii) Macro-MIDAS: quarterly differences of the Business Conditions Index values; and (iii) MacroHV-MIDAS: a combination of both (i) and (ii). Below each parameter estimate, we report its Bollerslev-Wooldridge standard error. The Bayesian information criterion (BIC) values account for the number of parameters in each model and for the length of the time series of S&P 500 returns between January 1968 and December 2007.

Short-run (SR) persistence and annualized volatility of variance (VoV) are $\alpha(1+\gamma^2)+\beta$ and the average of $\sqrt{252\,\alpha^2(2+4\gamma^2)h_{t+1}^2}$, respectively. For the long-run (LR) component, we approximate the persistence and volatility of variance by fitting an AR(1) to the fundamental volatility process, i.e.,

$$\tau_t = \phi_0 + \phi_1 \tau_{t-1} + \sqrt{v} e_t \,,$$

where e_t is white noise. The volatility of variance is approximated by $\sqrt{252v}$, while the long-run persistence is approximated by ϕ_1 and is very close to one for all MIDAS models; we here report $10^4 \times (1-\phi_1)$.

Table 4: *Option-Pricing Results*

	Moneyness	N	Avg.	NGARCH	MIDAS HV	MIDAS Macro	MIDAS MacroHV
Log-Likehood				33791.3	33806.7	33803	**33822**
IVRMSE		68923	18.41	5.23	5.14	5.21	**5.10**
DTM ≤ 45	**(0.33, 0.95)**	5478	29.83	12.61	12.52	12.49	**12.39**
	[0.95,1.00)	7464	17.81	4.29	4.19	4.18	**4.08**
	[1.00, 1.05)	8431	15.03	2.67	2.55	2.62	**2.47**
	[1.05, 1.87]	3147	18.28	3.81	3.57	3.72	**3.41**
45 < DTM ≤ 91	**(0.33, 0.95)**	3295	23.02	7.08	6.98	6.82	**6.77**
	[0.95,1.00)	3854	17.27	3.84	3.79	3.72	**3.67**
	[1.00, 1.05)	5282	15.13	3.01	2.96	3.02	**2.93**
	[1.05, 1.87]	4112	16.01	2.96	2.72	2.98	**2.72**
91 < DTM ≤ 182	**(0.33, 0.95)**	3221	21.54	5.36	5.43	**5.23**	5.34
	[0.95,1.00)	2519	17.88	3.68	3.72	**3.61**	3.63
	[1.00, 1.05)	2924	16.49	3.36	3.32	3.40	**3.31**
	[1.05, 1.87]	4018	16.02	3.37	3.11	3.43	**3.07**
DTM > 182	**(0.33, 0.95)**	2950	20.90	**4.95**	5.21	5.21	5.41
	[0.95,1.00)	2572	18.40	**4.06**	4.11	4.30	4.27
	[1.00, 1.05)	3105	17.89	4.08	**4.00**	4.35	4.18
	[1.05, 1.87]	6551	16.60	3.73	**3.31**	3.96	3.43
RMSE		68923	$40.44	39.34	32.35	38.23	**30.55**

For each model, we first recall its log-likehood as estimated in Section 3.3. Then, we report overall IVRMSE values, followed by IVRMSE values obtained over maturity/moneyness buckets of options. For completeness, we also report the overall RMSE values; IVRMSEs and RMSEs are computed as follows:

$$IVRMSE = \sqrt{\frac{1}{N}\sum_{t,k}\left(\mathrm{IV}\left(C_{t,k}\right) - \mathrm{IV}\left(C_{t,k}^{\mathrm{MODEL}}\right)\right)^2} \quad \textit{and} \quad RMSE = \sqrt{\frac{1}{N}\sum_{t,k}\left(\frac{C_{t,k} - C_{t,k}^{\mathrm{MODEL}}}{C_{t,k}}\right)^2}.$$

Apart from the average call price (40.44£), all entries are percentage points. In each row, the entry for the best performing model for that row is in bold font.

Table 5: *The MacroHV-MIDAS Model and Business Conditions*

Panel A: MacroHV-MIDAS IVRMSE Improvement Over the NGARCH Model

	Expansion	Overall	Recession	$x_t^{(63)} < 0$	$x_t^{(63)} < -1$	$x_t^{(63)} < -1.5$
# Weeks	948	1020	72	380	178	45
All Calls	3.35	4.50	19.73	12.88	22.18	29.75
	(0.69)	(0.68)	(2.58)	(1.02)	(1.60)	(2.64)
Short-Term	5.65	5.85	8.42	9.84	10.63	23.45
	(1.28)	(1.24)	(4.93)	(1.81)	(3.61)	(3.45)
Medium-Term	4.21	5.88	31.38	16.22	21.35	39.06
	(1.17)	(1.14)	(3.79)	(1.88)	(3.68)	(3.20)
Long-Term	-8.71	-6.12	28.01	11.95	30.07	42.92
	(1.27)	(1.25)	(3.61)	(1.60)	(1.87)	(3.09)
ITM Calls	0.77	1.89	16.65	9.85	16.65	24.21
	(0.67)	(0.66)	(2.46)	(0.95)	(1.64)	(2.68)
ATM Calls	-1.66	0.39	27.38	16.86	29.45	42.71
	(1.28)	(1.23)	(3.33)	(1.57)	(2.01)	(2.54)
OTM Calls	6.60	8.42	32.24	15.24	34.23	45.16
	(2.26)	(2.12)	(3.51)	(4.80)	(2.43)	(3.58)

Panel B: Synthetic 30-DTM, ATM Options

	Expansion	Overall	Recession	$x_t^{(63)} < 0$	$x_t^{(63)} < -1$	$x_t^{(63)} < -1.5$
# Days	8684	10068	1384	3833	1640	819
NGARCH IV	14.12	14.68	18.19	15.57	17.86	17.01
	(0.06)	(0.06)	(0.13)	(0.08)	(0.12)	(0.13)
IV Difference	0.06	0.13	0.56	0.57	1.02	1.60
	(0.01)	(0.01)	(0.04)	(0.02)	(0.04)	(0.05)
NGARCH Price ($)	9.90	9.44	6.56	8.88	11.41	6.73
	(0.11)	(0.10)	(0.25)	(0.16)	(0.27)	(0.29)
Relative Price	0.77	1.12	3.31	4.00	6.26	9.05
Difference (%)	(0.07)	(0.07)	(0.19)	(0.11)	(0.20)	(0.30)

In Panel A, using parameter estimates in Table 3, we price options and compute weekly IVRMSE measures for the NGARCH model, $IVRMSE_w^{\mathrm{NG}}$*, and for the MacroHV-MIDAS model,* $IVRMSE_w^{\mathrm{HV},\Delta x}$*. Panel A reports averages of the relative improvement resulting from using the latter model over the former, that is,* $(IVRMSE_w^{\mathrm{NG}} - IVRMSE_w^{\mathrm{HV},\Delta x})/IVRMSE_w^{\mathrm{NG}}$ *. The average over all 1020 weeks in the option data set is reported under the* Overall *column. The first and third columns report averages when restricting to weeks falling in periods of expansion or recession. The last three columns report the average when restricting to Wednesdays for which the centered quarterly moving average of the ADS Index* $\left(x_t^{(63)} = \frac{1}{63}\sum_{s=t-31}^{t+31} x_s\right)$ *is below 0, -1 or -1.5. All entries are in percentage points, and standard errors are parenthesized below each average. Similarly, Panel B reports the average NGARCH implied volatility of a synthetic 30-DTM, at-the-money call (obtained by setting the call's strike at the index level on each day in our sample), from January 1968 to December 2007 (the time series is reported in the upper-left panel of Figure 6). The difference between the MacroHV-MIDAS model's implied volatility and that of the NGARCH is then reported. For comparison, the NGARCH average price is reported for the different subsamples (the only entries in dollar terms), along with the MacroHV-MIDAS – NGARCH price difference in relative terms, i.e. based on* $\left(C_t^{\mathrm{HV},\Delta x} - C_t^{\mathrm{NG}}\right)\Big/C_t^{\mathrm{NG}}$ *.*

Table 6: *The MacroHV-MIDAS Model and the Marginal Impact of Business Conditions*

Panel A: MacroHV-MIDAS IVRMSE Improvement Over the HV-MIDAS Model

	$x_t^{(63)} > 1.5$	$x_t^{(63)} > 1$	$x_t^{(63)} > 0$	$x_t^{(63)} < 0$	$x_t^{(63)} < -1$	$x_t^{(63)} < -1.5$
# Weeks	22	52	640	380	178	45
All Calls	7.24	3.70	-1.16	7.09	10.94	19.82
	(2.87)	(1.29)	(0.34)	(0.84)	(1.27)	(2.27)
Short-Term	8.00	3.98	2.08	5.43	5.45	18.10
	(7.82)	(3.32)	(0.45)	(1.48)	(2.84)	(2.61)
Medium-Term	1.85	-0.96	-0.44	9.20	12.63	31.06
	(3.13)	(1.93)	(0.69)	(1.67)	(2.95)	(3.39)
Long-Term	20.92	9.36	-8.54	9.51	14.99	27.17
	(3.16)	(2.14)	(0.74)	(1.29)	(1.83)	(3.40)
ITM Calls	7.13	5.06	-1.45	6.67	10.05	17.97
	(2.79)	(1.29)	(0.36)	(0.82)	(1.37)	(2.60)
ATM Calls	6.04	0.78	-4.92	14.80	18.42	33.51
	(3.35)	(1.78)	(0.68)	(1.70)	(1.89)	(3.07)
OTM Calls	14.50	4.56	-3.33	18.10	16.76	28.50
	(4.34)	(2.29)	(0.83)	(7.67)	(2.10)	(3.45)

Panel B: Synthetic 30-DTM, ATM Options

	$x_t^{(63)} > 1.5$	$x_t^{(63)} > 1$	$x_t^{(63)} > 0$	$x_t^{(63)} < 0$	$x_t^{(63)} < -1$	$x_t^{(63)} < -1.5$
# Days	375	1131	6235	3833	1640	819
NGARCH IV	17.61	16.54	14.14	15.57	17.86	17.01
	(0.76)	(0.30)	(0.07)	(0.08)	(0.12)	(0.13)
IV Difference	-0.57	-0.50	-0.16	0.62	1.13	1.83
	(0.05)	(0.02)	(0.01)	(0.02)	(0.03)	(0.04)
NGARCH Price ($)	4.41	4.46	9.79	8.88	11.41	6.73
	(0.24)	(0.10)	(0.13)	(0.16)	(0.27)	(0.29)
Relative Price	-2.31	-2.18	-0.78	3.66	6.34	9.87
Difference (%)	(0.08)	(0.07)	(0.03)	(0.09)	(0.17)	(0.25)

This tables adds to the results reported in Table 5. Panel A reports averages of the relative improvement resulting from using the MacroHV-MIDAS model over the HV-MIDAS, that is, $(IVRMSE_w^{\mathrm{HV}} - IVRMSE_w^{\mathrm{HV},\Delta x})/IVRMSE_w^{\mathrm{NG}}$ *; the IVRMSE of the NGARCH is kept at the denominator to ease comparison with Table 5. The columns report the average when restricting to Wednesdays for which* $x_t^{(63)}$ *is above 1.5, 1 or 0, or below 0, -1 or -1.5. Similarly, Panel B reports, for a synthetic 30-DTM, at-the-money call, the difference between the MacroHV-MIDAS model's implied volatility and that of the HV-MIDAS, that is* $\left(\mathrm{IV}_t^{\mathrm{HV},\Delta x} - \mathrm{IV}_t^{\mathrm{HV}}\right)$. *The comparison is also performed in the price domain, i.e. based on* $\left(C_t^{\mathrm{HV},\Delta x} - C_t^{\mathrm{HV}}\right)/C_t^{\mathrm{NG}}$.

Table 7: *Model-Implied Volatility Premium*

Panel A: VIX Observations

	#	Average	Min	Median	Max
Daily	4936	19.03%	9.31%	17.90%	45.74%
Wednesdays	1020	19.00%	9.31%	17.95%	43.51%
Out-of-Sample	3916	19.03%	9.48%	17.89%	45.74%

Panel B: NLS Estimation from VIX Observations

	NGARCH		**MacroHV-MIDAS**	
λ	0.191		0.199	
	Wednesday (%)	**Out of Sample (%)**	**Wednesday (%)**	**Out of Sample (%)**
IVRMSE				
NLS	2.92	2.90	2.66	2.67
Benchmark	5.24	5.24	5.09	5.11
Improvement	44.34	44.59	47.63	47.71
IV Bias				
NLS	-0.23	-0.22	-0.06	-0.05
Benchmark	-4.26	-4.25	-4.05	-4.05
Improvement	94.50	94.81	98.45	98.72

Panel C: MacroHV-MIDAS Model-Implied Volatility Premium

		$\sqrt{\tau_t(HV)}$	$\sqrt{\tau_t(\Delta x)}$	$\sqrt{\tau_t}$					
cst	310.0	310.0	310.0	310.0	310.0	310.0	310.0	310.0	310.0
	(27.86)	**(32.05)**	**(28.91)**	**(34.82)**	**(35.08)**	**(52.44)**	**(57.08)**	**(197.63)**	**(201.55)**
$\sqrt{\tau_t}$		26.3	45.1	24.5	22.7			18.9	18.4
		(5.06)	**(4.56)**	**(7.56)**	**(6.33)**			**(28.58)**	**(26.93)**
$\sqrt{g_t}$						105.9	102.9	97.7	97.4
						(17.50)	**(16.99)**	**(32.16)**	**(31.05)**
NBER	138.5				54.3		79.4		14.3
	(4.59)				(1.63)		**(4.48)**		**(2.01)**
$\mathbf{R^2}$	8.9%	24.2%	12.9%	31.8%	33.0%	79.3%	82.2%	97.7%	97.8%

Panel A summarizes VIX observations between June 1988 and December 2007; for our NLS exercise, we retain one observation a week, on Wednesdays. Panel B reports the results obtained by minimizing, with respect to the price of risk parameter λ, squared differences between the VIX level and the expected one-month-ahead volatility implied by the NGARCH and MacroHV-MIDAS models under the risk-neutral measure. The benchmark measures are those obtained when using the λ value obtained by ML on asset returns (Table 3) and improvement measures are obtained by comparing the magnitude of IVRMSEs and biases to these benchmarks. The minimization is performed using only Wednesday observations; all measures of fit are also reported for non-Wednesday observations (out of sample) by way of validation. Panel C reports the results of nine linear regressions with, as regressand, the model-implied volatility premium of Figure 7, $\mathbb{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right] - \mathbb{E}_t^{\mathbb{P}}\left[\sqrt{h_{t+21}}\right]$, obtained by simulating daily physical and risk-adjusted volatility processes using the price of risk parameter reported in Panel B. The $\sqrt{\tau_t}$ regressor is the annualized fundamental volatility, in percentage points, under the physical measure; in the second and third columns, we restrict this process to the impact of historical volatilities ($\theta_m = 0$) and to that of changes in business conditions ($\theta_{hv} = 0$), respectively. The NBER regressor is a dummy variable that has value one when a day falls within an NBER recession, and zero otherwise. All regressors are demeaned, $\sqrt{g_t}$ is further standardized, and all loadings can be interpreted in basis points terms. The t-stats are computed using Newey-West standard errors with a lag of 63, corresponding to one quarter of trading days, and are bold whenever their magnitude is larger than 1.96.