DESAUTELS | McGill

Faculty of Management

Pricing Financial Derivatives: The Impact of Business Conditions and Systematic Risk

Christian Dorion
Desautels Faculty of Management
McGill University, Montreal

September 2010

A thesis submitted to McGill University in partial fulfilment of the requirements for the degree of Doctor of Philosophy.

September 9, 2010 • Final Version
May 7, 2010 • Original Version

McGill University
Desautels Faculty of Management

This thesis

**Pricing Financial Derivatives:
The Impact of Business Conditions and Systematic Risk**

was presented by

Christian Dorion

to a committee formed of

David RAGSDALE	Pro-dean
Saibal RAY,	Departmental chair
Peter CHRISTOFFERSEN,	Thesis co-supervisor
Kris JACOBS,	Thesis co-supervisor (absent)
Jeroen ROMBOUTS,	Internal Examiner
Jason WEI,	External Examiner
Lars STENTOFT,	External Member
Zvi SINGER,	External Member

Thesis submitted on May 7, 2010
Thesis defended on July 5, 2010
Final submission on September 9, 2010

Abstract

This thesis comprises of three essays on the pricing of financial derivatives. In the first essay, we assess the return fitting and option valuation performance of generalized autoregressive conditional heteroscedasticity (GARCH) models. We compare component versus GARCH(1,1) models, affine versus nonaffine GARCH models, and conditionally normal versus nonnormal GED models. We find that nonaffine models dominate affine models in terms of both fitting returns and option valuation. For the affine models, we find strong evidence in favor of the component structure for both returns and options; for the nonaffine models, the evidence is less convincing in option valuation. The evidence in favor of the nonnormal GED models is strong when fitting daily returns, but not when valuing options.

In the second essay, 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 in 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.

In the third essay, we study the impact of systematic risk on the pricing of two economically similar derivative contracts: credit default swaps and equity put options. We document, for roughly 130 firms that have been part of the CDX index between 2004 and 2007, that the greater proportion of a firm's volatility that is systematic, the more expensive it is to purchase insurance via both (i) put options and (ii) credit default swaps. We provide evidence that these two derivatives are influenced by systematic risk through the same channel.

Résumé

Cette thèse repose sur trois essais traitant de l'évaluation de produits dérivés financiers. Le premier essai porte sur la performance relative de différents modèles d'évaluation d'option à variance GARCH. Nous comparons ces modèles suivants trois axes : variance à deux composantes vs. GARCH(1,1), variance affine versus non affine, rendements conditionnellement Gaussiens ou non. Les différents modèles sont comparés sur la base de leur capacité à correctement expliquer la série des rendements (estimation) et à prédire le prix des options (prédiction). Les résultats obtenus favorisent largement les modèles à volatilité non affine sur ceux à volatilité affine, tant du point de vue de l'estimation que de la prédiction. Pour les modèles à volatilité affine, les modèle à composantes sont favorisée par les données, à l'estimation et à la prédiction ; pour les modèles à volatilité non affine, les résultats sont moins convaincants au point de vue de la prédiction. À l'estimation, la supériorité des modèles qui ne reposent pas sur une hypothèse de normalité conditionnelle est clairement démontrée, mais les résultats à la prédiction sont plus mitigés.

Dans le second essai, nous introduisons un modèle de volatilité dynamique suivant lequel la volatilité du marché varie autour d'un processus de volatilité fondamental dont le niveau est déterminé par une mesure de risque macroéconomique. Ce risque est quantifié à l'aide d'une structure de données échantillonnées à intervalles irréguliers (MIDAS) permettant de capturer les variations d'un nouvel indicateur macroéconomique, le *Aruoba-Diebold-Scotti (ADS) Business Conditions Index*. Le nouveau modèle performe mieux que le modèle de référence, tant en ce qui à trait à l'estimation qu'à la prédiction, tout particulièrement aux environs des récessions de 1990-1991 et de 2001. Peu importe la maturité ou le degré de parité des options, le modèle de référence présente une saisonnalité dans ses erreurs de prédictions ; elles s'accroisse quand les conditions macroéconomiques se détériorent. Le modèle ici introduit, en ajustant ses prédictions aux conditions macroéconomiques courantes, permet de corriger cette saisonnalité. L'effet du risque macroéconomique sur le modèle explique près de 13% de la variation la prime de volatilité implicite au modèle.

Dans le troisième essai, nous étudions l'influence du risque systématique sur deux types de contrats économiquement similaires : les swaps sur défaillance et les options sur actions. Nous documentons, pour un ensemble de 130 firmes ayant fait parti de l'indice CDX entre 2004 et 2007, qu'une plus grande part de risque systématique augmente le prime de risque encourue pour acheter autant les swaps que les options.

À mes parents, mes grands-parents, mon frère,
mon âme soeur et mes quelques petites soeurs

Acknowledgment

The finance faculty at Desautels is an extraordinary group. I wish to thank the professors who taught me so well, who provided guidance through the course of my doctoral degree, on academic or less academic issues. Thank you Benjamin Croitoru, Adolfo De Motta, Vihang Errunza, Peter Christoffersen, Kris Jacobs and Jan Ericsson.

I am especially indebted to Peter, Kris and Jan. Working with Jan has been and hopefully will continue to be an enriching experience. His support, especially through this crucial last year, has been dearly appreciated.

Peter and Kris probably form the best duo of thesis supervisors a Ph.D. student could hope for. Their invaluable support, insightful comments and knowledgeable pieces of advice were central to making my Ph.D. a success. Working with them undoubtedly made me a better researcher and gave me role models to look up to. I would be truly proud if I could be, to my future students, half as inspiring and resourceful as each of them were to me.

I also want to seize the oppotunity to thank my fellow students who contributed to making this experience so valuable to me, especially Hitesh Doshi, Aurelio Vasquez and Hai Ta, with whom I suffered the comps a few years ago, and the “elders”, Chayawat Ornthanalai and Redouane Elkamhi, for sharing their invaluable experience with us.

The secretarial staff has also been most kind and helpful, and I especially want to thank Susan Lovasik and Stella Scalia who’s precious help was sincerely appreciated.

Thanks to the *Fonds québécois de la recherche sur la nature et les technologies* and the *Institut de Finance Mathématique de Montréal* for the financial support that allowed me to fully devote to my studies.

Contributions of Authors

Chapter 2 is a joint collaboration between Peter Christoffersen, Christian Dorion, Kris Jacobs and Yintian Wang. Christoffersen and Jacobs are Associate Professors of finance at the Desautels Faculty of Management at McGill University. Wang is Assistant Professor at the School of Economics and Management at Tsinghua University. Christoffersen and Jacobs developed the main research ideas behind the chapter. Wang performed the data management and analysis that made up a first version of the chapter. Dorion performed the data management and analysis, and discussed with Christoffersen and Jacobs the substantial revisions that led to the current version of this paper, forthcoming in the Journal of Business and Economic Statistics.

Chapter 4 is a joint collaboration between Christian Dorion, Redouane Elkamhi and Jan Ericsson. Elkamhi is Assistant Professor of finance at the Henry B. Tippie College of Business at University of Iowa. Ericsson is Associate Professor of finance at the Desautels Faculty of Management at McGill University. All coauthors contributed equally to this paper.

Contents

List of Figures

List of Tables

1 Introduction

OVER THE LAST DECADE, financial derivative markets expanded at a roaring pace. According to the semiannual over-the-counter derivatives markets statistics from the Bank of International Settlements,* the notional value of all derivatives contract traded in the G10 grew from 72,134 billions of US dollars in 1998 to 604,622 billions in early 2009. The gross market value of these derivative contracts grew from 2,580 billions to 25,372 billions over the same period. While Black and Scholes' (1973) seminal paper contributed to the structuring and early growth of option markets, the growth of derivatives markets observed over the last decade is largely due to the introduction of new derivative contracts. These new products are often increasingly complex, and so are the models required to evaluate these products.

Yet, even the most sophisticated models rely on a variety of assumptions, of modeling choices made by the econometrician. When using these models for financial decision making, it is important to understand the implications of these assumptions on the economic properties of the models available to evaluate a given derivative contract. Proper understanding of the weaknesses of the models at hand is key to an investor's ability to manage the *model risk* inherent to the fact that his decision making process is influenced by the use of a given model or array of models. The LTCM debacle, in September 1998, is an obvious example of the possible consequences of underestimating the importance of model risk. Some will argue that erroneous models also played a part that the recent credit crisis.

In the option-valuation literature, it is now accepted that Black and Scholes (1973) model's constant volatility assumption is violated in the data. The Heston (1993) stochastic volatility model relaxes this assumption by assuming that variance follows a Cox, Ingersoll, and Ross (1985) process and reacts to shocks correlated with

*As of end-June 2009.

Figure 1.1: Gross Market Value

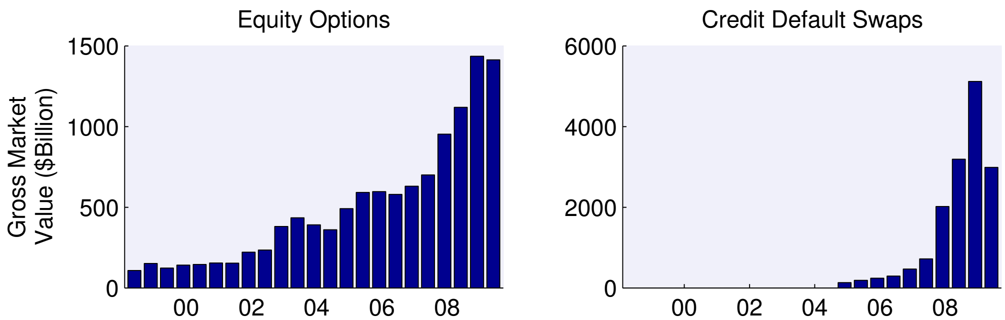

Evolution of the gross market value of equity options and credit default swaps on G10 markets.

those affecting the underlying's price process. Similarly to the Black and Scholes model, the Heston model provides a (quasi-) closed-form solution for option prices. Unfortunately, estimating the model is relatively difficult. This drawback is overcome by Heston and Nandi (2000) who propose a model in which the variance follows generalized autoregressive conditional heteroscedasticity (GARCH) process.

Option-valuation models that are based on GARCH variance processes are always straightforward to estimate. However, in order to obtain a closed-form solution for option prices, Heston and Nandi (2000) use a particular affine GARCH variance specification, a structure that implies that all moments of variance are affine in the current variance level. Yet, most GARCH models are of a nonaffine form.* Besides, the Heston and Nandi model assumes conditional normality, and the variance process is assumed to be a simple GARCH(1,1) as opposed to, for instance, a two-component model.† The first paper of this thesis, in Chapter 2, assesses the impact of these three typical modeling choices.

First, like Engle and Lee (1999), we find strong evidence in favour of component models from the standpoint of modeling daily return dynamics. When using option prices to assess the models, we also find strong evidence for the component structure in the affine GARCH models, but less so in the nonaffine models. Second, we

*See Christoffersen and Jacobs 2004.

†Christoffersen, Heston, and Jacobs (2006) provide an extension to Heston and Nandi's (2000) model that allows for Inverse Gaussian innovations to the return process. In continuous-time settings, Christoffersen, Heston, and Jacobs (2009) and Gauthier and Possamaï (2009) discuss two-factor versions of the Heston (1993) model.

consider nonaffine versus affine GARCH models. We compare the affine GARCH(1, 1) model with the nonaffine GARCH(1, 1) model of Hsieh and Ritchken (2005), who find strong support for the nonaffine specification. Our results support their findings, and we also find that the nonaffine models outperform affine models when allowing for component structures and nonnormal shocks. Third, we consider conditionally normal versus conditionally nonnormal models. We find that assuming GED shocks for the daily asset returns greatly improves the fit of all the models to daily returns, but the improvement in option valuation is much less evident. These results are most likely due to the choice of the nonnormal distribution used in this paper, the Generalized Error Distribution (GED). In light of the overwhelming literature highlighting the importance of accounting for nonnormality in option valuation, we believe that our results simply discredits the GED in option valuation settings.*

Hence, amongst many others, our results further support the growing consensus that two volatility components are required to properly describe the time-series properties of stock market volatility; one component accounting for transient shocks to the volatility process, another accounting for long-lasting shocks.† In Chapter 3, the second paper of this thesis focuses on providing economic insights into the macroeconomic determinants of this long-run volatility process. Most two-factor volatility models rely on latent, autoregressive volatility factors. However, to the extent that stock market volatility reflects the uncertainty or fear of market participants, volatility levels should be related to the current business conditions, which most likely affect expectations regarding *macroeconomic risk*.

Building on this insight, I develop a model, the MacroHV-MIDAS model, accounting for improving or deteriorating business conditions, and I show that they not only affect current volatility levels, but also the impact of volatility on option prices. 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 impli-

*Amongst many, many others, see Barndorff-Nielsen (1997), Bakshi, Cao, and Chen (1997), Bakshi, Kapadia, and Madan (2003), Carr, Geman, Madan, and Yor (2002), Chernov, Gallant, Ghysels, and Tauchen (2003), Carr and Wu (2003a,2003b,2004), Huang and Wu (2004), Eraker (2004), Bakshi, Carr, and Wu (2008), and Ornthanalai (2009).

†See, amongst others, Engle and Lee (1999), Andersen, Bollerslev, Diebold, and Ebens (2001), Alizadeh, Brandt, and Diebold (2002), and Christoffersen, Jacobs, Ornthanalai, and Wang (2008).

cations for financial derivatives. The model introduced in this paper is shown to outperform Duan's (1995) NGARCH 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.

These results, like those of Chapter 2, were obtained studying the behaviour of the volatility of the S&P 500 index. Assuming that the S&P 500 index proxies for the market portfolio, the index volatility can be seen as a proxy for market volatility.[*] In Chapter 4, the third paper of this thesis analyzes how the breakdown of a firm's volatility between its systematic and idiosyncratic components is reflected both in option prices and credit default swap spreads.

We show that two simple and internally consistent models for option prices and credit spreads, the Merton (1974) and Geske (1979) models, provide an intuitive prediction for both markets: options and credit default swaps on firms with more *systematic risk* exposure should be more expensive, all else equal. In other words, purchasing insurance on a firm with put options or credit default swaps should be costlier, after controlling for among other things total risk, leverage and risk free interest rates, the greater the systematic risk. This is a simple insight, a straightforward one in a CAPM world (Sharpe 1964; Lintner 1965; and Mossin 1966), yet options and credit derivatives are often viewed through the prism of risk-neutral (relative) pricing and, as a result, the wedge between risk-adjusted and physical probability return distributions tends to remain out of view.

Recently, Duan and Wei (2009) have shifted the focus onto the cross-sectional pricing of risk in stock option markets, relating prices to the proportions of systematic risk in equity volatilities. We document that their findings are robust to a more recent and broader data set, and then ask whether (as our comparative statics suggest) these findings have an analogue in the credit derivative markets. We find this to be the

[*] As discussed by Roll (1977), this assumption is a drastically simplifying one.

case: the proportion of the price of default insurance due to a risk premium is greater the larger the systematic risk exposure a firm has.

In summary, this thesis analyzes three closely interrelated themes of financial derivatives pricing: model risk, macroeconomic risk, and systematic risk. Better understanding and quantifying macroeconomic risk and its impact on financial derivatives could prove highly relevant in better understanding and hedging the risk inherent to option portfolios throughout the business cycle. Options on indexes can be used to hedge a portfolio's exposure to systematic risk, that is, of course, assuming that one is conscious of the impact of systematic risk on his portfolio. And the first step in quantifying and managing these risks is to know what model to trust or, more reasonably, when to trust which model.

2 Volatility Components, Affine Restrictions, and Nonnormal Innovations

Joint work with

Peter Christoffersen

Desautels Faculty of Management, McGill University, Montreal, Quebec, H3A 1G5, Canada, Copenhagen Business School, and CREATES (peter.christoffersen@mcgill.ca)

Kris Jacobs

C.T. Bauer College of Business, University of Houston, Houston, TX 77024, McGill University, and Tilburg University (kjacobs@bauer.uh.edu)

Yintian Wang

School of Economics and Management, Tsinghua University, Beijing 100083, China (wangyt2@sem.tsinghua.edu.cn)

Forthcoming in the Journal of Business and Economic Statistics

Abstract Here we assess the return fitting and option valuation performance of generalized autoregressive conditional heteroscedasticity (GARCH) models. We compare component versus GARCH$(1, 1)$ models, affine versus nonaffine GARCH models, and conditionally normal versus nonnormal GED models. We find that nonaffine models dominate affine models in terms of both fitting returns and option valuation. For the affine models, we find strong evidence in favor of the component structure for both returns and options; for the nonaffine models, the evidence is less convincing in option valuation. The evidence in favor of the nonnormal GED models is strong when fitting daily returns, but not when valuing options.

Keywords Affine; Component model; Generalized autoregressive conditional heteroscedasticity; Long memory; Normality; Option valuation; Volatility.

2.1 Introduction

Following the groundbreaking work of Engle (1982) and Bollerslev (1986), generalized autoregressive conditional heteroscedasticity (GARCH) models have become an ubiquitous tool kit in empirical finance. In this article we assess the ability of eight different GARCH models to fit daily return dynamics and their ability to match market prices of options in a sample of close to 22,000 contracts. The eight models that we investigate differ along three dimensions.

First, we consider component models versus GARCH(1,1) models. Engle and Lee (1999) were the first to develop a component GARCH model, which they built from the nonaffine NGARCH(1,1) model analyzed by Engle and Ng (1993) and Duan (1995). Component GARCH models can be viewed as a convenient way of incorporating long-memory-like features into a short-memory model, at least for the horizons relevant for option valuation. Bollerslev and Mikkelsen (1999) find support for a long-memory GARCH option valuation model applied to long-maturity options. We consider options with up to 1-year maturity, for which the component models are likely to provide good approximations to long-memory processes. Maheu (2005) presented Monte Carlo evidence indicating that a component model similar to those presented in this article can capture long-memory volatility dynamics. Adrian and Rosenberg (2008) demonstrated the relevance of the component volatility structure for cross-sectional asset pricing.

Second, we consider nonaffine versus affine GARCH models. Most GARCH models are of a nonaffine form (see Christoffersen and Jacobs 2004), but Heston and Nandi (2000) developed a class of affine GARCH models. From the affine GARCH(1,1) specification, Christoffersen, Jacobs, Ornthanalai, and Wang (2008) developed an affine GARCH component model, which we also consider herein. The affine GARCH(1,1) model was compared with the nonaffine NGARCH(1,1) model by Hsieh and Ritchken (2005), who found strong support for the nonaffine specification.

Third, we consider conditionally normal versus conditionally nonnormal models. In particular, we modify the foregoing four conditionally normal GARCH models by modeling the return shock using a generalized error distribution (GED). The GED distribution, suggested by Duan (1999) for its tractability in asset return modeling, conveniently nests the normal distribution. A skewed version of the GED distribution was developed by Theodossiou (2001) and has been used for option valuation by Lehnert (2003). Lehnert found support for a nonaffine EGARCH model with skewed GED shocks when comparing its option pricing performance with that of the affine GARCH(1,1) model of Heston and Nandi (2000), which has normal innovations. We extend Lehnert's work by analyzing whether the improvement comes from the nonaffine variance dynamic, from the nonnormal shocks, or from both features.

We estimate the resulting eight models using maximum likelihood (ML) estimation on S&P 500 returns. This empirical comparison allows us to compare the importance of three types of modeling assumptions: (a) the importance of the component structure versus the simpler and more parsimonious GARCH(1,1) structure, (b) the restrictions of the affine structure, and (c) the importance of nonnormal return innovations. We find that the likelihood criterion based on return data strongly favors the component models in all cases, as well as the nonnormal return innovations. Although the affine models are not nested in the nonaffine models, comparing the likelihoods suggests that the nonaffine models fit the return data the best.

Using the ML estimates, we characterize key properties of each model: multiday variance forecasting functions, conditional volatility of variance paths, and conditional correlations between returns and variance. We find important differences between affine and nonaffine models suggesting that the nonaffine structure provides more flexibility in a parsimonious fashion. We also find substantial differences between GARCH(1,1) and component models and between models with normal and nonnormal innovations.

When we use the estimated model parameters for option valuation, we again find strong support for the nonaffine variance specifications, but less evident support for the nonnormal return innovations. The component structure yields significant improvements in the affine class of models. In the nonaffine class, it yields improvements for long-maturity and out-of-the-money options.

The article is organized as follows. In Section 2 we develop the eight GARCH-based asset models that we investigate empirically in this work. In Section 3 we report the model estimates from daily returns and present some key dynamic properties of the models. In Section 4 we provide the theoretical mappings from physical to risk-neutral dynamics by applying the general approach of Duan (1999). In Section 5 we present the empirical results from using the GARCH models in option valuation, as well as an economic analysis of the option pricing errors. In Section 6 we conclude and suggest some promising avenues for future research.

2.2 Asset Return Models

In this section we introduce the eight GARCH models to be used for option valuation. The eight models cover all of the possibilities in our three-way comparison: GARCH(1,1) versus component GARCH, affine versus nonaffine GARCH, and normal versus GED-distributed return shocks.

2.2.1 The Affine GARCH (1,1) Model With Normal Shocks

We first introduce the affine normal GARCH(1,1) model of Heston and Nandi (2000). The return dynamics on the underlying asset are

$$R_{t+1} \equiv \ln\left(\frac{S_{t+1}}{S_t}\right) = r + \lambda h_{t+1} + \sqrt{h_{t+1}} z_{t+1}, \tag{2.1}$$
$$h_{t+1} = w + b' h_t + a\left(z_t - c\sqrt{h_t}\right)^2,$$

where S_{t+1} is the underlying asset price on the close of day $t+1$, r is the risk-free rate, λ is the price of risk, z_t is the iid $N(0,1)$ return shock, and h_{t+1} is the daily variance on day $t+1$ that is known at the end of day t. We refer to this model as the AGARCH(1, 1)-N model.

Note that c renders the variance response asymmetric to positive versus negative innovations in returns. If c is 0, then the variance dynamic is symmetric in z_t, and the conditional distribution of returns will be largely symmetric at all horizons, because the distribution of z_t is symmetric as well. In that case, the only source of asymmetry is the conditional mean return,

$$\mathrm{E}_t\left[\, R_{t+1} \,\right] = r + \lambda h_{t+1}, \tag{2.2}$$

and this effect typically is very small in magnitude.

We next derive some features of the model that are particularly important for its performance in option valuation. The model's unconditional variance can be derived as

$$\mathrm{E}\left[\, h_{t+1} \,\right] = \sigma^2 = \frac{w + a}{1 - ac^2 - b'}. \tag{2.3}$$

If we use this expression to substitute for w and also define variance persistence as $b = b' + ac^2$, then we can write

$$h_{t+1} = \sigma^2 + b\left(h_t - \sigma^2\right) + a\left(z_t^2 - 1 - 2c\sqrt{h_t}z_t\right). \tag{2.4}$$

The options that we analyze herein have maturities of between 1 week and 1 year. Thus it is important to gauge the model's properties at multiday horizons. In this regard, consider the conditional variance k days ahead,

$$\mathrm{E}_t\left[\, h_{t+k} \,\right] = \sigma^2 + b^{k-1}\left(h_{t+1} - \sigma^2\right). \tag{2.5}$$

Although by design, the 1-day-ahead variance is deterministic in the GARCH models, the multiday variance is stochastic, and its distribution is important for option pricing as well. For 2 days ahead, the conditional variance of variance is easily derived from (2.4) as

$$\mathrm{Var}_t\left[h_{t+2}\right] = 2a^2 + 4a^2c^2h_{t+1}. \tag{2.6}$$

Note that the variance of h_{t+2} is linear in h_{t+1}, which is a defining characteristic of the affine GARCH model. Note further that if c is 0, then the future variance will have constant conditional variance. This is at odds with the empirical evidence reported by, for example, Jones (2003), who found that the volatility of implied options volatility is higher when the level of implied options volatility is larger. Thus in the affine normal GARCH(1,1) model, the c parameter is needed both to provide conditional variance of variance dynamics and to provide substantial conditional distribution asymmetry. This double duty may cause a tension in the model; we revisit this later.

The relationship between future variance and return is also of interest for option valuation. The so-called "leverage effect" was noted by Black (1976), who observed a negative correlation between volatility and returns. To describe this relationship, we consider the conditional covariance

$$\mathrm{Cov}_t\left[R_{t+1}, h_{t+2}\right] = E_t\left[\sqrt{h_{t+1}}z_{t+1}a\left(z_{t+1}^2 - 1 - 2c\sqrt{h_{t+1}}z_{t+1}\right)\right] = -2ach_{t+1}. \tag{2.7}$$

Note that because a must be strictly positive to ensure that the GARCH process is identified and positive, the sign of the leverage effect is driven entirely by c. From this conditional covariance, the conditional correlation is easily derived as

$$\mathrm{Corr}_t\left[R_{t+1}, h_{t+2}\right] = \frac{-2c\sqrt{h_{t+1}}}{\sqrt{2+4c^2 h_{t+1}}}. \tag{2.8}$$

Note that this conditional correlation is time-varying, which is a relatively unique property of the affine GARCH model.

2.2.2 The Affine GARCH Component Model With Normal Shocks

Many investigators (see, e.g., Bollerslev and Mikkelsen 1999) have found that simple exponential decay of the conditional variance to its unconditional value, in (2.5), is too fast. This motivates the affine normal GARCH component model developed by Christoffersen, Jacobs, Ornthanalai, and Wang (2008), who built on the work of Engle and Lee (1999) and Heston and Nandi (2000). The return and variance dynamics are now

$$\begin{aligned} R_{t+1} &= r + \lambda h_{t+1} + \sqrt{h_{t+1}} z_{t+1}, \qquad (2.9) \\ h_{t+1} &= q_{t+1} + \beta(h_t - q_t) + \alpha\left(z_t^2 - 1 - 2\gamma_1\sqrt{h_t} z_t\right), \\ q_{t+1} &= \sigma^2 + \rho\left(q_t - \sigma^2\right) + \varphi\left(z_t^2 - 1 - 2\gamma_2\sqrt{h_t} z_t\right). \end{aligned}$$

Instead of mean-reverting to a constant unconditional variance, the conditional variance, h_{t+1}, now moves around a long-run component, q_{t+1}, which itself mean-reverts to the constant unconditional variance, σ^2. Furthermore, the two parameters γ_1 and γ_2 in the component model allow for a different degree of asymmetry in the two components, $(h_{t+1} - q_{t+1})$ and q_{t+1}. We refer to this model as the AGARCH(C)-N model.

The added dynamics in this model chiefly serve to generate more flexible dynamics in the multi-day-ahead conditional variance. We now have

$$\begin{aligned} \mathrm{E}_t\left[h_{t+k}\right] &= \mathrm{E}_t\left[q_{t+k} + (h_{t+k} - q_{t+k})\right] \\ &= \sigma^2 + \rho^{k-1}\left(q_{t+1} - \sigma^2\right) + \beta^{k-1}\left(h_{t+1} - q_{t+1}\right), \end{aligned} \tag{2.10}$$

which clearly allows for slower mean-reversion than (2.5). We refer to ρ as long-run persistence and to β as short-run persistence.

This component model also offers greater flexibility in the conditional variance of variance dynamic, which is now

$$\mathrm{Var}_t\left[h_{t+2}\right] = 2\left(\alpha+\varphi\right)^2 + 4\left(\alpha\gamma_1+\varphi\gamma_2\right)^2 h_{t+1} \tag{2.11}$$

so that the affine structure is preserved. The conditional covariance and correlation are

$$\mathrm{Cov}_t\left[R_{t+1}, h_{t+2}\right] = -2\left(\alpha\gamma_1+\varphi\gamma_2\right) h_{t+1} \tag{2.12}$$

and

$$\mathrm{Corr}_t\left[R_{t+1}, h_{t+2}\right] = \frac{-2\left(\alpha\gamma_1+\varphi\gamma_2\right)\sqrt{h_{t+1}}}{\sqrt{2\left(\alpha+\varphi\right)^2 + 4\left(\alpha\gamma_1+\varphi\gamma_2\right)^2 h_{t+1}}}. \tag{2.13}$$

Note that in these formulas, $(\alpha + \varphi)$ has replaced a in the AGARCH(1,1)-N model's formula, and similarly $(\gamma_1\alpha + \gamma_2\varphi)$ has replaced c. Thus, whereas c had to perform double duty (creating asymmetry and variance of variance dynamics) in the AGARCH(1,1)-N model, the component model offers much added flexibility. The return asymmetry and variance of variance are now driven by two sources, parameterized by γ_1 and γ_2.

2.2.3 The Nonaffine GARCH(1,1) Model With Normal Shocks

The benchmark NGARCH(1,1)-N model of Engle and Ng (1993), which also was used for option valuation by Duan (1995), is defined as

$$\begin{aligned} R_{t+1} &= r + \lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1} + \sqrt{h_{t+1}}z_{t+1}, \\ h_{t+1} &= w + b'h_t + ah_t\left(z_t - c\right)^2. \end{aligned} \tag{2.14}$$

The parameter c again renders the variance response asymmetric to positive versus negative return shocks and creates asymmetry in the conditional distribution of multiday returns beyond that created by the conditional return mean,

$$\mathrm{E}_t\left[\,R_{t+1}\,\right] = r + \lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1}. \tag{2.15}$$

Although this conditional mean specification differs from that used in the affine model, we use it because it will generate a risk-neutral conditional variance specification that is similar to the physical one, as we describe in Section 4. Similarly, the affine conditional mean specification in (2.2) will generate an affine conditional variance under the risk-neutral measure. We discuss this issue in more detail later in this article.

The unconditional variance is now

$$E\left[h_{t+1}\right] = \sigma^2 = \frac{w}{1 - b' - a\left(1 + c^2\right)}. \tag{2.16}$$

Defining variance persistence as $b = b' + a(1 + c^2)$, we can rewrite the conditional variance as

$$h_{t+1} = \sigma^2 + b\left(h_t - \sigma^2\right) + ah_t\left(z_t^2 - 1 - 2cz_t\right). \tag{2.17}$$

The conditional variance k days ahead has the same form as in the affine model,

$$\mathrm{E}_t\left[\, h_{t+k} \,\right] = \sigma^2 + b^{k-1}\left(h_{t+1} - \sigma^2\right), \tag{2.18}$$

and the conditional variance of variance can be derived from (2.17) as

$$\mathrm{Var}_t\left[h_{t+2}\right] = a^2\left(2 + 4c^2\right) h_{t+1}^2. \tag{2.19}$$

Thus the variance of h_{t+2} is now quadratic in h_{t+1}, whereas it was linear in the affine model. Note also that even if c is 0 in the nonaffine model, the future variance will still have a dynamic variance, now driven by a.

The conditional covariance in this model is

$$\mathrm{Cov}_t\left[R_{t+1}, h_{t+2}\right] = -2ach_{t+1}^{3/2}. \tag{2.20}$$

Thus the leverage effect again is driven by c, but now the covariance is nonlinear in h_{t+1}. Given the foregoing formula for conditional covariance, the conditional correlation is simply

$$\mathrm{Corr}_t\left[R_{t+1}, h_{t+2}\right] = \frac{-2c}{\sqrt{2 + 4c^2}}. \tag{2.21}$$

Note that conditional correlation in the nonaffine model is constant, whereas it was time-varying in the affine model. Thus, along this dimension, the affine model seemingly offers more flexibility.

2.2.4 The Nonaffine GARCH Component Model With Normal Shocks

As described in Section 2.2, this model is obtained by replacing the constant σ^2 in the NGARCH(1,1)-N model with a time-varying, long-run component q_{t+1}. We write

$$R_{t+1} = r + \lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1} + \sqrt{h_{t+1}}z_{t+1},$$
$$h_{t+1} = q_{t+1} + \beta(h_t - q_t) + \alpha h_t\left(z_t^2 - 1 - 2\gamma_1 z_t\right), \tag{2.22}$$
$$q_{t+1} = \sigma^2 + \rho\left(q_t - \sigma^2\right) + \varphi h_t\left(z_t^2 - 1 - 2\gamma_2 z_t\right),$$

and refer to this model as the NGARCH(C)-N model.

Once again, the added dynamics in this model chiefly serve to generate more flexible dynamics in the multi-day-ahead conditional variance. Here the multiday conditional variance is

$$\begin{aligned}\mathrm{E}_t\left[h_{t+k}\right] &= \mathrm{E}_t\left[q_{t+k} + (h_{t+k} - q_{t+k})\right] \\ &= \sigma^2 + \rho^{k-1}\left(q_{t+1} - \sigma^2\right) + \beta^{k-1}\left(h_{t+1} - q_{t+1}\right).\end{aligned} \tag{2.23}$$

The conditional variance of the variance dynamic is

$$\mathrm{Var}_t\left[h_{t+2}\right] = \left[2\left(\alpha+\varphi\right)^2 + 4\left(\alpha\gamma_1 + \varphi\gamma_2\right)^2\right]h_{t+1}^2, \tag{2.24}$$

which has contributions from both components and again is quadratic in h_{t+1}. The conditional covariance and correlation are

$$\mathrm{Cov}_t\left[R_{t+1}, h_{t+2}\right] = -2\left(\alpha\gamma_1 + \varphi\gamma_2\right)h_{t+1}^{3/2}, \tag{2.25}$$

which is nonlinear in h_t, and

$$\mathrm{Corr}_t\left[R_{t+1}, h_{t+2}\right] = \frac{-2\left(\alpha\gamma_1 + \varphi\gamma_2\right)}{\sqrt{2\left(\alpha+\varphi\right)^2 + 4\left(\alpha\gamma_1 + \varphi\gamma_2\right)^2}}, \tag{2.26}$$

which again is constant.

2.2.5 Generalized Error Distribution Shocks

The assumption that the daily return shock, z_t, is normally distributed is typically rejected empirically for daily asset returns. Our empirical analysis here is no exception. Note, however, that while the conditional 1-day distribution is normal when z_t is normal, the multiday distribution is not normal, and neither is the unconditional distribution. Thus the effect of the normal innovation assumption on option valuation in a GARCH model is not straightforward. Our analysis investigates whether these dynamics are sufficient to fit the underlying asset return as well as the option prices on the underlying asset, or whether the conditional normality assumption should be relaxed.

Following Duan (1999), we assume that the iid return shock, denoted by z_t in the normal case earlier, now follows the GED and is denoted by ζ_t. Once it is normalized to get a 0 mean and unit variance, we have the probability density function

$$g_\nu(\zeta) = \frac{\nu}{2^{1+\frac{1}{\nu}}\theta(\nu)\Gamma\left(\frac{1}{\nu}\right)} \exp\left(-\frac{1}{2}\left|\frac{\zeta}{\theta(\nu)}\right|^\nu\right) \qquad \text{for } 0 < \nu \leq \infty,$$

where Γ is the gamma function and $\theta(\nu) = (2^{-2/\nu}\Gamma(\frac{1}{\nu})/\Gamma(\frac{3}{\nu}))^{1/2}$.

The parameter ν determines the thinness of the density tails. For $\nu < 2$, the density function has fatter tails than those of the normal distribution, and the opposite is true for $\nu > 2$. The expected simple return exists as long as $\nu > 1$, which thus is a natural lower bound in financial return applications.

The GED innovation ζ has a skewness of 0 and a kurtosis of $\kappa(\nu) = \Gamma(\frac{5}{\nu})\Gamma(\frac{1}{\nu})/\Gamma(\frac{3}{\nu})^2$. In the special case where $\nu = 2$, we get $\kappa(2) = 3$, and because $\Gamma(\frac{1}{2}) = \sqrt{\pi}$, we get

$$g_2(\zeta) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\zeta^2\right)$$

so that the standardized GED conveniently nests the standard normal distribution that obtains when $\nu = 2$. Nelson (1991), Hamilton (1994), and Duan (1999) have provided more details on the properties of the GED distribution.

Replacing the normal distribution by the GED distribution in each of the four foregoing models provides four new models. The resulting eight models allow us to study the three dimensions of modeling in which we are interested: GARCH(1,1) versus component GARCH, affine versus nonaffine GARCH, and normal versus nonnormal return shocks.

The GED distribution does not directly affect the variance persistence and, consequently, the multiday conditional variance in the first four models that we considered. The functional form for the conditional covariance is also unchanged in these models. But the excess kurtosis of the GED distribution does affect the conditional variance of variance in the four GED models; we now have

$$\begin{aligned}
\text{AGARCH(1, 1)-GED:}\ & \mathrm{Var}_t(h_{t+2}) = (\kappa(\nu)-1)a^2 + 4a^2c^2h_{t+1}, \\
\text{AGARCH(C)-GED:}\ & \mathrm{Var}_t(h_{t+2}) = (\kappa(\nu)-1)(\alpha+\varphi)^2 + 4(\alpha\gamma_1+\varphi\gamma_2)^2h_{t+1}, \\
\text{NGARCH(1, 1)-GED:}\ & \mathrm{Var}_t(h_{t+2}) = (\kappa(\nu)-1+4c^2)a^2h_{t+1}^2, \\
\text{NGARCH(C)-GED:}\ & \mathrm{Var}_t(h_{t+2}) = \left[(\kappa(\nu)-1)(\alpha+\varphi)^2 + 4(\alpha\gamma_1+\varphi\gamma_2)^2\right]h_{t+1}^2,
\end{aligned}$$

where $\kappa(\nu) = \frac{\Gamma(\frac{5}{\nu})\Gamma(\frac{1}{\nu})}{\Gamma(\frac{3}{\nu})^2}$ denotes the kurtosis under the GED distribution. This in turn affects the conditional correlation between return and volatility. In each of the conditional correlation formulas for the first four models, the 2 in the denominator is replaced by a $(\kappa(\nu)-1)$ term.

2.3 Asset Return Empirics

In this section we present the empirical results from fitting our GARCH models to daily returns. First, we use ML estimation on a long time series of S&P 500 return data to estimate eight models: AGARCH(1,1)-N, AGARCH(C)-N, NGARCH(1,1)-N, NGARCH(C)-N, and the four GED-based models. We then discuss the parameter estimates and their implications for the salient properties of the models. The eight models allow us to make three types of comparisons: component models versus GARCH(1, 1) models, affine models versus nonaffine models, and nonnormal innovations versus normal innovations.

2.3.1 Parameter Estimates From Daily Return Data

Table 2.1 presents the ML estimation results obtained using daily returns data from July 1, 1962 through December 31, 2001, obtained from CRSP. Standard errors, calculated as done by Bollerslev and Wooldridge (1992), are given in parentheses. The table reports the physical conditional variance parameters, as well as the price of risk, λ. The estimates of λ must be positive to guarantee positive excess log returns.

We use variance targeting to control the unconditional variance level across models, which is important for the subsequent option valuation exercise. We thus force

Table 2.1: Maximum Likelihood Estimation on Daily S&P 500 Returns, 1962–2001

	AGARCH-N		**NGARCH-N**		**AGARCH-GED**		**NGARCH-GED**	
	GARCH(1,1)	Components	GARCH(1,1)	Components	GARCH(1,1)	Components	GARCH(1,1)	Components
w	8.89E−21	7.02E−07	5.90E−07	1.74E−07	1.27E−09	5.89E−07	5.39E−07	1.64E−07
λ	0.00002	1.00495	0.03768	0.03390	0.56677	1.78607	0.03984	0.03674
	(1.83E−05)	(3.53E−06)	(8.62E−03)	(8.98E−03)	(1.42E−05)	(1.64E−04)	(8.30E−03)	(9.67E−03)
a, α	3.342E−06	2.132E−06	6.253E−02	3.696E−02	3.105E−06	1.705E−06	5.982E−02	3.071E−02
	(1.72E−07)	(1.40E−07)	(3.95E−03)	(3.45E−03)	(2.52E−04)	(1.35E−07)	(4.12E−03)	(4.72E−03)
b', β	0.89921	0.74928	0.90825	0.89262	0.90297	0.83454	0.91133	0.91320
	(2.47E−02)	(1.71E−04)	(5.30E−03)	(2.20E−02)	(6.90E−06)	(2.61E−05)	(5.71E−03)	(1.89E−02)
c, γ_1	135.7520	297.2247	0.5972	1.6588	139.7188	313.8362	0.6136	1.7759
	(2.47E−02)	(5.42E−05)	(4.67E−02)	(4.97E−02)	(1.70E−06)	(8.63E−04)	(5.77E−02)	(1.98E−01)
φ		1.739E−06		3.393E−02		1.524E−06		3.341E−02
		(4.84E−08)		(3.36E−03)		(1.10E−07)		(4.15E−03)
ρ		0.99176		0.99796		0.99309		0.99807
		(2.34E−06)		(5.21E−04)		(4.31E−04)		(6.13E−04)
γ_2		71.40695		0.38247		57.94967		0.38521
		(2.94E−05)		(4.76E−02)		(1.22E−03)		(5.64E−02)
v					1.34637	1.41600	1.43298	1.45868
					(1.17E−07)	(3.04E−03)	(2.66E−02)	(2.84E−02)
Properties								
Log-likelihood	33,954	34,129	34,130	34,201	34,192	34,310	34,309	34,352
Annual volatility target	14.66	14.66	14.66	14.66	14.66	14.66	14.66	14.66
Total variance persistence	0.9608	0.9979	0.9931	0.9998	0.9636	0.9989	0.9937	0.9998
Empirical z kurtosis	8.9123	7.4070	6.6153	5.7752	9.1408	7.9727	6.6575	5.9424
Model z kurtosis	3.0000	3.0000	3.0000	3.0000	4.1797	3.9732	3.9274	3.8612
Annual vol. of variance	0.2289	0.3507	0.2497	0.3852	0.2341	0.3016	0.2725	0.3726
Average correlation	−0.8318	−0.8949	−0.6452	−0.8289	−0.7784	−0.8557	−0.5828	−0.7792

We use daily total returns from July 1, 1962 to December 31, 2001 on the S&P 500 index to estimate the GARCH models using Maximum Likelihood. We use variance targeting to fix w so that the annual volatility in each model is 14.66%. Standard errors are calculated using Bollerslev and Wooldridge (1992). Annual vol. of variance refers to the sample mean of the annualized conditional volatility of variance path in each model. Correlation refers to the sample mean of the time-varying conditional correlation between variance and return in the affine models and to the constant model implied correlation in the nonaffine models.

the annualized return standard deviation to be 14.66%. This technique fixes the parameter w in each model, and thus we do not report standard errors for w in Table 2.1.

In the AGARCH(1,1) cases, the unconditional variance is defined by $\sigma^2 = \frac{w+a}{1-b}$. Thus, if the data warrant a high a (perhaps to match variance of variance), then w will be small, to match σ^2. If left unconstrained, the w estimate may be negative, which yields the possibility of a negative conditional variance and thus causes problems in the subsequent Monte Carlo computation of option prices. Therefore, we constrain w to be positive in the estimation.

Table 2.1 reports the total variance persistence in each model. If we substitute out q_{t+1} and q_t from the h_{t+1} equation in the component models, then persistence can be computed as the sum of the coefficients on h_t and h_{t-1}. Thus the total persistence in the component models is $\rho + (1-\rho)\beta$. In the GARCH(1,1) models, persistence is simply b. Note from Table 2.1 that while the GARCH(1, 1) models have high persistence, the persistence is even higher for each corresponding component model; this is particularly true for the affine models. The very large component variance persistence is driven by a large long-run component persistence, ρ, plus the contribution from [$(1-\rho)$ times] the less persistent short-run component, β.

In the GARCH(1,1) models, the correlation between return and conditional variance is driven by c, which is, as expected, significantly positive in all cases. In the component models, the correlation is driven by a combination of γ_1 and γ_2, both of which are significantly positive in all four component models. Thus both the long-run and short-run components contribute to the overall correlation with the expected sign. The average conditional correlations between the return and conditional variance, reported in the third-to-last row of the table, are all negative, as expected. The results demonstrate that for each set of models, the component model displays a more pronounced leverage effect than its GARCH(1,1) counterpart, in that the average correlation is more negative.

The variance of variance is driven mainly by the a parameter in the GARCH(1,1) models and by the α and ρ parameters in the component models. In the GED models, the ν parameter also contributes. Table 2.1 also reports the overall unconditional volatility of variance (annualized square root of variance of variance). Note again that in each case, the component model has a larger volatility of variance than its GARCH(1,1) counterpart. Thus three important empirical regularities emerge when comparing component models with their GARCH(1,1) counterparts: the component models allow us to (simultaneously) capture a larger variance persistence, a larger leverage effect, and a larger volatility of variance than their GARCH(1, 1) counterparts.

Finally, Table 2.1 presents log-likelihood values for each model. In all cases, the component model has a significantly larger log-likelihood than the nested GARCH(1,1) model. Comparing the GED models with their normal counterparts also shows that the GED-based models have significantly larger log-likelihood values. Thus this return-based analysis strongly favors the component models over the GARCH(1,1) models and favors the GED models over the normal models. Although the affine and nonaffine models are not nested, a casual comparison of the log-likelihoods suggests that the nonaffine GARCH models also are strongly preferred over the affine GARCH models in all four cases.

2.3.2 Time Series Properties

To explore the asset return models further, here we plot various key dynamic properties of the models for the period 1989-2001. This period includes the dates for the option valuation exercise presented in Section 5.

Figure 2.1 plots the conditional volatility for the period 1989-2001. To be exact, the annualized conditional standard deviation is plotted as a percentage, that is, $100 * \sqrt{252 h_{t+1}}$. Note that the conditional volatility patterns across the four GARCH(1,1) models in the left column and the corresponding four component models in the right column display some similarities. The models all capture the low volatility during the equity market runup in 1993-1998, preceded by higher volatility during the first Gulf War and the 1990-1991 recession. The LTCM and Russia crises in the fall of 1998 are evident, as is the higher volatility during the dot-com bust and the 2001 recession in the later part of the sample.

But Figure 2.1 also reveals some important differences among the models. The nonaffine models (in the second and fourth rows) appear to display much more

Figure 2.1: Conditional Volatility Paths

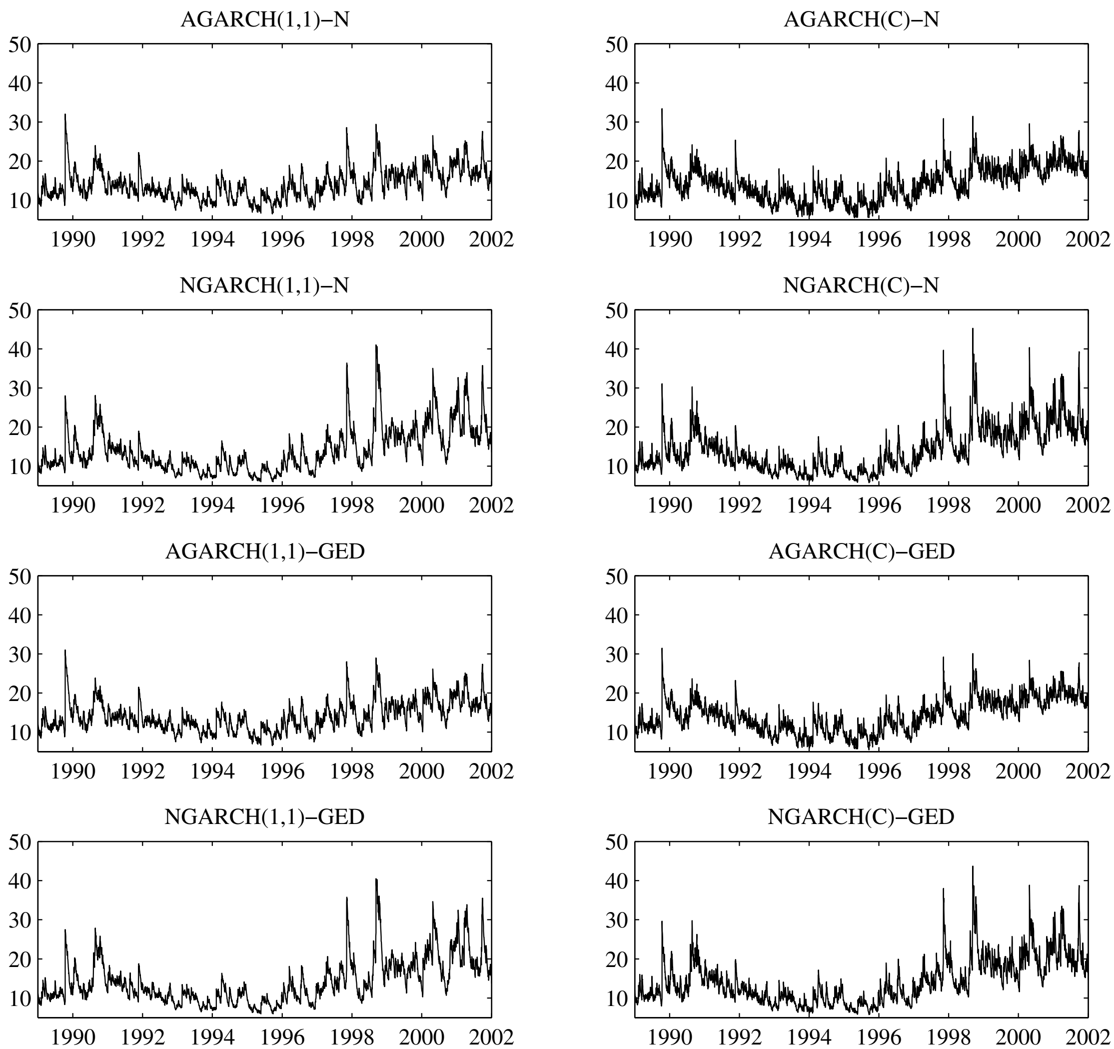

The annualized conditional volatility in percent, $100 * \sqrt{252h_{t+1}}$*, is plotted for each of the eight models that we consider. The parameter values for the underlying GARCH models are obtained from ML estimation on daily S&P 500 returns, as reported in Table 2.1.*

Figure 2.2: Conditional Volatility of Variance and Conditional Correlation Paths – Conditionally Normal Models

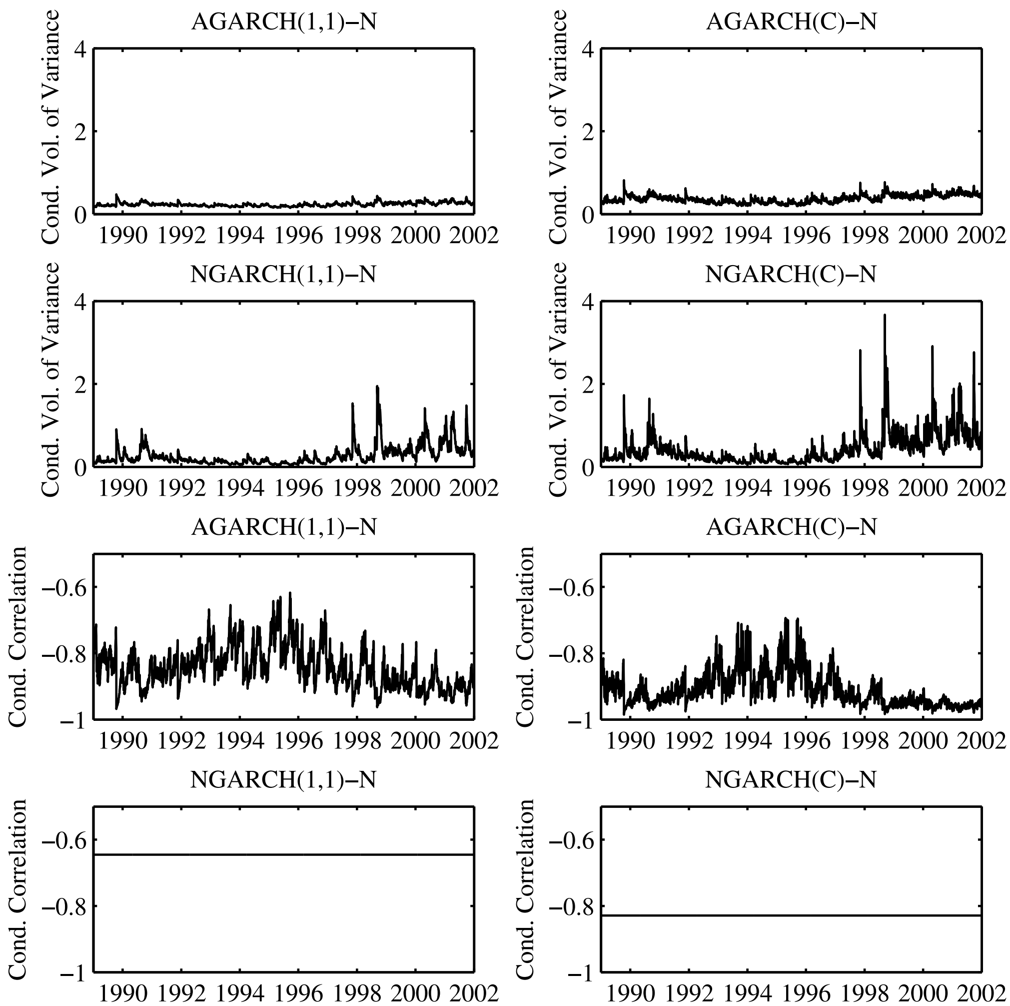

We plot the annualized conditional volatility of variance path in percent, $100 * 252 * \sqrt{Var_t\left(h_{t+2}\right)}$ *(top four panels) and the conditional correlation path,* $Corr_t\left(R_{t+1}, h_{t+2}\right)$ *(bottom four panels) for the four conditionally normal models. The parameter values for the underlying GARCH models are from Table 2.1.*

variation in the conditional volatility during the second half of the sample than do the two affine models (in the first and third rows).

Figure 2.2 plots the annualized conditional volatility of variance path, $100 * 252\sqrt{\text{Var}_t(h_{t+2})}$ for each of the four conditionally normal models in the top four panels. These plots confirm the findings presented in Figure 2.1. The nonaffine models in the second row of Figure 2.2 display a much larger volatility of variance than the affine models in the first row. This is true for both the GARCH(1,1) models in the left column and the component models in the right column. The nonaffine models display a much larger volatility of variance during the first Gulf War and the 1990-1991 recession, and especially during the LTCM and Russia crises in the fall of 1998 and the dot-com bust in 2000-2002.

The bottom two rows of Figure 2.2 plot the conditional correlation path, $\text{Corr}_t(R_{t+1}, h_{t+2})$, for each of our four conditionally normal models. Note that, as derived earlier, the nonaffine models imply a constant conditional correlation and thus exhibit a flat line in the plot. The affine models instead have time-varying correlation and imply a conditional correlation very close to -1 when economic events drive volatility up, as during, for example, the 1990-1991 recession and from 1999 onward. During the equity market runup in the mid 1990s, the conditional correlation implied by the affine models is much lower in magnitude.

Table 2.2: Model Properties and RMSE from Option Valuation

	AGARCH-N		**NGARCH-N**		**AGARCH-GED**		**NGARCH-GED**	
Correlation	GARCH(1,1)	Components	GARCH(1,1)	Components	GARCH(1,1)	Components	GARCH(1,1)	Components
Conditional								
- Volatility	1.0000	0.9297	0.9373	0.9297	0.9995	0.9306	0.9365	0.9337
- Vol. of variance	1.0000	0.9267	0.7499	0.7571	0.9993	0.9283	0.7549	0.7667
- Correlation	1.0000	0.9113	0.0000	0.0000	0.9980	0.9270	0.0000	0.0000
Option RMSE								
Overall RMSE ($)	2.6927	1.8138	1.5875	1.3820	2.6819	1.7404	1.4599	1.4270
RMSE / Avr call price	0.0965	0.0650	0.0569	0.0495	0.0961	0.0623	0.0523	0.0511

We compute the correlation between the AGARCH(1,1) and each of the other models for the conditional volatility, the conditional volatility of variance, and the conditional correlation between return and variance. We then use the MLE parameters from Table 2.1 to risk neutralize the models and compute option prices. From the model option prices we compute the root mean squared error (RMSE) in dollars as well as divided by the average option price for the sample.

The top row of Table 2.2 provides the correlation between conditional volatility in the AGARCH(1,1)-N model and each of the other models. As Figure 2.1 suggests, these correlations are very high. Not surprisingly, the correlation with the

AGARCH(1,1)-GED model is virtually 1. The second row in Table 2.2 gives the correlation between the conditional volatility of variance in the AGARCH(1, 1)-N model and the other models. This correlation is still high when computed for other affine models but is now considerably lower for nonaffine models. The third row in Table 2.2 computes the correlation between the conditional return-variance correlation in the AGARCH(1, 1)-N model and the other models. Note again that this correlation is very high when computed for other affine models. It is, of course, zero for the nonaffine models, which have a constant conditional correlation over time.

2.3.3 Variance Term Structure Properties

While Figures 2.1 and 2.2 depict various aspects of the dynamics of the 1-day-ahead conditional distribution, Figure 2.3 captures the properties of the variance dynamics across longer horizons. In this figure, we plot the expected future conditional variance from 1 to 252 days ahead. The dashed lines in Figure 2.3 denote multiday variance forecasts starting from a low current spot variance corresponding to the 25th percentile of the path of variances from 1962 to 2001. The solid lines in Figure 2.3 denote multiday variance forecasts starting from a high current spot variance corresponding to the 75th percentile of the path of variances from 1962 to 2001. In the component models, the 25th and 75th percentiles are used for the current spot values for both components. The forecasts are normalized by the unconditional variance and shown in percentage terms, so that for each model, we are plotting $100 * (\mathrm{E}_t [h_{t+k}] - \sigma^2)/\sigma^2$ against horizon k.

Figure 2.3 shows that the variance term structure properties vary strongly across models. In the AGARCH(1,1) models, the relatively low daily variance persistence implies that the conditional variance forecasts converge to their unconditional levels after around 100 days. This is true for both the normal and GED version of the model. The AGARCH(C) models display some variation in variance forecasts across initial spot variance up until 252 days ahead, comparable to the variation displayed by the nonaffine GARCH(1,1) models. The nonaffine component models display the most variation in variance forecasts at long horizons. These differences across models should have important implications for the option pricing properties.

Figure 2.3: Variance Forecasts across Forecast Horizons

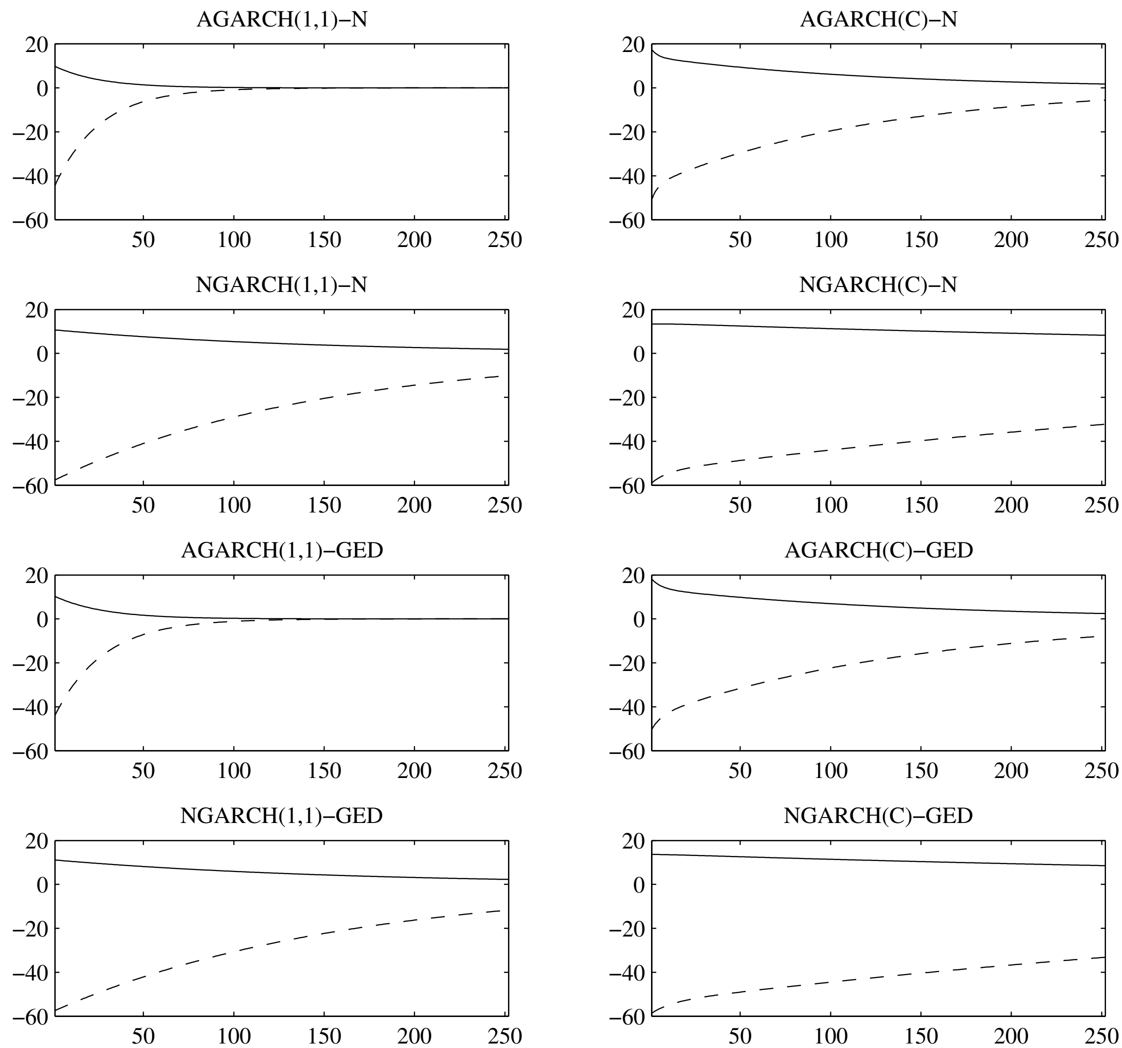

We plot the normalized k-day ahead variance forecast $100 * \left(\mathrm{E}_t \left[h_{t+k} \right] - \sigma^2 \right) / \sigma^2$ *for a low initial spot variance (dashed line) and a high initial spot variance (solid line). See the text for details.*

2.4 Option Valuation Methodology

To price options using the models developed earlier, we need to know the mapping between the physical return shocks, z_t, and the risk-neutral return shocks. This mapping will allow us to use the physical asset return models developed earlier to simulate future stock prices from the risk-neutral distribution. These risk-adjusted prices can in turn be used to compute simulated option payoffs that once averaged and discounted using the risk-free rate, yield the current model option price.

Duan (1999) extended the normal case of Duan (1995) and derived a generalized local risk-neutral framework for option valuation in conditionally nonnormal GARCH models. As before, let z_{t+1} be iid normal under the physical measure and let z^*_{t+1} be iid normal under the risk-neutral measure. Define the $t+1$ mean-shift between the two measures by

$$\eta_{t+1} = z^*_{t+1} - z_{t+1}. \tag{2.27}$$

For a GED-distributed ζ_{t+1} shock, we can write the mapping as

$$\eta_{t+1} = z^*_{t+1} - \Phi^{-1}\left(G_\nu\left(\zeta_{t+1}\right)\right),$$

where Φ^{-1} is the standard normal inverse cumulative distribution function (cdf) so that $z_{t+1} = \Phi^{-1}(G_\nu(\zeta_{t+1}))$ is normally distributed. We can then rewrite the linear normal mapping in (2.27) as a nonlinear GED mapping given by

$$\zeta_{t+1} = G_\nu^{-1}\left(\Phi\left(z^*_{t+1} - \eta_{t+1}\right)\right). \tag{2.28}$$

The derivation of the risk-neutral model requires solving for η_{t+1}. This is done by setting the conditionally expected risk-neutral asset return in each period equal to the risk-free rate. In general, we can write

$$\exp(r) = \mathrm{E}_t^{\mathbb{Q}}\left[\exp\left\{\mathrm{E}_t\left[R_{t+1}\right] + \sqrt{h_{t+1}}G_\nu^{-1}\left(\Phi\left(z^*_{t+1} - \eta_{t+1}\right)\right)\right\}\right],$$

where the normal case obtains for $\nu = 2$.

2.4.1 The Affine-Normal Models

Recall from before that in the affine models, the conditional return mean is defined as

$$\mathrm{E}_t\left[R_{t+1}\right] = r + \lambda h_{t+1}. \tag{2.29}$$

In the normal case, of course, we have $G_2^{-1} = \Phi^{-1}$, so that the solution for η_{t+1} can be found as

$$\begin{aligned}
\exp(r) = \; & \mathrm{E}_t^{\mathbb{Q}}\left[\exp\left\{r + \lambda h_{t+1} + \sqrt{h_{t+1}}\left(z_{t+1}^{*} - \eta_{t+1}\right)\right\}\right] \\
\Leftrightarrow 1 = & \exp\left(\lambda h_{t+1}\right)\exp\left(-\eta_{t+1}\sqrt{h_{t+1}}\right)\exp\left(\frac{1}{2}h_{t+1}\right) \\
\Leftrightarrow \eta_{t+1} = & \left(\lambda + \frac{1}{2}\right)\sqrt{h_{t+1}}
\end{aligned}$$

so that

$$z_{t+1} = z_{t+1}^{*} - \left(\lambda + \frac{1}{2}\right)\sqrt{h_{t+1}}, \tag{2.30}$$

which corresponds to the mapping of Heston and Nandi (2000).

Future stock returns can now be simulated under the risk-neutral measure by substituting the shock transformation in (2.30) into the asset return models in (2.1) and (2.9). In the affine GARCH(1,1) model, we get

$$\begin{aligned}
R_{t+1}^{*} &= r - \frac{1}{2}h_{t+1} + \sqrt{h_{t+1}}z_{t+1}^{*}, \\
h_{t+1} &= w + b' h_t + a\left(z_t^{*} - c^{*}\sqrt{h_t}\right)^2,
\end{aligned} \tag{2.31}$$

where $z_{t+1}^{*} \sim N(0,1)$ and $c^{*} = c + \lambda + \frac{1}{2}$. Note how the structure of the expected return, via its impact on the mapping from z_{t+1} to z_{t+1}^{*}, ultimately provides a risk-neutral volatility dynamic similar to the physical one.

2.4.2 The Nonaffine-Normal Models

In these models, we have

$$\mathrm{E}_t\left[R_{t+1}\right] = r + \lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1} \tag{2.32}$$

so that the solution for η_{t+1} is found to be $\eta_{t+1} = \lambda$, which gives the mapping

$$z_{t+1} = z_{t+1}^{*} - \lambda, \tag{2.33}$$

which in turn corresponds to the mapping of Duan (1995).

Future stock returns can be simulated under the risk-neutral measure by substituting the shock transformation in (2.33) into the asset return models in (2.14) and

(2.22). We get

$$R_{t+1}^* = r - \frac{1}{2}h_{t+1} + \sqrt{h_{t+1}}z_{t+1}^*, \qquad (2.34)$$
$$h_{t+1} = w + b'h_t + ah_t\left(z_t^* - c^*\right)^2,$$

where $z_{t+1}^* \sim N(0,1)$ and $c^* = c + \lambda$. Note how once again, the structure of the expected return ultimately provides a nonaffine risk-neutral volatility dynamic similar to the nonaffine physical volatility dynamic.

2.4.3 The Affine GED Models

In the nonnormal case, an exact solution for η_{t+1} involves a prohibitively cumbersome numerical solution for η_{t+1} on every day and on every Monte Carlo path. Consequently, we develop the following approximation.

In the GED case, the parameter ν determines the degree of nonnormality in the cdf G_ν. When $\nu = 2$, we get normality, and when $\nu < 2$, we get fat tails. In the normal special case, we have $G_2^{-1}(\Phi(z)) = z$ for all z. Because the normal and GED are both symmetric, we know that $G_\nu^{-1}(\Phi(0)) = 0$ for all ν. We use this to suggest the linear approximation

$$G_\nu^{-1}\left(\Phi\left(z\right)\right) \approx b_\nu z,$$

where b_ν is easily found for a given value of ν by fitting $\zeta_i = G_\nu^{-1}(\Phi(z_i))$ to z_i for a wide grid of z_i values. This approximation is motivated by the fact that the probability integral transform, $\zeta = G_\nu^{-1}(\Phi(z))$, is very close to linear for ν values around the empirical estimates in Table 2.1.

With this approximation, we can write

$$\begin{aligned} 1 &= \exp\left(\lambda h_{t+1}\right) \mathrm{E}_t^{\mathbb{Q}}\left[\exp\left\{\sqrt{h_{t+1}}G^{-1}\left(\Phi\left(z_{t+1}^* - \eta_{t+1}\right)\right)\right\}\right] \\ &\approx \exp\left(\lambda h_{t+1}\right) \mathrm{E}_t^{\mathbb{Q}}\left[\exp\left\{\sqrt{h_{t+1}}b_\nu\left(z_{t+1}^* - \eta_{t+1}\right)\right\}\right] \\ &= \exp\left(\lambda h_{t+1}\right)\exp\left(-\eta_{t+1}\sqrt{h_{t+1}}b_\nu\right)\exp\left(\frac{1}{2}h_{t+1}b_\nu^2\right). \end{aligned}$$

Taking logs yields

$$0 = \lambda h_{t+1} - \eta_{t+1}\sqrt{h_{t+1}}b_\nu + \frac{1}{2}h_{t+1}b_\nu^2,$$

and solving for η_{t+1} yields

$$\eta_{t+1} = \left(\frac{\lambda}{b_\nu} + \frac{1}{2}b_\nu\right)\sqrt{h_{t+1}},$$

where, of course, the normal case obtains when $b_\nu = 1$.

The mapping between the physical GED and the risk-neutral normal shocks is now

$$\zeta_{t+1} = G_\nu^{-1}\left(\Phi\left(z_{t+1}^* - \left(\frac{\lambda}{b_\nu} + \frac{1}{2}b_\nu\right)\sqrt{h_{t+1}}\right)\right), \tag{2.35}$$

which can be substituted into the return dynamics for the AGARCH(1,1)-GED and AGARCH(C)-GED models to get the risk-neutral processes.

Note that whereas the linear approximation greatly facilitates computing η_{t+1} in the GED models, computing option prices is still more cumbersome with GED innovations because of the frequent inversion of the GED cumulative distribution function in (2.35).

2.4.4 The Nonaffine GED Model

Using the nonaffine return drift but the same approximation of $G_\nu^{-1}(\Phi(\cdot))$ as before, we can write

$$\begin{aligned}
1 &= \exp\left(\lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1}\right) \mathrm{E}_t^{\mathbb{Q}}\left[\exp\left\{\sqrt{h_{t+1}}G^{-1}\left(\Phi\left(z_{t+1}^* - \eta_{t+1}\right)\right)\right\}\right] \\
&\approx \exp\left(\lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1}\right) \mathrm{E}_t^{\mathbb{Q}}\left[\exp\left\{\sqrt{h_{t+1}}b_\nu\left(z_{t+1}^* - \eta_{t+1}\right)\right\}\right] \\
&= \exp\left(\lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1}\right) \exp\left(-\eta_{t+1}\sqrt{h_{t+1}}b_\nu\right) \exp\left(\frac{1}{2}h_{t+1}b_\nu^2\right).
\end{aligned}$$

Taking logs yields

$$0 = \lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1} - \eta_{t+1}\sqrt{h_{t+1}}b_\nu + \frac{1}{2}h_{t+1}b_\nu^2,$$

and solving for η_{t+1} yields

$$\eta_{t+1} = \frac{\lambda}{b_\nu} + \left(\frac{b_\nu}{2} - \frac{1}{2b_\nu}\right)\sqrt{h_{t+1}},$$

where, of course, the normal case again obtains when $b_\nu = 1$. The mapping between the shocks is then

$$\zeta_{t+1} = G_\nu^{-1}\left(\Phi\left(z_{t+1}^* - \frac{\lambda}{b_\nu} - \left(\frac{b_\nu}{2} - \frac{1}{2b_\nu}\right)\sqrt{h_{t+1}}\right)\right),$$

which can be substituted into the return dynamics for the NGARCH(1,1)-GED and NGARCH(C)-GED models to get the risk-neutral processes.

2.4.5 Monte Carlo Simulation

The European call option prices are computed via Monte Carlo, simulating the risk-neutral return process and computing the sample analog of the discounted risk neutral expectation. For a call option, $C_{t,T}$, quoted at the close of day t with maturity on day T and strike price X, we have

$$\begin{aligned} C_{t,T} &= \exp\left(-r(T-t)\right) E_t^*[Max(S_T - X, 0)] \\ &\approx \exp\left(-r(T-t)\right) \frac{1}{MC} \sum_{i=1}^{MC} \left[Max\left(S_t \exp\left(\sum_{\tau=1}^{T-t} R_{i,t+\tau}^* \right) - X, 0 \right) \right], \end{aligned}$$

where $R_{i,t+\tau}^*$ denotes future daily log returns simulated under the risk-neutral measure. The subscript i refers to the ith out of a total of MC simulated paths.

To compute the option prices numerically, we use 20 affine matrix scrambles of 5,000 Sobol sequences each, for a total of 100,000 simulated paths. We use the empirical martingale method of Duan and Simonato (1998) to increase numerical efficiency.

To assess the accuracy of our Monte Carlo implementation, we use the model of Heston and Nandi (2000), for which quasi-analytical option prices are available via Fourier inversion of the conditional characteristic function. We compare these analytical option prices with the Monte Carlo prices for various strike prices, maturities, and volatility levels. We use the parameterization of the Heston-Nandi variance process from Table 2.1.

Figure 2.4 reports the results. The figure presents results for two maturities (1 month and 3 months) and for three spot volatility levels [6.48%, 12.11%, and 24.00%, equal to the minimum, median, and maximum model volatility of the AGARCH(1, 1)-N between January 1990 and December 1995]. In all cases, the Monte Carlo prices, denoted by "+,", offer a very good approximation to the analytical prices, denoted by solid lines.

2.4.6 Option Valuation: Model Mechanics

Before turning to an empirical investigation of the various option valuation models, we analyze how the models differ in terms of option valuation by generating synthetic option prices for standardized moneyness, maturity, and volatility. This exercise

Figure 2.4: Accuracy of Monte Carlo Simulated Prices

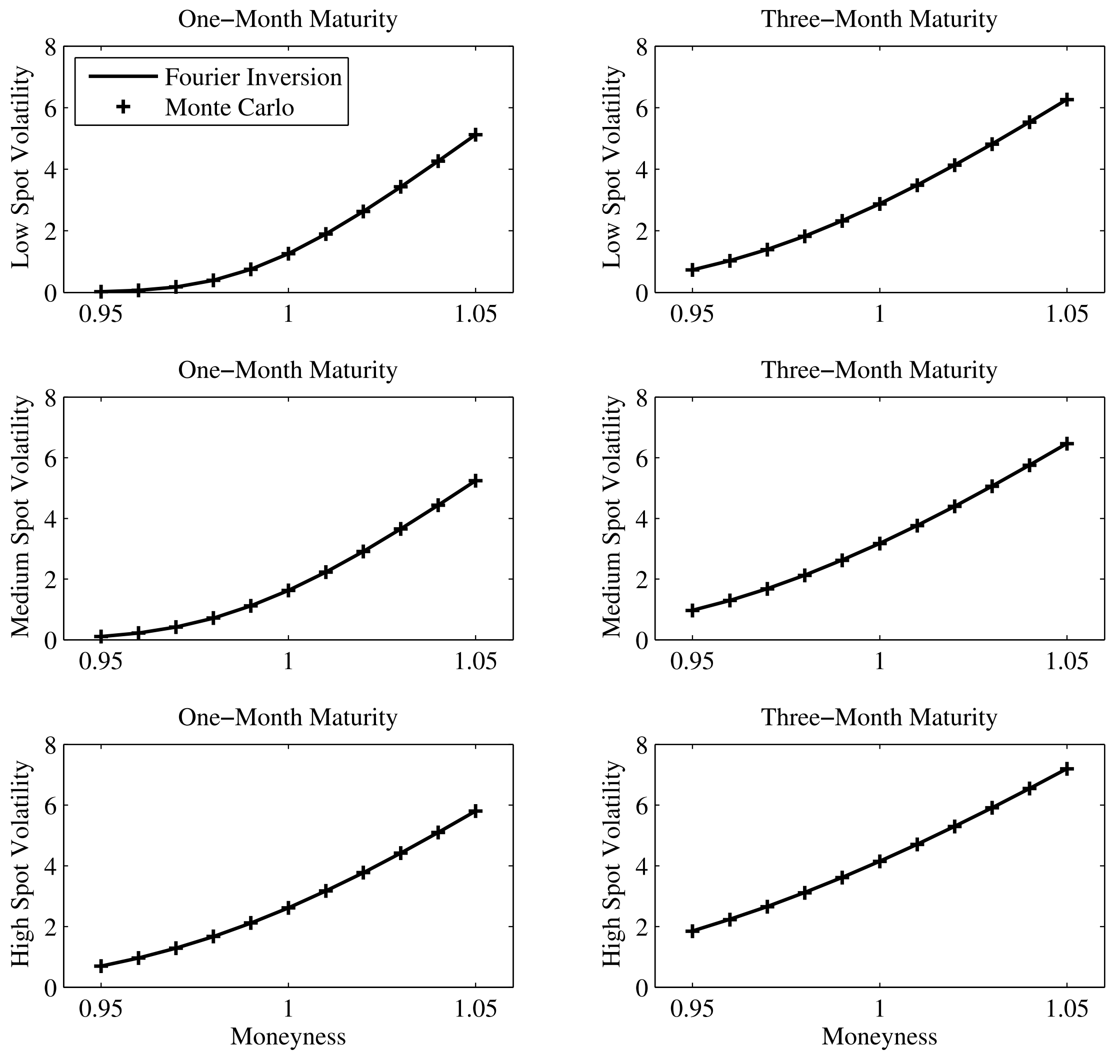

Using the AGARCH(1,1)-N model and the parameters from Table 2.1, we compare quasi-analytical option prices from Fourier inversion (solid) to the Monte Carlo prices ("+") for various strike prices, maturities, and volatility levels. The three spot volatility levels are 6.48%, 12.11%, and 24.00%, corresponding to the minimum, median, and maximum model volatility of the model between 1990 and 1995.

provides us with intuition on how the various models are able to match option prices, and indicates in which dimensions different models fall short. We generate synthetic prices using the parameter estimates in Table 2.1, which reflect model properties at the optimum.

Figure 2.5 presents weekly synthetic prices for the four models with normal innovations. The results are very similar for the GED models, and thus we do not report these here. To include multiple maturities in the same figure, we convert the prices to Black-Scholes implied volatilities. The light-gray band plots results for a 1-month option, and the dark-gray band plots results for a 1-year option. The thickness of each band indicates the range of option prices generated by each model across the moneyness of the options. Each week, we value an option with moneyness (S/X) equal to 0.95 and 1.05; the thickness of the band indicates the difference between these two prices. Option prices are monotonic in moneyness, and all other determinants of price are kept constant; thus the width of the band is a function of moneyness only. Therefore, Figure 2.5 allows us to comment on differences between the models as a function of moneyness and maturity, as well as differences in model price over time, which are related to the volatility level.

We draw the following conclusions. First, the thickness of the bands clearly indicates that for longer-maturity options, the nonaffine models generate much more price variation across moneyness than the affine models. Second, differences between the models are more pronounced for longer-maturity options, which is not surprising; however, the nonaffine models value the 1-month options somewhat differently than the affine models. For comparison, we indicate the time path of the VIX by a black line. For the nonaffine models, the time path of the band for the 1-month option follows the VIX much more closely.

Third, some models—most notably the AGARCH(1, 1)-N—generate very little variation in long-maturity prices over the 1990-1995 sample. The nonaffine models and component models perform very differently in this respect. For example, it is striking that the time paths for the long-maturity and short-maturity bands are much more correlated for the NGARCH(C)-N model than for the other models. Thus we conclude that models strongly differ in their pricing results under changing (volatility) conditions. The nonaffine models allow for more variation in prices as a

Figure 2.5: Synthetic Option Implied Volatilities across Time

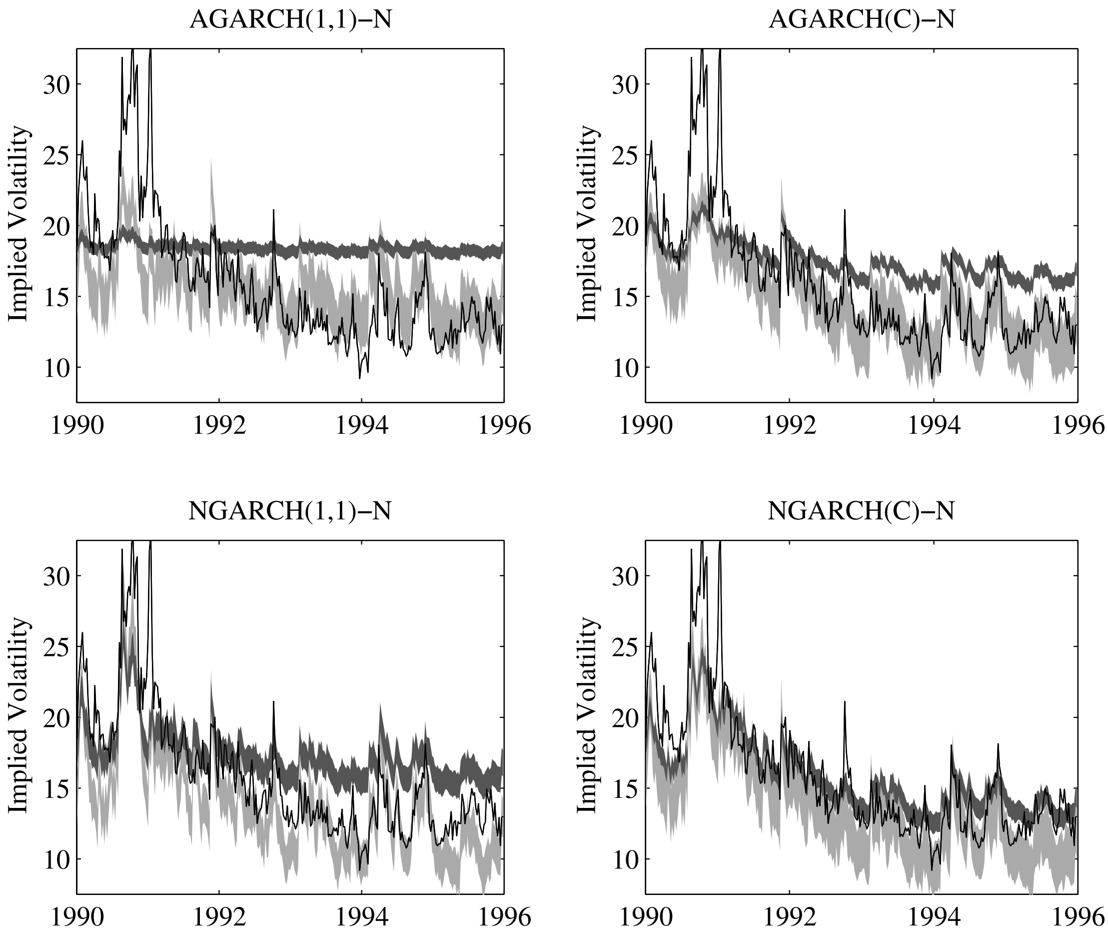

We plot weekly synthetic implied Black-Scholes volatilities for the four models with normal innovations. The light grey band plots the results for a 1-month option, and the dark grey band plots the results for a one-year option. The black line denotes the VIX index. Each week, we value an option with moneyness (S/X) equal to 0.95 *and* 1.05*; the thickness of the band indicates the difference between these two prices.*

function of volatility, such as, for example, higher prices in the more volatile 1990-1991 period. But obtaining prices for long-maturity options that are more correlated with those of shorter-maturity options requires component models.

Figure 2.6 provides further insight by presenting plots of synthetic 1-month option prices as a function of moneyness and volatility. To accentuate differences between models, we plot the differences between model prices and Black-Scholes prices. Because of space constraints, we provide results only for the nonaffine models with normal innovations.

Figure 2.6: Synthetic NGARCH-N Option Prices across Moneyness (S/X) and Volatility Levels

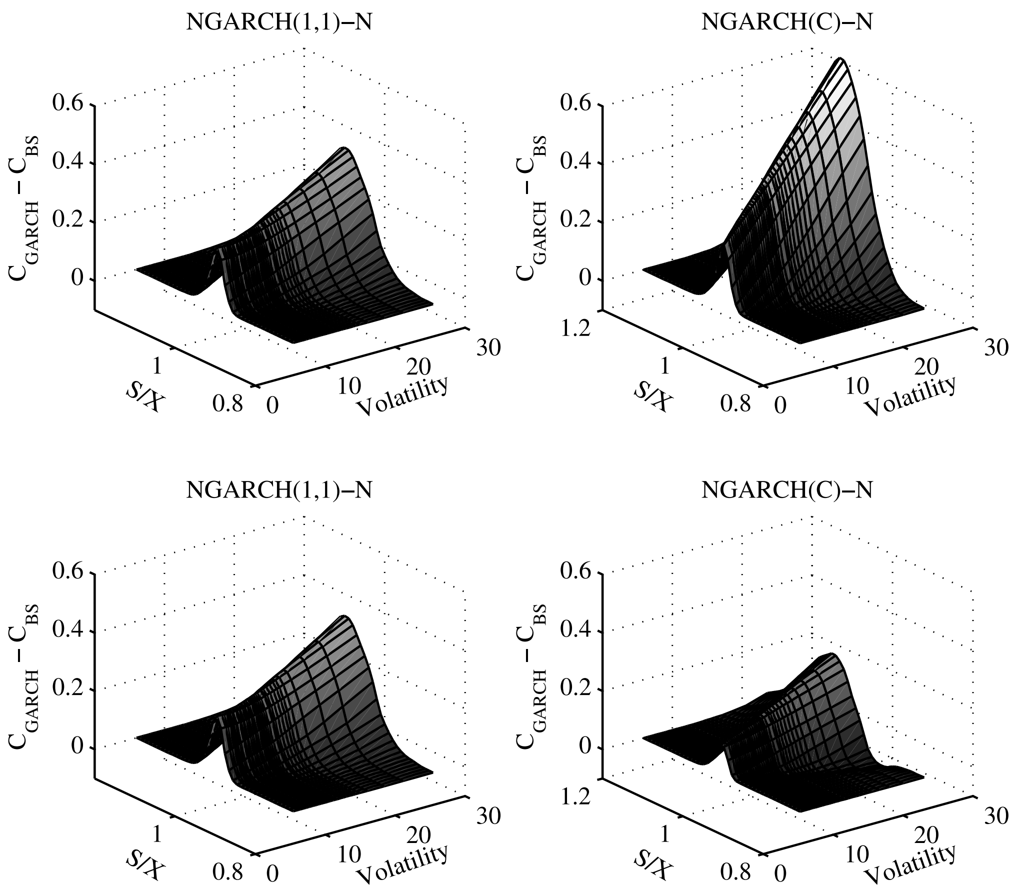

For the NGARCH normal models, we plot the synthetic model option price less the Black-Scholes price as a function of moneyness (S/X) and volatility for options with one month to maturity. For the component model, we set the proportion of the long-run variance factor, q_t, to the total variance, h_t, equal to the 90th percentile (top-right) and the 10th percentile (bottom-right) of its historical distribution.

While the volatility level on the horizontal axis in the figure fixes the spot variance h_t in the models, we still need to fix the long-run variance factor, q_t, in the component models. We proceed as follows. In the top right panel of Figure 2.6, we set q_t, so that the ratio q_t/h_t corresponds to the 90th percentile of the historical ratio, thus generating a relatively high long-run variance factor. In the bottom right panel of Figure 2.6, we set q_t, so that the ratio q_t/h_t corresponds to the 10th percentile of the historical ratio, generating a low long-run variance factor. Note again that the overall variance, h_t, is kept the same across models. Thus the values for the two GARCH(1,1) models in the left panels are identical. Finally, note that while option model prices increase as a function of volatility, of course the same need not apply to the model price net of the Black-Scholes price.

Figure 2.6 shows large differences between the GARCH(1,1) and the component models when the long-run variance factor is high. The nonaffine GARCH(1,1) model generates larger differences with Black-Scholes when volatility is high; however, this difference is dwarfed by the difference between the GARCH(1,1) and component models when the long-run variance component is high.

The bottom-right panel of Figure 2.6 shows the case where the long-run variance factor is low. Note that we have kept the overall volatility the same as in the top-right panel. When the long-run variance factor is relatively low, the component models generate prices that are even closer to Black-Scholes than those generated by the GARCH(1,1) models.

The overall conclusion from Figure 2.6 is that the component structure gives additional flexibility in generating option prices that vary not only as a function of the overall level of volatility, but also as a function of the composition of overall volatility into short-run and long-run components.

2.5 Option Valuation Empirics

We are now ready to use the eight models estimated in Section 2 and the transformation to risk-neutrality in Section 4 to assess the performance of the models for option valuation. In this section, we first introduce the options data, then use each of our eight models to price the option contracts and compare model and market prices for various maturities, strike prices, and sample years. We then report an economic analysis of the errors, and finally investigate alternative estimation strategies.

2.5.1 Option Data

We use 6 years of S&P 500 call option data covering the period 1990-1995. Starting from the raw data from the Berkeley option data base, we apply standard filters following Bakshi, Cao, and Chen (1997). We use only options with more than 7 days to maturity. We also use only Wednesday options data. Wednesday is the day of the week least likely to be a holiday, and is also less likely than other days, such as Monday and Friday, to be affected by day-of-the-week effects. If Wednesday is a holiday, then we use the next trading day. Using only Wednesday data allows us to study a fairly long time series, which is useful considering the highly persistent volatility processes.

Our sample comprises 21,752 options with a wide range of moneyness and maturity. The average overall price is $27.91. The data have been described in more detail by Christoffersen, Jacobs, Ornthanalai, and Wang (2008).

The implied Black-Scholes volatilities display strong evidence of a "smirk" or "skew" across strikes, with higher implied volatility for in-the-money calls than at-the-money calls, and this holds for all maturities. This empirical regularity illustrates that the Black-Scholes option valuation formula, which assumes a constant per period volatility across time, maturity, and strike prices, will result in systematic pricing errors, which motivates the use of GARCH models for option valuation.

2.5.2 Overall Option Valuation Results

When calculating option prices according to the eight GARCH models, we use the ML estimation parameters in Table 2.1 transformed to the risk-neutral measure. We use these risk-neutral parameters, along with the conditional variance paths from Figure 2.1 as inputs into the option pricing formula. In the case of the nonaffine and/or nonnormal models, Monte Carlo simulation is required to calculate the price. In the case of the normal affine models, numerical integration solutions exist; however, to ensure that the results are not driven by the numerical pricing technique, we use Monte Carlo simulation using the same set of random numbers for all models.

The overall RMSEs for the eight GARCH models are reported in the next to last row of Table 2.2. The root mean squared error (RMSE) is computed as

$$RMSE = \sqrt{\frac{1}{N}\sum_{i,t}\left(C_{i,t}^{Market} - C_{i,t}^{Model}\right)^2},$$

where the summation is over contract i observed on day t and where N is equal to 21,752, the total number of option contracts in the sample. The last row in Table 2.2 normalizes the $RMSE$ by dividing it by the average call price in the sample.

Note first that the best overall model (i.e., that with the lowest $RMSE$) is the NGARCH(C)-N, with an $RMSE$ of 1.38, followed closely by the NGARCH(C)-GED, with an $RMSE$ of 1.43. The two nonaffine GARCH(1,1) models also perform relatively well, with $RMSE$s of 1.46 in the GED case and 1.59 in the normal case. The affine models as a group perform worse than the nonaffine models. The $RMSE$ of the AGARCH(C)-GED model is 1.74, and that of the AGARCH(C)-N model is 1.81. The two affine GARCH(1, 1) models perform the worst, with an RMSE of 2.68 in the GED case and 2.69 in the normal case.

These overall $RMSE$ results allow us to make comparisons in three dimensions: affine versus nonaffine variance dynamics, GED versus normal shocks, and GARCH(1,1) versus component variance models.

First, as noted earlier, we see that nonaffine models perform much better than affine models. This is true both for GARCH(1,1) and component models and for GED and normal shocks. Thus our results confirm and extend those of Hsieh and Ritchken (2005) who compared an affine model and a nonaffine model in the GARCH(1,1) case with normal shocks.

Second, we see that GED models perform only marginally better than normal models. The greatest improvement is in the NGARCH(1,1) case, where the $RMSE$ drops from 1.59 to 1.46 when going from GED to normal shocks. In the other pairwise comparisons, the difference between the GED and the normal $RMSE$ is around 5 cents.

Third, we see that the component structure offers large improvements in fit for the affine class of models but more modest improvements in the nonaffine class of models. For the affine models, the $RMSE$ drops from 2.69 to 1.81 for normal shocks and from 2.68 to 1.74 for GED shocks, whereas for the nonaffine models, the $RMSE$ drops from 1.59 to 1.38 for normal shocks and from 1.46 to 1.43 for GED shocks.

Recall now the main findings in terms of daily return log-likelihood values in Table 2.1. In all cases, the component model has a significantly larger log-likelihood than the nested GARCH(1,1). Comparing the GED models with their normal counterparts reveals that the GED-based models have significantly larger log-likelihood values. The log-likelihood values also suggest that the nonaffine GARCH models are strongly preferred over the affine GARCH models. The option-based results support the return-based improvement of nonaffine models over affine models. They also support the component model improvements over GARCH(1, 1) for affine models, but somewhat less so for nonaffine models. While the GED offers drastic improvements in the return-based likelihood analysis, the improvements offered in option valuation are much more modest. The normal GARCH models may offer sufficient nonnormality in the multiday distribution, or, alternatively, the GED specification may not be adequate for the purpose of option valuation.

2.5.3 Results by Moneyness, Maturity, and Volatility Level

Figure 2.7 provides more intuition for the models' performance by plotting the $RMSE$ and bias results as a function of moneyness (i.e., index value over strike value), maturity, and volatility level. Note that the scales differ across panels in Figure 2.7, to focus on the differences between models.

The plots in the top row of Figure 2.7 show $RMSE$ and bias for each model using six moneyness bins. $RMSE$ is presented in the left panel, and bias is shown in the right panel. Bias is defined as average market price minus model price. In each panel, the solid line represents the model with normal innovations and the dashed line represents the model with GED innovations. The bias plots indicate that model prices are generally larger than market prices, especially for out-of-the money call options. An important observation is that while the GED models outperform the models with normal innovations, the differences in performance are small, and for many models are almost nonexistent in most of the bins. This confirms the results presented in Section 2.5.2. But the GED models slightly outperform the models with normal innovations for out-of-the-money options, where nonnormality has more impact. Furthermore, while affine models perform similar to nonaffine models for in-the-money options, nonaffine models perform much better for out-of-

Figure 2.7: Option $RMSE$ and Bias against Moneyness, Maturity, and VIX Level

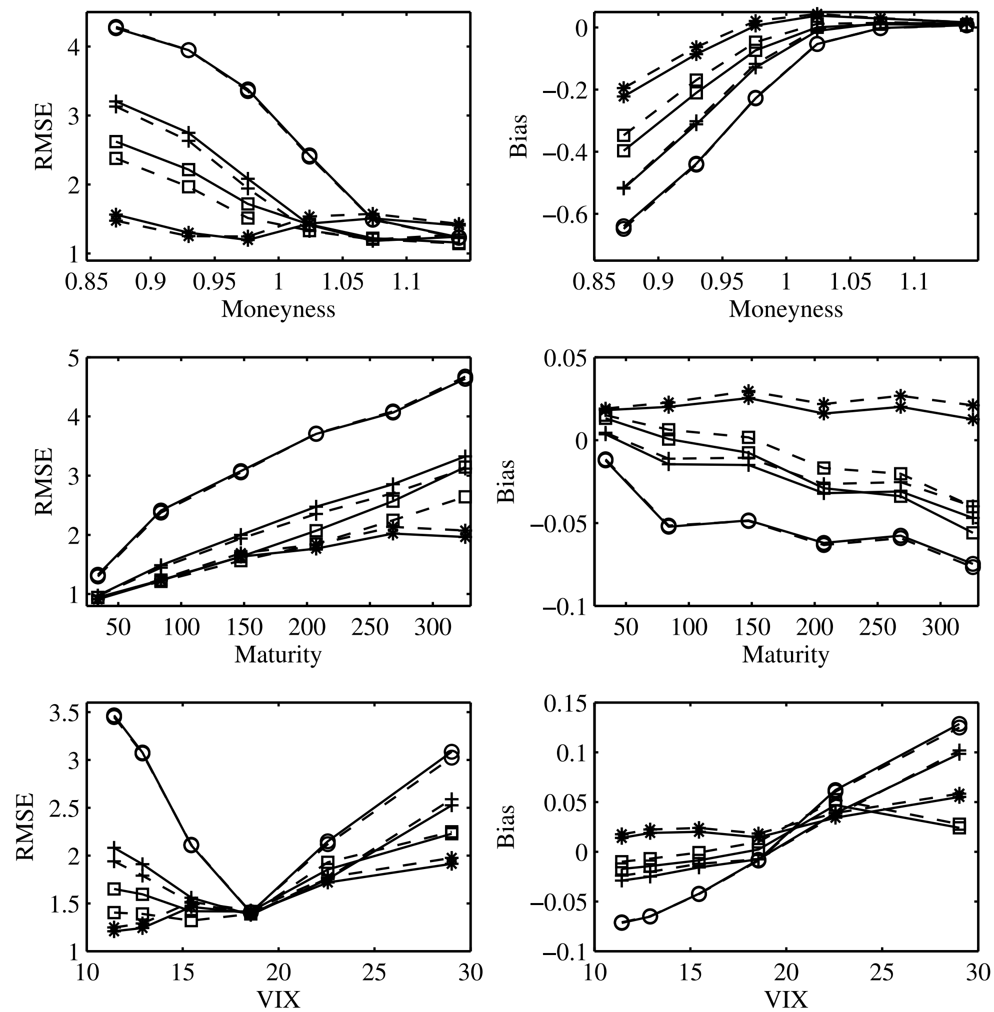

We plot RMSE (left column) and bias (right column) for each model using six moneyness bins (top row), maturity bins (middle row) and VIX levels (bottom row). Normal models are solid and GED models are dashed lines. AGARCH(1,1) is marked with "○", NGARCH(1,1) with "□", AGARCH(C) with "+" and NGARCH(C) with "".*

the money and at-the-money options. The affine models perform particularly poorly for out-of-the money options, perhaps not providing sufficient nonnormality at the relevant horizons. Differences across models also are generally larger for out-of-the-money options, driving the overall $RMSE$ result. Finally, the nonaffine component models significantly outperform the NGARCH(1, 1) models, suggesting that in the moneyness dimension, both the nonaffine structure and the component structure are needed.

The plots in the middle row of Figure 2.7 strongly suggest that both components and a nonaffine structure are also needed in the maturity dimension. The plots show $RMSE$ and bias for each model using six maturity bins. The nonaffine component models significantly outperform the affine component models and the nonaffine GARCH(1,1) models, suggesting that the nonaffine structure and the component structure by themselves are not sufficient to capture the richness of the data. The other important conclusion is, of course, that differences in model performance appear mainly at long maturities. The affine component models do well for the shortest maturities, but the nonaffine component and GARCH(1,1) models perform better than the affine models at longer maturities. The relative lack of flexibility in the longer-term variance dynamics shown in Figure 2.3 seems to hurt the affine models in the valuation of long-maturity options.

Finally, the plots in the bottom row of Figure 2.7 show $RMSE$ and bias for each model using eight volatility bins. To ensure meaningful comparisons, we define the volatility on a given day as the level of the CBOE's VIX volatility index. These plots illustrate the problems with the AGARCH(1,1) models even more starkly. In particular, the $RMSE$ plots indicate that these models perform similarly to the other models only when the VIX index is near its average over the sample period (i.e., 15.98%). Interestingly, the differences between the NGARCH(C) models on the one hand and the AGARCH(C) and NGARCH(1,1) models on the other hand seem to be smaller than those in the moneyness and maturity dimensions depicted in the top and middle rows. Most notably, even with the nonaffine component models, a substantial bias clearly remains at high volatility levels.

2.5.4 Economic Assessment of Option Valuation Performance

We now turn to a more detailed analysis of the option valuation performance of the models over time. Toward this end, we regress the weekly option bias (defined as the weekly average market price less the weekly average option price) on some key economic variables: the VIX index, the weekly S&P 500 return, the weekly crude

Brent Oil price, the 3-month T-bill rate, the credit spread (defined as the yield on corporate bonds rated Baa less the yield on Aaa bonds as rated by Moody's), and the term spread, defined as the difference between the yield on 10-year T-bond and the 3-month T-bill rate. We also include as regressors the weekly average moneyness, defined as index value over strike price (S/X), and the weekly average maturity in years. For these variables, averages are taken each week across the option contracts observed on the Wednesday of that week. The final regressor is the model-specific variance forecast for each model, defined as the average of the model's variance forecasts for the next month, as implied by the parameters given in Table 2.1.

Table 2.3: Regressing Weekly Bias on Economic Variables

	AGARCH-N		**NGARCH-N**		**AGARCH-GED**		**NGARCH-GED**	
	GARCH(1,1)	Components	GARCH(1,1)	Components	GARCH(1,1)	Components	GARCH(1,1)	Components
Avr weekly bias	-0.9684	-0.3031	-0.1260	0.5500	-0.9745	-0.2186	0.0700	0.6467
Regressor	**t-Statistics using White's Robust Standard Errors**							
Constant	**-3.983**	**-3.247**	-1.497	-1.537	**-3.970**	**-3.117**	-1.581	-1.237
VIX	**10.559**	**10.345**	**9.520**	**10.978**	**10.326**	**10.947**	**9.511**	**11.142**
S&P500 weekly return	-1.269	**-2.215**	-1.734	**-4.023**	-1.250	**-2.870**	-1.755	**-4.011**
Oil price	1.651	**3.648**	**6.077**	**5.745**	1.725	**3.494**	**6.281**	**5.873**
3-month T-bill rate	**2.993**	0.731	0.309	**-3.247**	**3.008**	0.322	0.054	**-3.408**
Credit spread	**7.548**	**7.233**	**4.724**	1.646	**7.564**	**6.794**	**4.553**	0.887
Term spread	**3.376**	0.790	0.076	**-2.748**	**3.292**	0.543	-0.261	-2.772
Avr moneyness (S/X)	**2.049**	1.766	0.115	0.685	**2.053**	1.717	0.238	0.402
Avr maturity (YTM)	-0.356	0.580	-1.844	**3.233**	-0.562	1.308	-0.914	**3.721**
Model variance forecast	**-5.950**	**-12.011**	**-15.830**	**-14.017**	**-5.787**	**-12.911**	**-16.046**	**-14.194**
Regression R-squared	0.8881	0.7582	0.7405	0.6124	0.8856	0.7432	0.7298	0.6112

For each model we regress the weekly bias on a constant and various weekly economic variables, as well as the average weekly moneyness and maturity of the options in the sample. We also regress on the model-specific variance forecast. We report the t-statistic for each regressor using White's robust standard errors. Numbers in bold are larger than two in absolute value. The top row shows the average weekly bias and the bottom row reports the R-squared regression fit.

Table 2.3 reports the results of the regressions of weekly option bias on the economic variables. The top row reports the average weekly bias (across the whole sample) for each model. t-statistics for each regressor are given, with the standard deviation for each regressor computed using White's robust variance matrix. Any t-statistic exceeding 2 in absolute value is in bold type. The bottom row gives the regression fit via the R-squared statistic. Most importantly, note that the R-squared values are quite high, suggesting that for these models, misspecification can be detected quite easily by analyzing model bias. Note that the explanatory power is particularly high in the affine GARCH(1,1) models and lowest in the nonaffine component models.

The VIX is positive and significant for all models, and the variance forecast is negative and significant. The coefficient on the S&P 500 return is estimated to be negative for all models (the opposite sign of the VIX), but it is not always significant. The increase in the short-term rate, as well as the decline in the term and credit spreads, seem capable of explaining some of the upward bias in the affine model prices seen in the second half of the sample.

In summary, this specification analysis highlights some interesting differences between the models and provides insight into the models' strengths and weaknesses. These findings also suggest that building option pricing models with variance dynamics driven by key economic variables, such as credit spreads and oil prices, is a viable avenue for future research.

2.5.5 Alternative Estimation Strategies

So far, we have estimated the GARCH models on daily returns only and then used them for option valuation without letting the model parameters be driven in any way by the observed option prices. We now want to check the robustness of our results in two ways. First, because the return sample period (1962-2001) used here overlaps with the option sample period (1990-1995), we shorten the return sample period to end in 1989, just before the start of the option sample period. Second, we use options to estimate the weekly spot variance, h_t, minimizing the weekly option $RMSE$ while keeping the model parameters fixed at the ML estimate values estimated on the 1962-1989 sample of daily index returns. These results are presented in Table 2.4 and Figure 2.8.

Panel A of Table 2.4 summarizes the valuation results using the ML estimation parameters from 1962-2001. The $RMSE$ in the top row of Table 2.4 is simply repeated from the next to last row of Table 2.2. The second row of Table 2.4 reports the overall bias, which is close to 0 for all models. The third row reports the average spot volatility across the 313 option sample days. Because of the variance targeting used in ML estimation, these average volatilities are quite similar across models. Finally, the fourth row reports the standard deviation of the 313 spot variances. In keeping with the results in Figure 2.2, the volatility of variance is highest in the nonaffine models.

Table 2.4: Alternative Estimation Strategies

	AGARCH-N		**NGARCH-N**		**AGARCH-GED**		**NGARCH-GED**	
	GARCH(1,1)	Components	GARCH(1,1)	Components	GARCH(1,1)	Components	GARCH(1,1)	Components
Panel A: Results from MLE on 1962-2001 Sample								
Option RMSE, 90-95	2.6927	1.8138	1.5875	1.3820	2.6819	1.7404	1.4599	1.4270
Option bias, 90-95	-0.0423	-0.0149	-0.0077	0.0193	-0.0425	-0.0114	-0.0004	0.0229
Ann avr vol % 90-95	12.367	11.884	11.293	11.153	12.398	11.907	11.292	11.157
Ann vol of var % 90-95	0.7904	0.8763	1.0064	0.9758	0.7819	0.8587	1.0081	0.9753
Panel B: Results from MLE on 1962-1989 Sample								
Option RMSE, 90-95	2.4750	1.6805	1.4347	1.4782	2.4163	1.6438	1.4021	1.5179
Option Bias, 90-95	-0.0314	-0.0060	0.0043	0.0268	-0.0296	-0.0046	0.0095	0.0296
Ann avr vol % 90-95	12.442	11.871	11.218	11.086	12.428	11.904	11.235	11.107
Ann vol of var % 90-95	0.7485	0.8481	1.0122	0.9932	0.7283	0.8324	1.0087	0.9876
Panel C: Results from NLS Estimation of GARCH Spot Variance								
Option RMSE, 90-95	2.0362	1.3555	1.0416	1.0076	1.9490	1.3731	1.0273	1.0143
Option bias, 90-95	-0.0174	0.0107	0.0058	0.0061	-0.0137	0.0118	0.0054	0.0055
Ann avr vol % 90-95	7.097	10.255	11.306	12.050	7.275	10.500	11.715	12.261
Ann vol of var % 90-95	2.5754	1.3651	1.1766	0.9814	2.4127	1.3083	1.2098	1.0068
RMSE of ann vol %	9.3485	3.7099	2.3055	1.6955	9.0147	3.4056	2.3297	1.7829
Option RMSE 90-92	1.3531	1.1555	0.8042	0.9170	1.3226	1.2479	0.8390	0.9249
Option RMSE 93-95	2.4156	1.4845	1.1852	1.0689	2.3008	1.4577	1.1453	1.0748
Option bias 90-92	0.1499	0.1772	0.0684	0.0987	0.1611	0.2336	0.0544	0.0856
Option bias 93-95	-0.9482	0.3863	0.2317	0.2225	-0.7770	0.3982	0.2232	0.2027

Using the GARCH MLE paramaters from Table 1, Panel A reports the overall option RMSE and bias, as well as the average model volatility on the option data days and the standard deviation of the model variance path on the option data days. Panel B reports the same statistics but using GARCH parameters estimated on the 1962-1989 sample instead. In Panel C, we use NLS to estimate the RMSE optimal spot variance each week but rely on the GARCH MLE parameters from the 1962-2001 sample of returns. We also report the RMSE distance between the annualized MLE volatility path and the NLS optimal volatilities, as well as the option RMSE and bias split up into the first and second half of the sample.

Panel B of Table 2.4 reports the same set of results using ML estimates of the GARCH parameters from daily returns from 1962 through 1989 rather than through 2001. Note that the option $RMSE$ often is lower when parameters are estimated on returns observed through 1989 than when estimated on returns observed through 2001. Compared with the 2001 estimates, the 1989 estimates result in higher option $RMSE$s only in the two nonaffine component models. Whereas the evidence in favor of nonaffine component models weakens, the other overall conclusions from the 2001 estimates remain intact. The poorer option valuation performance of the nonaffine component models when using the shorter sample could be driven by the fact that the component models require a longer return sample to properly identify the components.

Panel C of Table 2.4 uses the GARCH parameters from ML estimation on returns up through 1989, but estimates the GARCH spot variance, h_t, each week by minimizing that week's option $RMSE$ using a nonlinear least squares (NLS) technique. Comparing the pure MLE $RMSE$s in Panel B with the hybrid $RMSE$s in panel C shows that the reduction in $RMSE$ is dramatic in all models. The four nonaffine models now all have an $RMSE$ of around \$1, compared with around \$1.36 for the two affine component models and around \$2 for the GARCH(1,1) models. The overall ranking of models from the pure ML estimation analysis remains largely intact.

Figure 2.8 elaborates further on this finding by plotting the weekly spot volatility from NLS (dots), along with the corresponding ML estimation-based spot volatilities (solid lines). The differences between the ML estimation-based and NLS-based results in Figure 2.8 are quite striking. In the affine GARCH(1,1) models, the NLS optimizer forces the spot volatility to zero in much of the second half of the sample to decrease the overpricing.

Table 2.4 also reports the $RMSE$ between the annualized $RMSE$-optimal volatility and its ML-optimal counterpart. These numbers confirm the visual impression from Figure 2.8 and also corroborate some of our earlier findings. First, the nonaffine models have a much closer correspondence between option-implied and purely return-generated spot volatility. Second, the component structure reduces the volatility $RMSE$ by well over 50% in the affine models and also quite substantially in the

Figure 2.8: $RMSE$-Optimal Spot Volatility Estimated Weekly by NLS

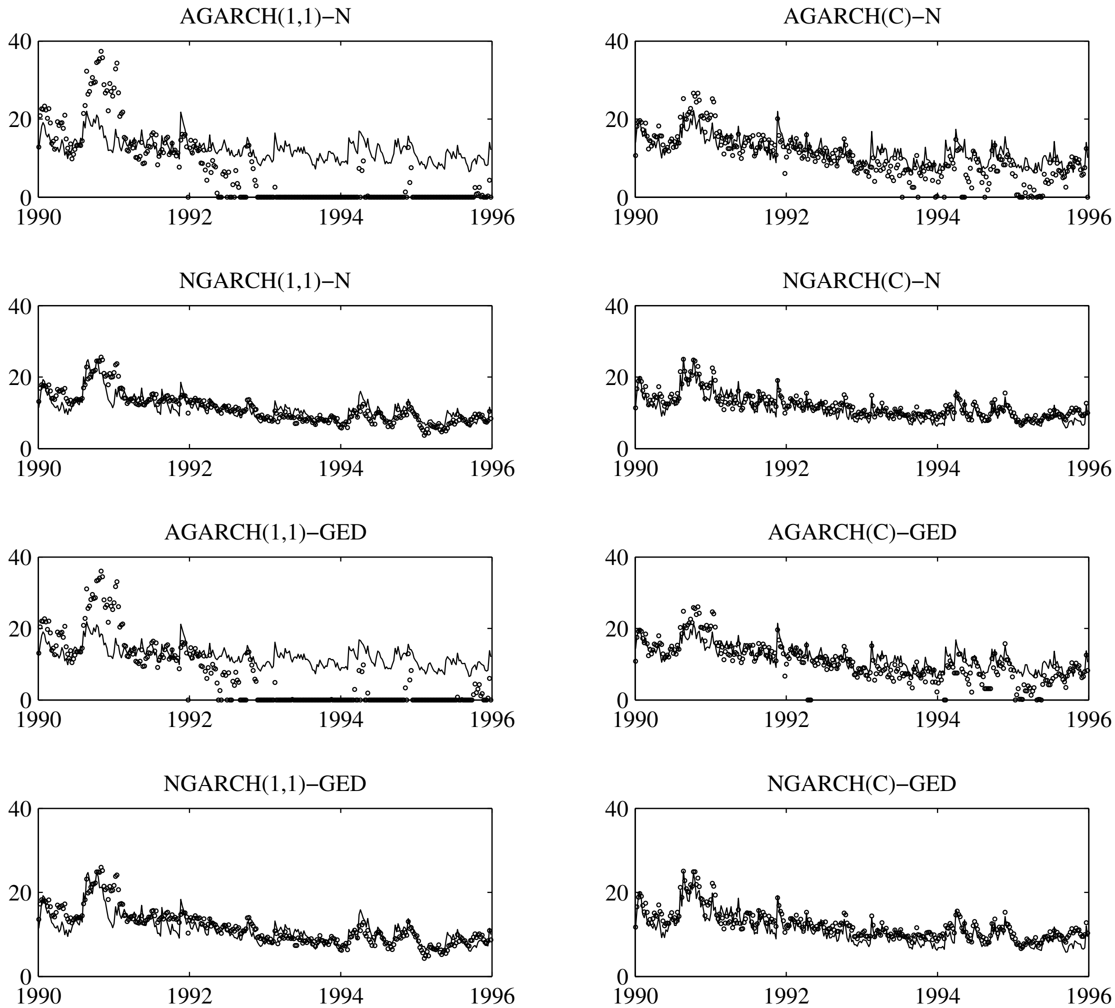

The dots show the weekly $RMSE$-optimal GARCH annualized spot volatilities, $100 * \sqrt{252h_t}$, *estimated by NLS. The annualized ML estimation-optimal volatility path is show in solid lines.*

nonaffine models. Third, the GED shocks do not have much effect when judged by this metric either.

In the last four rows of Table 2.4, we split the 1990-1995 sample into two halves and report the option $RMSE$ and bias for each half, using the NLS estimation results. Note that the AGARCH(1, 1) models are characterized by large negative bias in the second half of the sample, which leads to large $RMSE$s. Thus the overpricing persists even when the NLS optimization drives the spot variances to zero in the second half of the sample.

Finally, it must be emphasized that the smaller NLS bias cannot be interpreted as suggesting the superiority of this type of exercise. The MLE setup clearly provides more challenges for any model, thereby automatically leading to a higher bias.

2.6 Conclusion and Directions for Future Work

We have assessed the ability of eight different GARCH models to fit daily return dynamics and their ability to match market prices of options. First, we considered component models versus GARCH(1,1) models. Like Engle and Lee (1999), we found strong evidence in favor of component models from the standpoint of modeling daily return dynamics. When using option prices to assess the models, we also found strong evidence for the component structure in the affine GARCH models, but less so in the nonaffine models. Second, we considered nonaffine versus affine GARCH models. We compared the affine GARCH(1, 1) model with the nonaffine NGARCH(1, 1) model of Hsieh and Ritchken (2005), who found strong support for the nonaffine specification. Our results support their findings, and we also found that the nonaffine models outperform affine models when allowing for component structures and nonnormal shocks. Third, we considered conditionally normal versus conditionally nonnormal models. We found that assuming GED shocks for the daily asset returns greatly improves the fit of the all models to daily returns, but the improvement in option valuation is much less evident.

The empirical results suggest some viable directions for future research. First, it remains to be seen whether the differences in performance between models are confirmed when using model parameters estimated from option prices, or when using an integrated analysis that uses option prices as well as underlying returns (see Bates 2000; Chernov and Ghysels 2000; Eraker 2004). The analysis in Table 2.4 suggests that the relative performance of the models is comparable when the spot volatility is estimated from options rather than filtered from returns.

Second, it would be interesting to expand the analysis of nonnormal shocks to a wider class of distributions. Toward this end, Christoffersen, Heston, and Jacobs (2006) developed an inverse-Gaussian GARCH model, Duan, Ritchken, and Sun (2006) suggested augmenting GARCH models with jumps, and Lehnert (2003) applied an EGARCH model with skewed GED shocks.

Third, we have restricted attention to European-style options on the S&P 500 index. It would be interesting to apply the GARCH modeling framework to some of the many American-style contracts traded in the derivatives markets. Ritchken and Trevor (1999) and Stentoft (2005) have provided fast numerical techniques for GARCH option valuation with early exercise.

Finally, it would be interesting to compare the range of discrete-time GARCH models considered here with continuous-time stochastic volatility models. Bates (1996), Bakshi, Cao, and Chen (1997), and Eraker (2004) have studied stochastic volatility models with jumps; Taylor and Xu (1994) have studied multifactor stochastic volatility models; and Bates (2000) has analyzed models with Poisson jumps and multiple volatility factors. Comparing GARCH and SV models for the purpose of option valuation may provide more insight into the strengths and weaknesses of the various models.

3 Business Conditions, Market Volatility, and Option Prices

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 MIDAS structure to account for changes in the recently introduced 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.

3.1 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.[*] These authors find that this fundamental volatility process is significantly related to such factors as inflation and industrial production growth.[†] 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 (2009; henceforth ADS). Using this index within a mixed data sampling (MIDAS) model,[‡] 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.[§]

However, two-factor models mostly rely on two latent, autoregressive volatility factors. Whereas the actual drivers of these processes are usually left unidentified,

---

[*] 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.

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

[‡] 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).

[§] 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).

our fundamental volatility process acknowledges that macroeconomic determinants do impact 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.*

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.

*Our results are thus consistent with those of, for instance, Christoffersen, Jacobs, Ornthanalai, and Wang (2008), and Christoffersen, Dorion, Jacobs, and Wang (2008). 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 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 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.*

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 3.2 presents the MacroHV-MIDAS model. Section 3.3 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 3.4 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 3.5 refines the results of Section 3.4 and studies the time-series dynamics of the volatility premium implied by the MacroHV-MIDAS model. Finally, Section 3.6 concludes.

*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.

3.2 The MacroHV-MIDAS Model

3.2.1 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. 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

$$R_{t+1} = \mu_{t+1} + \sqrt{h_{t+1}}\varepsilon_{t+1} , \qquad \varepsilon_{t+1} \overset{\mathbb{P}}{\sim} \mathcal{N}(0,1) , \tag{3.1}$$

$$h_{t+1} = f\,(\,\cdot\mid\Theta,\,\mathcal{F}_t\,)\ , \tag{3.2}$$

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 (2008; henceforth EGS) 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:

$$h_{t+1} = g_{t+1}\tau_{t+1} , \tag{3.3}$$

$$g_{t+1} = (1-\alpha-\beta) + \alpha g_t \varepsilon_t^2 + \beta g_t , \tag{3.4}$$

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,

$$HV_t = \sum_{n=0}^{N-1} R_{t-n}^2 \,, \tag{3.5}$$

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

$$\log(\tau_{t+1}) = m + \theta_{hv} \sum_{k=0}^{K-1} \phi_k(w_{hv})\, HV_{t-k} \,, \tag{3.6}$$

where historical volatilities are computed on a daily basis using the N last daily returns observed on the market.† 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.‡ While w can be estimated through maximum likelihood, the number of observations used, N in Equation (3.5), just as the number of lags considered, K in Equation (3.6), 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

*EGS refer to the estimate of Equation (3.5) 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 (3.5). 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).

†Equation (3.6) 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.

‡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.

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.

3.2.2 MacroHV-MIDAS: Recent Volatility Levels and Business Conditions

This paper studies the following model:

$$R_{t+1} = r + \lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1} + \sqrt{h_{t+1}}\varepsilon_{t+1} \ , \tag{3.7}$$

$$h_{t+1} = g_{t+1}\tau_{t+1} \ , \tag{3.8}$$

$$g_{t+1} = (1 - \alpha(1+\gamma^2) - \beta) + \alpha g_t(\varepsilon_t - \gamma)^2 + \beta g_t \ , \tag{3.9}$$

$$\log(\tau_{t+1}) = m + \theta_{hv}\sum_{k=0}^{K-1}\phi_k(w_{hv})\,HV_{t-k} + \theta_m\sum_{k=0}^{K-1}\phi_k(w_m)\,\Delta x_{t-k} \ , \tag{3.10}$$

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 3.3, 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 (3.10) 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 (3.3)–(3.6) 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 (3.9) 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.* 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 around a constant expected variance level parameterized by ω,

$$h_{t+1} = \omega + \alpha h_t(\varepsilon_t - \gamma)^2 + \beta h_t \,. \tag{3.11}$$

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 (2008; henceforth CDJW) 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.

3.3 Estimating the Model using the ADS Business Conditions Index

The fundamental variance process of the MacroHV-MIDAS model, defined in Equation (3.10), requires a daily measure of business conditions. Before discussing the estimation of the model, this section presents the ADS Business Conditions Index

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

and argues that it is well suited to fulfill the role implied by our characterization of the fundamental variance process.

3.3.1 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 (2009; henceforth ADS). 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.*

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

$$x_t = \varphi x_{t-1} + v_t\,, \qquad v_t \overset{\mathbb{P}}{\sim} \mathcal{N}(0,1)\,, \tag{3.12}$$

that is referred to as the business conditions index. The average value of the ADS index, $\mathrm{E}\left[\,x_t\,\right]$, 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 3.1 reports summary statistics on the index, and the lower right panel of Figure 3.1 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 (3.12) for the x_t business conditions index; in particular, the homoskedasticity assumption is neces-

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

sary for identification. The second column of Table 3.1, 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 (3.12),

$$v_t = \sqrt{h_t^x} u_t \,, \qquad u_t \overset{\mathbb{P}}{\sim} \mathcal{N}(0,1) \,, \tag{3.13}$$

$$h_t^x = \omega_x + \alpha_x v_{t-1}^2 + \beta_x h_{t-1}^x \,. \tag{3.14}$$

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 (3.12)–(3.14) is, thus, preferable for forecasting purposes.

Table 3.1: Descriptive Statistics on the Business Conditions Index

			GARCH	
	Index: $\mathbf{x_t}$	**Residuals: $\mathbf{v_t}$**	**Residuals: $\mathbf{u_t}$**	**Variances: $\mathbf{h_t} \times$ 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 (3.12), (3.13) *and* (3.14) *:*

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

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.

3.3.2 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 3.1 reports some key macroeconomic series throughout the time period considered in this paper.

Figure 3.1: Macroeconomic Series

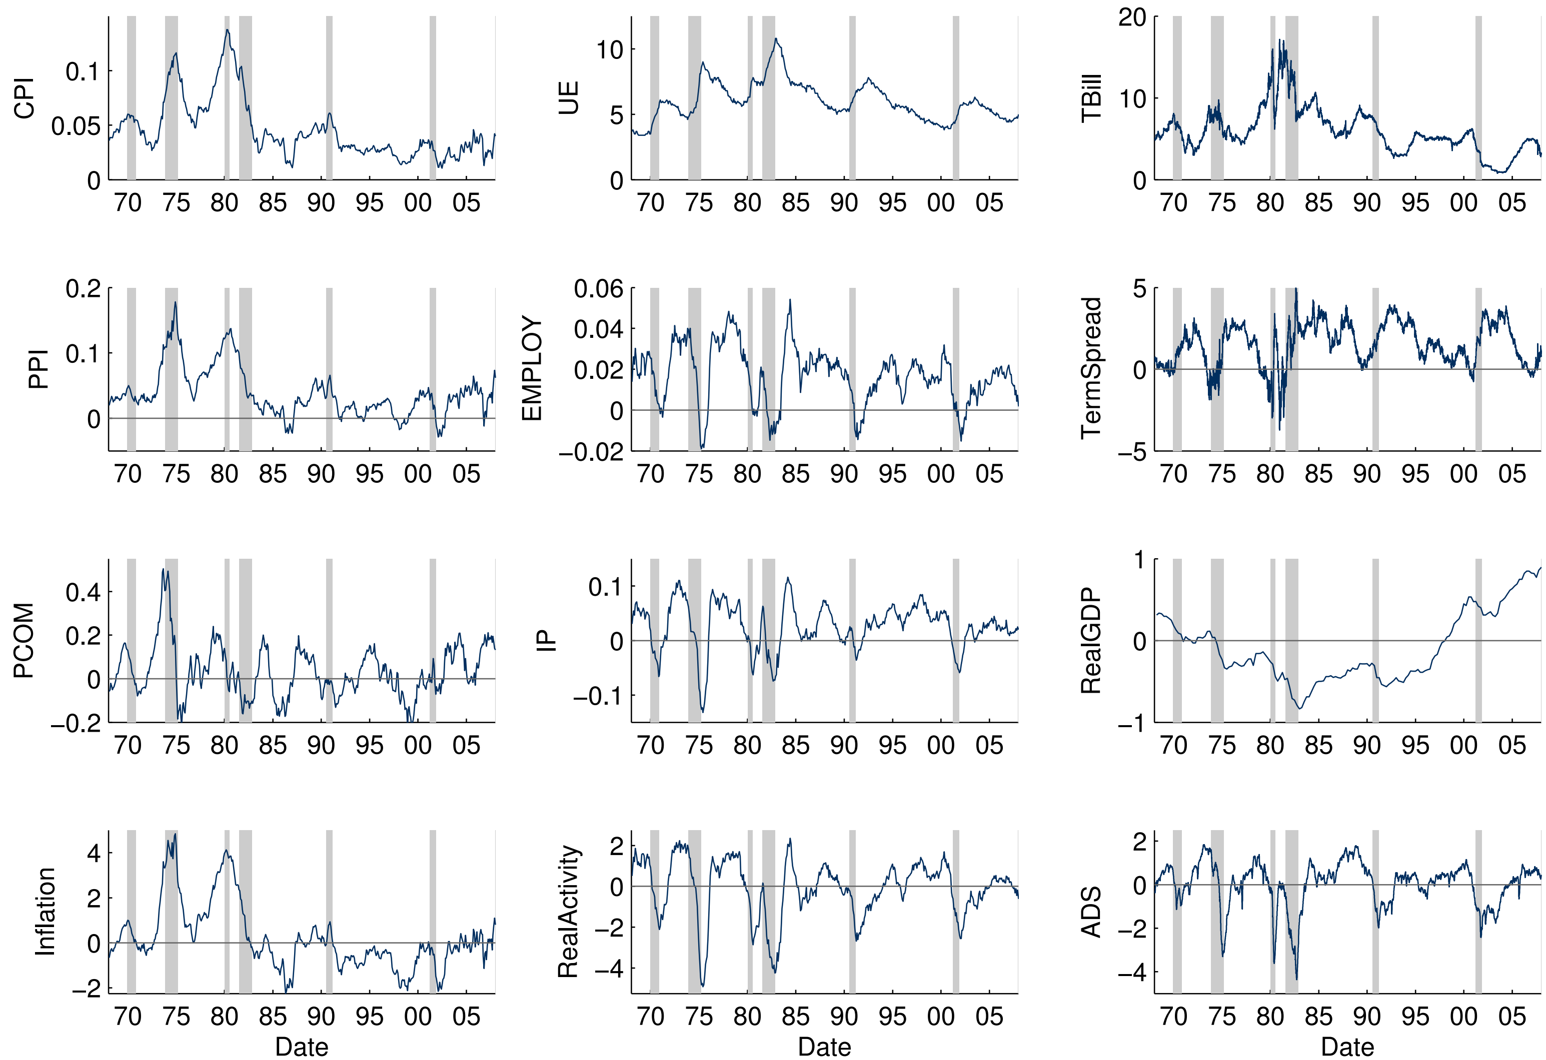

In the bottom row of the left and central columns are Ang & Piazzesi's monthly inflation and real activity indicators; their respective input series are reported in the first three rows. The right column displays the T-Bill rates, the term spread, the detrended real GDP, and the business conditions index.

Table 3.2: 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 3.1. 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.*

The lower left panel of Figure 3.1 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 3.1. 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.*

Table 3.2 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, † and Ang and Piazzesi's factors.

*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.

†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.

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, 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.

3.3.3 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 (3.7)–(3.10). Table 3.3 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).*

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

* 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.

Table 3.3: 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)
$\boldsymbol{\omega}, \mathbf{m}$	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$(\mathbf{R}_{\mathbf{t+1}}, \mathbf{h}_{\mathbf{t+2}})$	-69.70%	-71.75%	-71.72%	-74.30%
Corr$(\boldsymbol{\varepsilon}_{\mathbf{t+1}}, \mathbf{u}_{\mathbf{t+1}})$			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.,*

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

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

0.70 basis points ($0.70 \times\ 10^{-4}$) 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 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 3.3 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 3.2 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 $\sqrt{252\,\mathrm{E}\left[\,h_t\,\right]} = \left(252\omega/(1-\alpha(1+\gamma^2)+\beta)\right)^{\frac{1}{2}}$. 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 3.3 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 3.3.2, 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.* 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

*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.

Figure 3.2: Fundamental Variance Processes

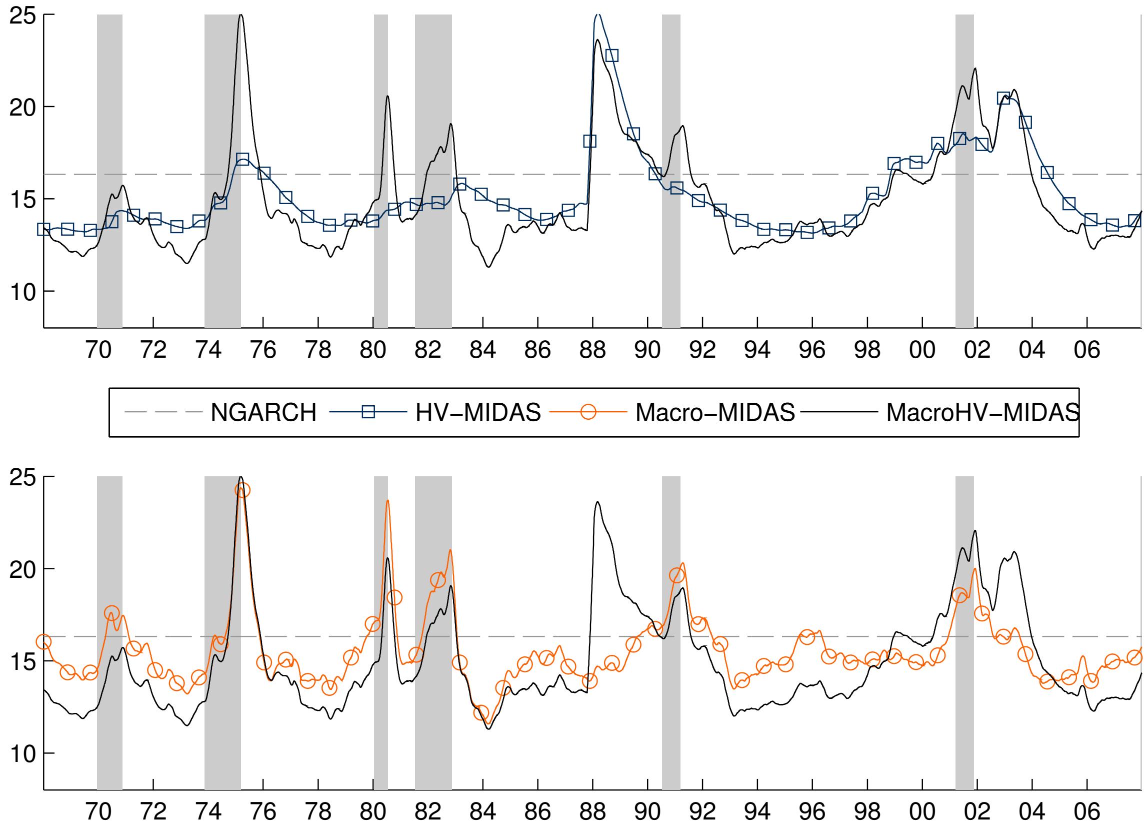

In both panels, we plot, as a solid black line, the annualized fundamental volatility level of the MacroHV-MIDAS model along with a dashed horizontal line at 16.31%, which corresponds to the NGARCH model's expected volatility, $\sqrt{252\omega/(1-\alpha(1+\gamma^2)-\beta)}$*. In the upper panel, we superimpose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

index reflect heightened expectations about the state of the economy rather than the arrival of unexpected positive news.

3.4 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.

3.4.1 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 (3.1). 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 (3.13). Moreover, as observed in Section 3.3.3, the correlation between the two innovation processes, $\text{Corr}_t(\varepsilon_{t+1}, u_{t+1})$, 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

$$\varepsilon_{t+1} = \rho u_{t+1} + \sqrt{1-\rho^2} z_{t+1} \, , \tag{3.15}$$

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

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 (3.7). We show in the appendix that, under this assumption, the correlation structure of Equation (3.15) leads to the following risk neutralization of the macroeconomic and pure-market innovation processes:

$$u_t^* = u_t + \rho\lambda \tag{3.16}$$

$$z_t^* = z_t + \sqrt{1-\rho^2}\lambda\,. \tag{3.17}$$

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 (3.16) and (3.17), the risk-adjusted returns of the MacroHV-MIDAS model are given by

$$R_{t+1} = r - \frac{1}{2}\sqrt{\tau_{t+1}g_{t+1}} + \sqrt{\tau_{t+1}g_{t+1}}\left(\rho u_{t+1}^* + \sqrt{1-\rho^2}z_{t+1}^*\right) \tag{3.18}$$

$$g_{t+1} = (1-\alpha(1+\gamma^2)-\beta) + \alpha g_t\left(\rho u_t^* + \sqrt{1-\rho^2}z_t^* - \gamma - \lambda\right)^2 + \beta g_t \tag{3.19}$$

$$\log\left(\tau_{t+1}\right) = m + \theta\sum_{k=0}^{K-1}\phi_k(w_{hv})\,HV_{t-k} + \theta_m\sum_{k=0}^{K-1}\phi_k(w_m)\,x_{t-k} \tag{3.20}$$

$$x_t = \varphi x_{t-1} + \sqrt{h_t^x}\left(u_t^* - \rho\lambda\right) \tag{3.21}$$

$$h_t^x = \omega_x + \alpha_x h_{t-1}^x\left(u_{t-1}^* - \rho\lambda\right)^2 + \beta_x h_{t-1}^x\,, \tag{3.22}$$

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

*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.

Table 3.4: Option-Pricing Results

					MIDAS		
	Moneyness	**N**	**Avg.**	**NGARCH**	**HV**	**Macro**	**MacroHV**
Log-Likehood				33791.3	33806.7	33803	**33822**
IVRMSE		68923	18.41	5.23	5.14	5.21	**5.10**
DTM $\leq$ 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 $\leq$ 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 $\leq$ 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 3.3.3. 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:

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

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.

3.4.2 Option Valuation Results

Using parameter estimates of Table 3.3 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 (2008); 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.[*] 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 3.3.3. 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 3.4. 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 3.3 sheds some light on these results by casting them in a time-series perspective. The upper panel reports yearly IVRMSE values. First, the MacroHV-MIDAS

[*] Christoffersen, Dorion, Jacobs, and Wang (2008) 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).

Figure 3.3: Option Valuation: A Time-Series Perspective

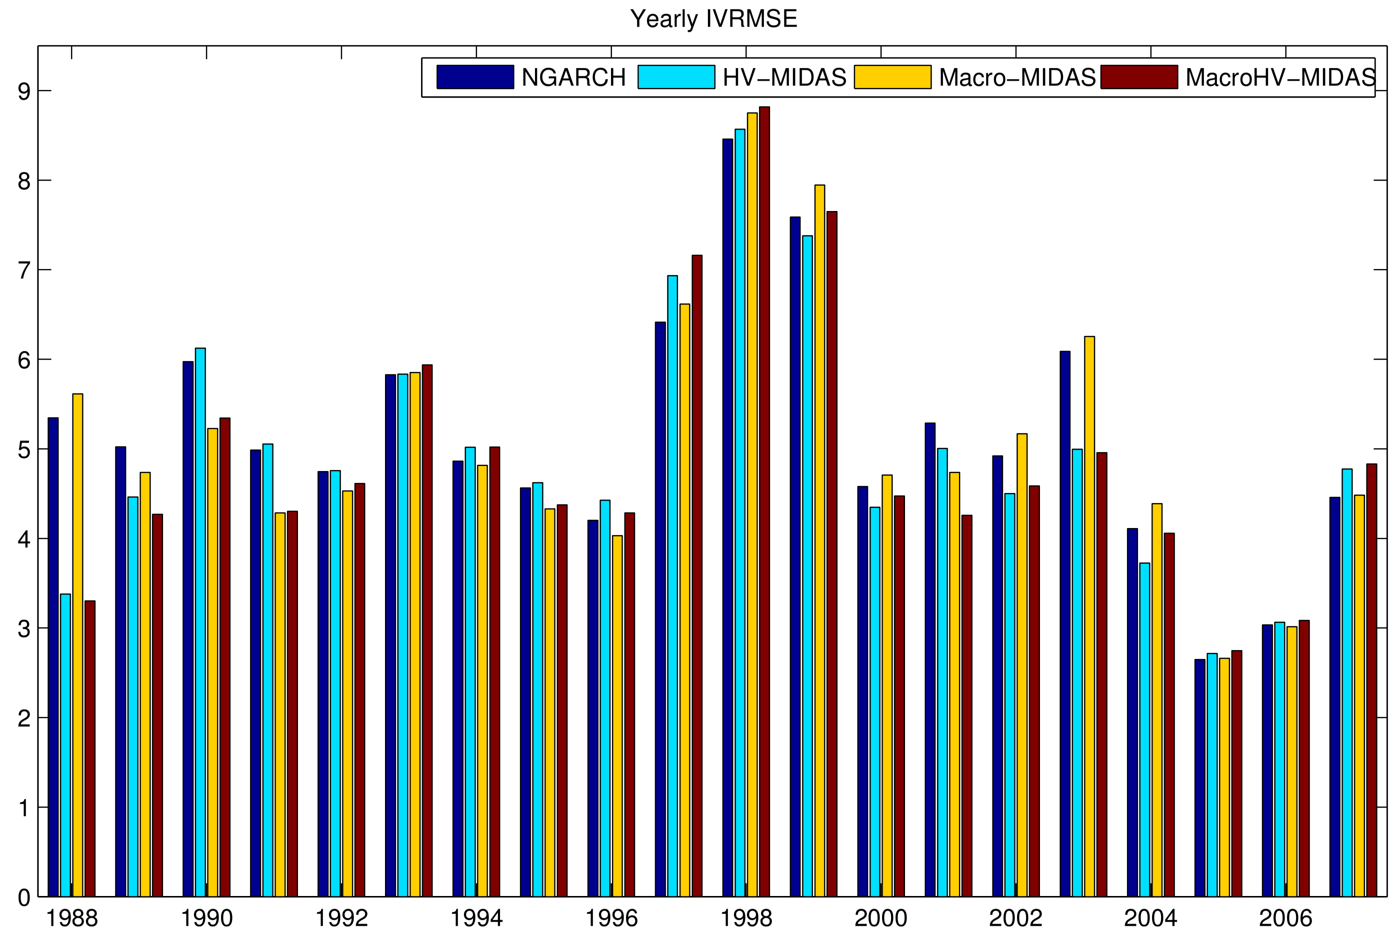

This figure reports yearly IVRMSE measures for the NGARCH model along with those for the MacroHV-MIDAS model and the nested HV-MIDAS and Macro-MIDAS models.

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 3.4 and 3.5 further illustrate how these results unfold through time and break up the results along the maturity and moneyness dimensions. Figure 3.4 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 3.5 similarly reports the MacroHV-MIDAS model's bias at forecasting implied volatilities.* 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 3.4, 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 3.3, is shown in Figure 3.5 to be due to large negative bias at all maturities and over all moneyness levels during this period.

A second look at Figure 3.2 can provide us with an intuition why this is so. Indeed, while the VIX reaches all-time highs during the Russian/LTCM crisis, the

*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.

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, 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 3.5. 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.

3.4.3 The Impact of Business Conditions

Along with Figure 3.4, 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 3.5, 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 3.5 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,

$$x_t^{(63)} = \frac{1}{63} \sum_{s=t-31}^{t+31} x_s \,. \tag{3.23}$$

Figure 3.4: MacroHV-MIDAS Model's IVRMSE Improvement Over the NGARCH Model

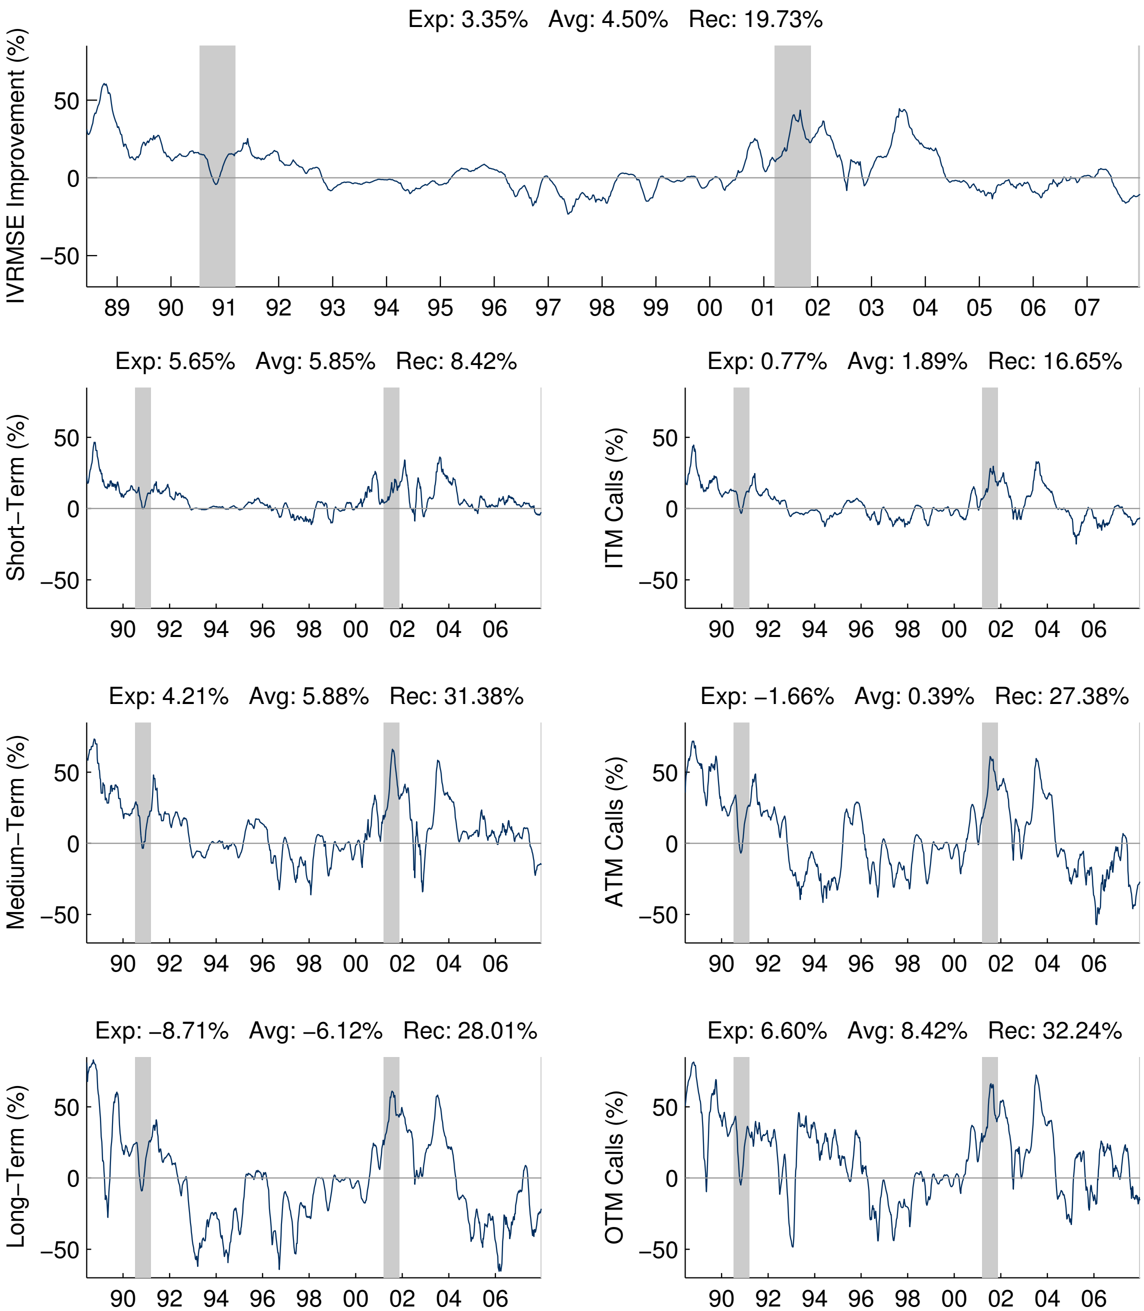

Using parameter estimates in Table 3.3, we price options and compute weekly IVRMSE measures for the NGARCH model, $IVRMSE_w^{\text{NG}}$*, and for the MacroHV-MIDAS model,* $IVRMSE_w^{\text{HV},\Delta x}$*. This figure reports the relative improvement of using the latter model over the former, that is,* $(IVRMSE_w^{\text{NG}} - IVRMSE_w^{\text{HV},\Delta x})/IVRMSE_w^{\text{NG}}$ *. 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 3.5: Implied-Volatility-Forecasting Bias of the MacroHV-MIDAS Model

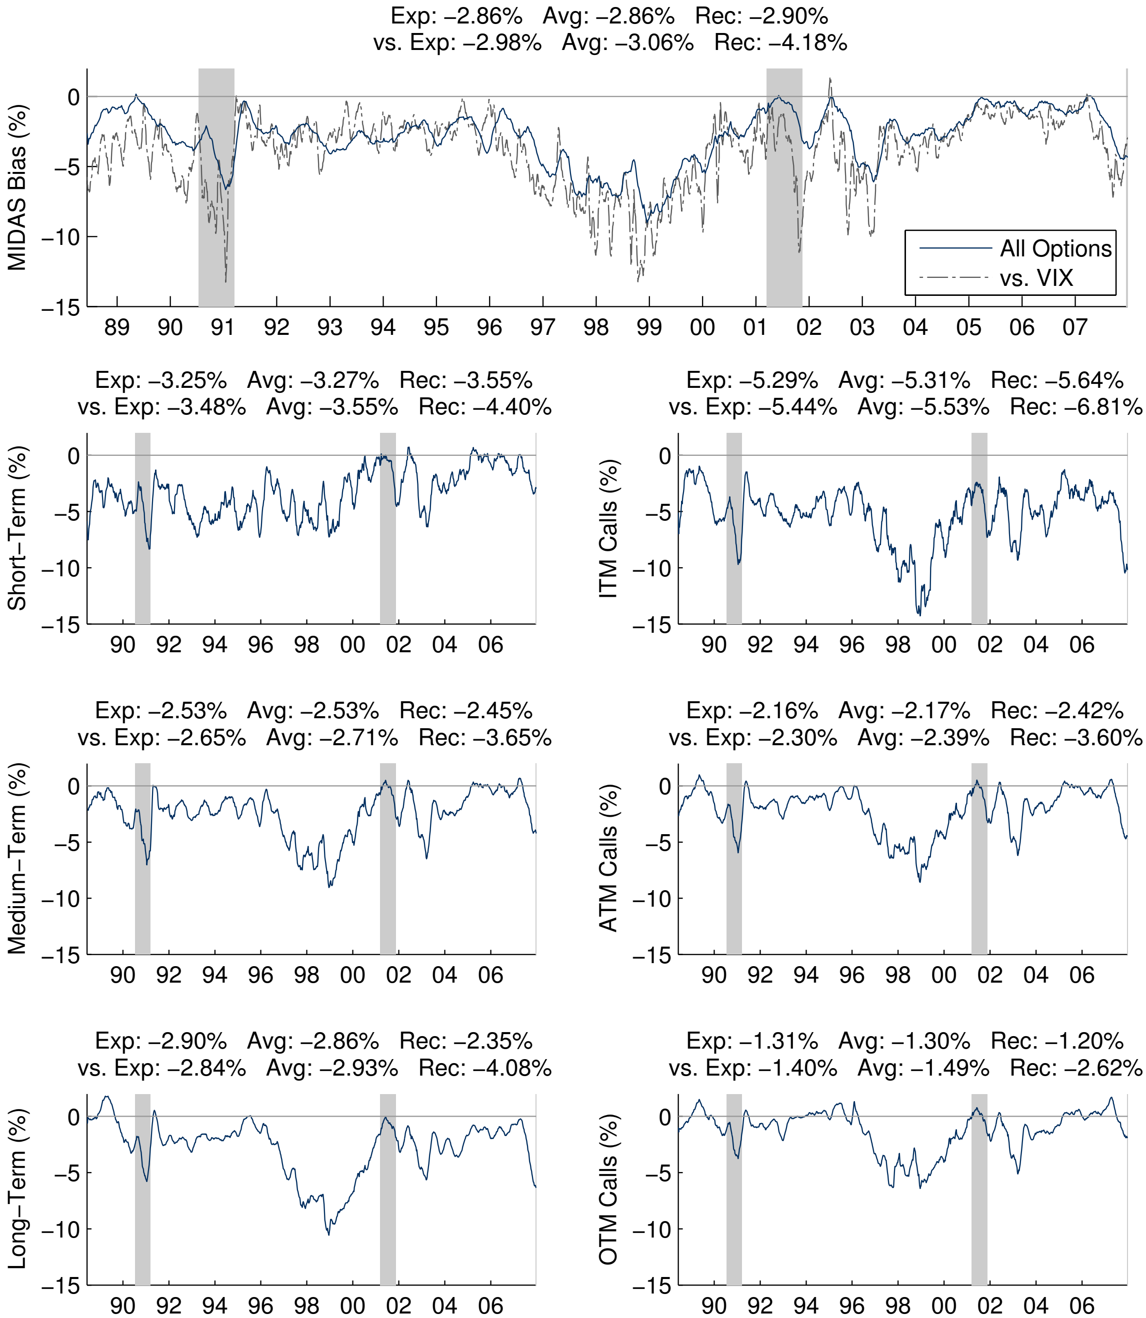

This figure reports the time series of implied-volatility biases for the MacroHV-MIDAS model. On the left-hand side, results are divided along the options' maturity: 45 days to maturity (DTM) or less, between 46 and 90 DTM, or more than 90 DTM. On the right-hand side, results are divided along options' 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).

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 3.5 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 3.6, and statistics for this time series are reported in the first row of Table 3.5'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 3.6. The right-column panels of Figure 3.6 and the remainder of Table 3.5's Panel B report the same results but in the price domain.*

As we have seen in Figure 3.5, 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 3.5'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 3.5 shows that business conditions indeed play a key role the ability of MacroHV-MIDAS model to outperform its benchmark.

Table 3.6 sheds further light on the role played by business conditions in the

*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.

Figure 3.6: Synthetic Options

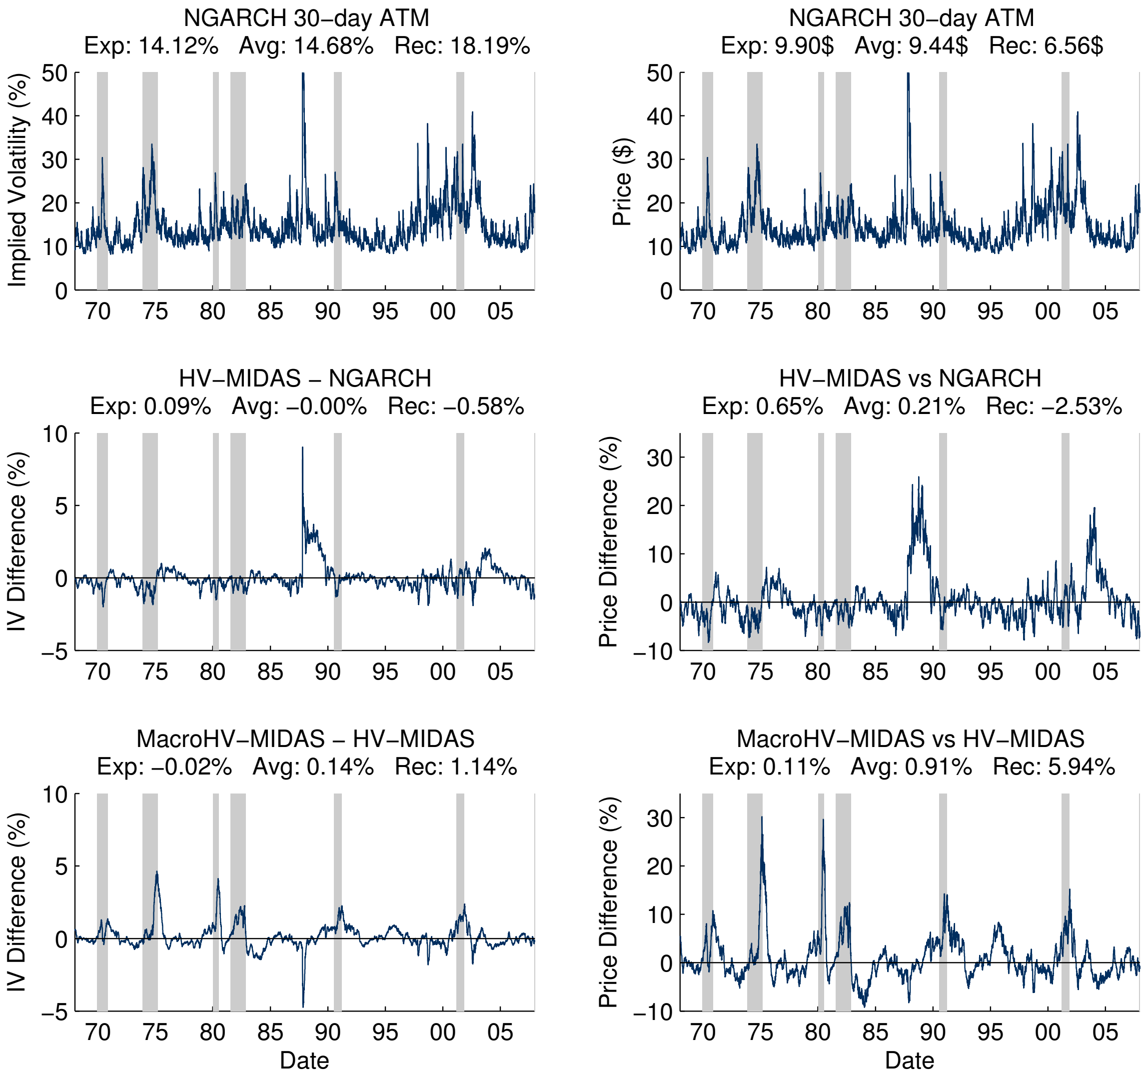

The upper-left panel of this figure reports the time series of implied volatilities obtained using the NGARCH model to price synthetic, at-the-money options with 30 days to maturity, daily from January 1968 to December 2007. These options are created assuming that their strike is equal to the index value 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. $\left(C_t^{\text{HV}} - C_t^{\text{NGARCH}}\right)/C_t^{\text{NGARCH}}$.

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%. 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 3.6 illustrates that this pattern is rather robust to maturities and moneyness levels.

As highlighted by Panel B of Table 3.6, 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.

3.5 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

Table 3.5: 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 Difference (%)	0.77	1.12	3.31	4.00	6.26	9.05
	(0.07)	(0.07)	(0.19)	(0.11)	(0.20)	(0.30)

In Panel A, using parameter estimates in Table 3.3, we price options and compute weekly IVRMSE measures for the NGARCH model, $IVRMSE_w^{\text{NG}}$, and for the MacroHV-MIDAS model, $IVRMSE_w^{\text{HV},\Delta x}$. Panel A reports averages of the relative improvement resulting from using the latter model over the former, that is, $(IVRMSE_w^{\text{NG}} - IVRMSE_w^{\text{HV},\Delta x})/IVRMSE_w^{\text{NG}}$. 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 3.6). 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 $\left(C_t^{\text{HV},\Delta x} - C_t^{\text{NG}}\right)/C_t^{\text{NG}}$.*

Table 3.6: 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 3.5. Panel A reports averages of the relative improvement resulting from using the MacroHV-MIDAS model over the HV-MIDAS, that is, $(IVRMSE_w^{HV} - IVRMSE_w^{HV,\Delta x})/IVRMSE_w^{NG}$ *; the IVRMSE of the NGARCH is kept at the denominator to ease comparison with Table 3.5. 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* $\left(\mathrm{IV}_t^{HV,\Delta x} - \mathrm{IV}_t^{HV}\right)$. *The comparison is also performed in the price domain, i.e. based on* $\left(C_t^{HV,\Delta x} - C_t^{HV}\right)/C_t^{NG}$.

directly related to the risk aversion of a representative agent.* 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.† 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 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 3.4 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 3.B for details. The resulting daily volatility premium series is displayed in the lower panel of Figure 3.7.

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 3.4.2.

---

*See, amongst others, Heston (1993), Eraker (2007), and Bollerslev, Gibson, and Zhou (2009).

†See, for instance, Bollerslev, Gibson, and Zhou (2009), Corradi, Distaso, and Mele (2009), and Bollerslev, Tauchen, and Zhou (2009).

Figure 3.7: The Volatility Premium

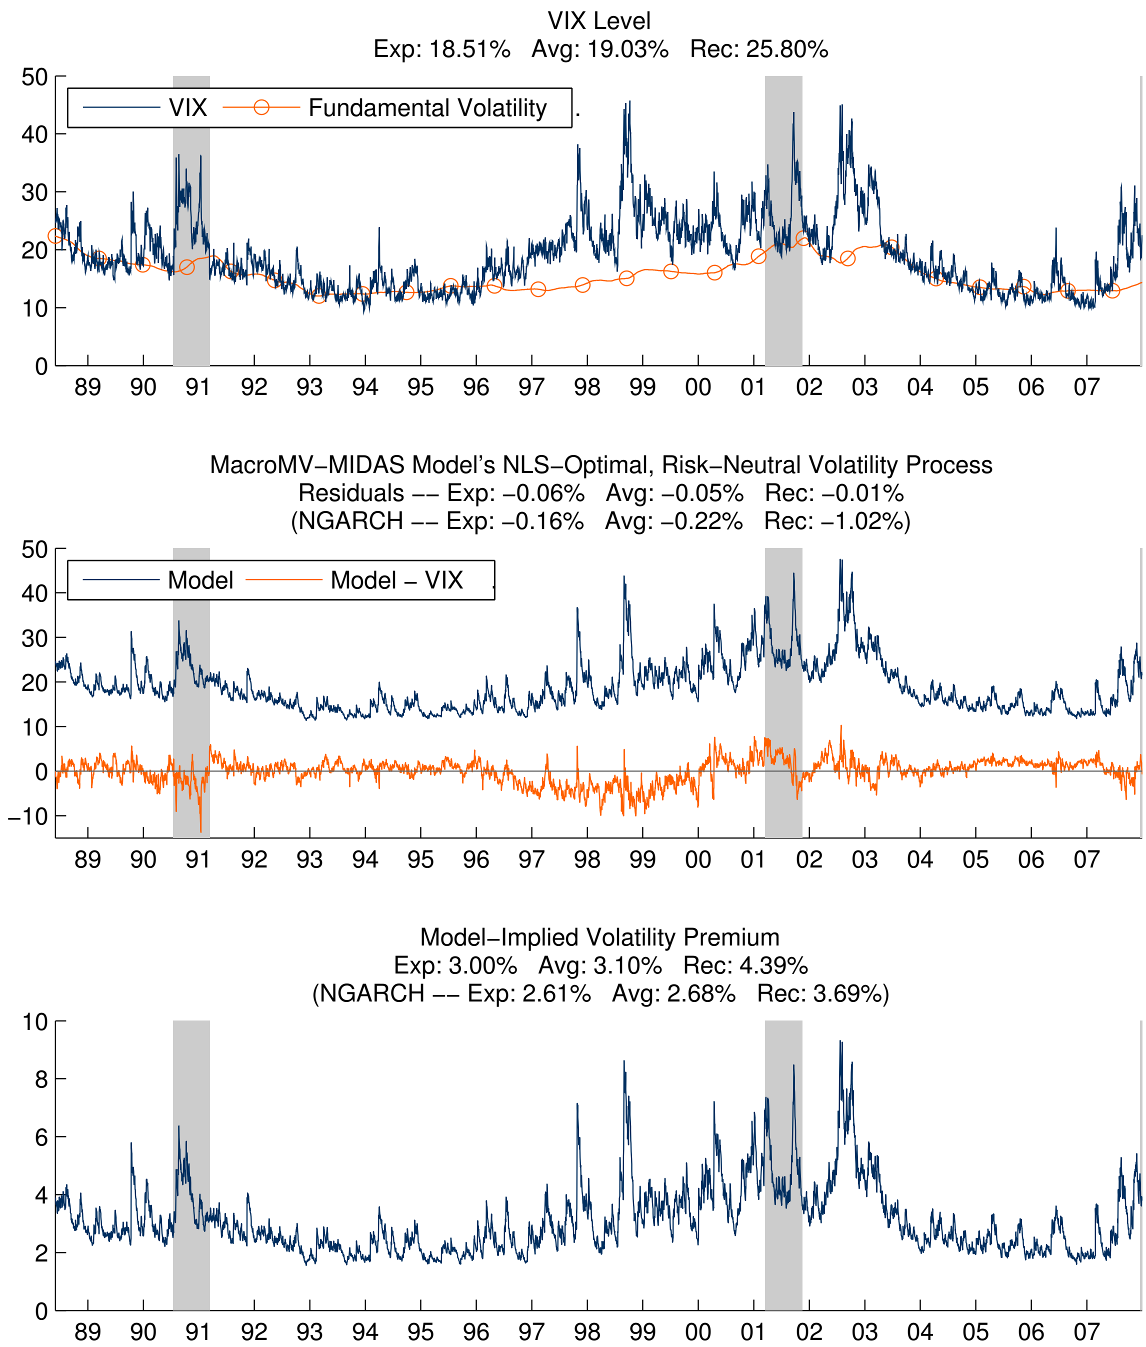

In the upper panel, we plot the VIX through time and report its overall average level as well as that through expansion and recession periods (Exp/Avg/Rec). For comparison sakes, we also plot the MacroHV-MIDAS model's fundamental volatility process; note that the latter is the long-run mean-reversion 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.

Panel C of Table 3.7 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.* 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 volatility 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%

*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).

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.

3.6 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.

Appendix

3.A Risk Neutralization

We consider a GARCH model of the form

$$R_{t+1} = r + \lambda_{t+1}\sqrt{h_{t+1}} - \tfrac{1}{2}h_{t+1} + \sqrt{h_{t+1}}\big(\rho u_{t+1} + \sqrt{1-\rho^2}z_{t+1}\big) \tag{3.A.1}$$

$$h_{t+1} = f\left(\,\cdot\, \mid \Theta,\, \mathcal{F}_t \right) , \tag{3.A.2}$$

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 (3.16) and (3.17), we here draw on Christoffersen, Elkamhi, Feunou, and Jacobs (2009; henceforth CEFJ) 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:

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

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

$$\Psi_t^{u,z}\left(\eta_u, \eta_z\right) = \tfrac{1}{2}\left(\eta_u^2 + \eta_z^2\right) . \tag{3.A.4}$$

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

$$\begin{aligned} 1 &= \mathrm{E}_{t-1}^{\mathbb{Q}}\left[\frac{S_t}{S_{t-1}} \bigg/ \frac{B_t}{B_{t-1}} \right] = \mathrm{E}_{t-1}^{\mathbb{P}}\left[\frac{\xi_t}{\xi_{t-1}} \frac{S_t}{S_{t-1}} \bigg/ \frac{B_t}{B_{t-1}} \right] \\ &= \mathrm{E}_{t-1}^{\mathbb{P}}\left[\exp\left\{-\eta_{u,t}u_t - \eta_{z,t}z_t - \Psi_t^{u,z}\left(\eta_{u,t}, \eta_{z,t}\right)\right\} \exp\left\{ \lambda_t\sqrt{h_t} - \tfrac{1}{2}h_t + \sqrt{h_t}\big(\rho u_t + \sqrt{1-\rho^2}z_t\big)\right\}\right] \end{aligned} \tag{3.A.5}$$

or, equivalently,

$$0 = \Psi_t^{u,z}\left(\eta_{u,t} - \rho\sqrt{h_t}\,, \eta_{z,t} - \sqrt{(1-\rho^2)h_t}\right) - \Psi_t^{u,z}\left(\eta_{u,t}, \eta_{z,t}\right) + \lambda_t\sqrt{h_t} - \tfrac{1}{2}h_t\,, \quad (3.A.6)$$

which boils down to

$$\rho\eta_{u,t} + \sqrt{1-\rho^2}\eta_{z,t} = \lambda_t\,. \quad (3.A.7)$$

Equation (3.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 (3.A.1) is equivalent to

$$R_{t+1} = r + \lambda_{t+1}\sqrt{h_{t+1}} - \tfrac{1}{2}h_{t+1} + \sqrt{h_{t+1}}\varepsilon_{t+1}\,. \quad (3.A.8)$$

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

$$\xi_\tau \equiv \left.\frac{\mathrm{d}\mathbb{Q}}{\mathrm{d}\mathbb{P}}\right|\mathcal{F}_\tau = \exp\left\{-\sum_{t=1}^{\tau}\left(\eta_t\varepsilon_t + \Psi_t^{\varepsilon}\left(\eta_t\right)\right)\right\}\,, \quad (3.A.9)$$

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 (3.A.9) to be an EMM, it must be that

$$\begin{aligned}
1 =&\ \mathrm{E}_{t-1}^{\mathbb{Q}}\left[\frac{S_t}{S_{t-1}}\Big/\frac{B_t}{B_{t-1}}\right] = \mathrm{E}_{t-1}^{\mathbb{P}}\left[\frac{\xi_t}{\xi_{t-1}}\frac{S_t}{S_{t-1}}\Big/\frac{B_t}{B_{t-1}}\right] \\
=&\ \mathrm{E}_{t-1}^{\mathbb{P}}\left[\exp\left\{-\eta_t\varepsilon_t - \Psi_t^{\varepsilon}\left(\eta_t\right)\right\}\exp\left\{\lambda_t\sqrt{h_t} - \tfrac{1}{2}h_t + \sqrt{h_t}\varepsilon_t\right\}\right] \quad (3.A.10)\\
\Leftrightarrow 0 =&\ \Psi_t^{\varepsilon}\left(\eta_t - \sqrt{h_t}\right) - \Psi_t^{\varepsilon}\left(\eta_t\right) + \lambda_t\sqrt{h_t} - \tfrac{1}{2}h_t\,, \quad (3.A.11)
\end{aligned}$$

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

$$\begin{aligned}
\xi_\tau &= \exp\left\{-\sum_{t=1}^{\tau}\left(\lambda_t\varepsilon_t + \tfrac{1}{2}\lambda_t^2\right)\right\} \\
&= \exp\left\{-\sum_{t=1}^{\tau}\left(\lambda_t\rho u_t + \lambda_t\sqrt{1-\rho^2}z_t + \tfrac{1}{2}\overbrace{\left(\rho^2\eta_{u,t}^2 + 2\rho\sqrt{1-\rho^2}\eta_{u,t}\eta_{z,t} + (1-\rho^2)\eta_{z,t}^2\right)}^{\text{Using Equation (3.A.7)}}\right)\right\} \\
&= \exp\left\{-\sum_{t=1}^{\tau}\left(\eta_{u,t}u_t + \eta_{z,t}z_t + \Psi_t^{u,z}\left(\eta_{u,t}, \eta_{z,t}\right)\right)\right\}
\end{aligned}$$

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

$$\eta_{u,t} = \rho\lambda_t \quad \text{and} \quad \eta_{z,t} = \sqrt{1-\rho^2}\lambda_t \,. \tag{3.A.12}$$

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

$$\varepsilon_t^* = \varepsilon_t + \lambda_t = \rho u_t^* + \sqrt{1-\rho^2} z_t^* \,. \tag{3.A.13}$$

3.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 $\mathrm{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right]$ and can be easily computed using Monte-Carlo integration. In order to study the properties of the volatility premium, $\mathrm{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right] - \mathrm{E}_t^{\mathbb{P}}\left[\sqrt{h_{t+21}}\right]$, implied by our model, we must first ensure that the bias of its implied volatility process is minimized; that is, $\mathrm{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right]$ should track the VIX level as closely as possible. However, as can be observed in Figure 3.4, 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

$$\log \mathrm{E}_t^{\mathbb{P}}\left[e^{-r}\frac{S_{t+1}}{S_t}\right] = \log \mathrm{E}_t^{\mathbb{P}}\left[\exp\left\{\lambda\sqrt{h_{t+1}} - \frac{1}{2}h_{t+1} + \sqrt{h_{t+1}}\varepsilon_{t+1}\right\}\right] = \lambda\sqrt{h_{t+1}}. \tag{3.B.1}$$

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.* Hence, an interesting alternative that has been pursued by many authors is

*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.

to estimate a forward-looking price of risk parameter λ from option data.[*] 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,

$$\min_{\lambda} \sum_{t \in \mathcal{T}} \left(\mathrm{E}_t^{\mathbb{Q}} \left[\sqrt{h_{t+21}} \right] - \mathrm{VIX}_t \right)^2 . \tag{3.B.2}$$

A nonlinear least squares (NLS) estimation of λ is performed using the implied-volatility sum of squared errors criterion in Equation (3.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.

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 (3.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.[†] 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 $\mathrm{E}_t^{\mathbb{P}} \left[\sqrt{h_{t+21}} \right]$; 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 3.7, Panel A summarizes VIX observations retained for our NLS estimation exercise, while Panel B reports the results obtained by minimizing Equa-

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

[†]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 3.4.2. 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.

Table 3.7: 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

		$\sqrt{\tau_t}(\mathbf{HV})$	$\sqrt{\tau_t}(\boldsymbol{\Delta}\mathbf{x})$	$\sqrt{\tau_t}$					
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 3.3) 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 3.7, $\mathrm{E}_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right] - \mathrm{E}_t^{\mathbb{P}}\left[\sqrt{h_{t+21}}\right]$, 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.

tion (3.B.2). Note that these latter results are reported in implied volatility root mean squared error (IVRMSE) terms, that is, using

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

The IVRMSE is strictly monotone in the objective function of Equation (3.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 3.3), 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-MIDAS model, the benchmark IVRMSE is lower at 5.09% and yet the improvement is greater at 47.6%.*

As seen in Figure 3.4, 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 bi-

*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 3.7. 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.

ases 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, $E_t^{\mathbb{Q}}\left[\sqrt{h_{t+21}}\right]$, is close to bias-free throughout the twenty years of data we consider.

4 Credit Default Swaps, Options, and Systematic Risk

Joint work with

Redouane Elkamhi

Henry B. Tippie College of Business, University of Iowa

Iowa City, IA 52242-1994, US (redouane-elkamhi@uiowa.edu)

Jan Ericsson

Desautels Faculty of Management, McGill University,

Montreal, Quebec, H3A 1G5, Canada (jan.ericsson@mcgill.ca)

Abstract We here study the impact of systematic risk on the pricing of two economically similar derivative contracts: credit default swaps and equity put options. We document, for roughly 130 firms that have been part of the CDX index between 2004 and 2007, that the greater proportion of a firm's volatility that is systematic, the more expensive it is to purchase insurance via both (i) put options and (ii) credit default swaps. We provide evidence that these two derivatives are influenced by systematic risk through the same channel.

Keywords Systematic risk; Spreads; Credit default swaps; Option valuation; Merton model; Geske model; Volatility.

4.1 Introduction

Recent work has documented the existence of a significant component in corporate bond and default swap credit spreads over and above the actuarially fair compensation for expected losses.* Although, by now, it has been documented that a part of this component compensates for illiquidity and other risks, it is quite plausible that the the bulk of this additional spread represents compensation for systematic market risk (Elton, Gruber, Agrawal, and Mann 2001).† The purpose of this paper is twofold. The first is to examine whether the cross section of risk premia in default swap markets can, in fact, be related to firm-specific measures of systematic risk. Second, in the light of recent related evidence in stock option markets (Duan and Wei 2009), we seek to establish a link between risk premia in these two markets.

To lay the foundation for our empirical work, we consider risk premia in credit and stock option markets through the lens of the Merton (1974) and Geske (1979) models. Although many important improvements to both of these models have been suggested in the more than three decades since their publication, once augmented with the CAPM, these models are rich enough to provide both intuition and empirical implications which are internally consistent across both markets.‡

Augmented with the CAPM, Merton's framework provides two fundamental insights: (i) firms with very different systematic risk profiles can have the same physical default probability but very different spreads; and (ii) firms with higher default rates will tend to have proportionally less compensation for systematic risk in their spreads. These insights have important practical implications. First, the debt of firms within a given rating category may trade at very different spreads depending on the exposure to market risk. This is clearly a crucial aspect for bond portfolio

*We will throughout the paper view credit default swaps and corporate bond spreads as economically similar risk measures. This is justifiable exactly by means of a simple buy-and-hold arbitrage argument for floating rate corporate bonds vis-à-vis CDS contracts. In practice, CDS spreads and corporate bond spreads are not exactly equivalent. For example, Kim, Li, and Zhang (2009) provide an interesting study of the differential, known as the basis. Nevertheless, we find it plausible that our findings in relation to risk premia in CDS markets are instructive for bond markets as well.

†See e.g. Ericsson and Renault (2006), Chen, Lesmond, and Wei (2007), Tang and Yan (2007), Berndt, Jarrow, and Kang (2007), Arora, Gandhi, and Longstaff (2009).

‡Interesting extensions of the Merton and Geske models include Black and Cox (1976), Leland (1994), Longstaff and Schwartz (1995), Collin-Dufresne and Goldstein (2001), Bhamra, Kuehn, and Strebulaev (2009).

management - if bonds are selected on the basis of providing a high yield for a given rating, then this should tend to bias the portfolio towards greater amounts of systematic risk. A similar point can be made for capital standards, like those imposed by Basel II. If capital requirements are set according to risk buckets defined in terms of credit ratings, then institutions will have an incentive to increase the retun on equity by opting for higher spread exposures within rating classes. Given the findings in this paper, this would tend to yield a greater concentration of systematic risk.* Point (ii) suggests that more highly rated debt will proportionally provide more compensation for systematic risk. In other words, spreads on firms near default, will compensate more for expected losses and less for systematic risk. Highly rated firms will compensate mostly for systematic risk.

The volatility premium in stock option markets can be analyzed along similar arguments. We define the volatility premium as the difference between the risk-neutral and physical implied volatility levels. Risk-neutral implied volatility may sound pleonastic as implied volatility is usually thought of as risk-neutral concept. Yet, augmenting the Geske model with the CAPM under the objective measure, one can derive a physical counterpart to implied volatility, which is equal to the risk-neutral one for a zero-beta firm and, interestingly, decreases as beta increases. Hence, the augmented model predicts that the volatility premium explains a larger proportion of the price of options on stocks facing greater systematic risk. These predictions are consistent with recent empirical findings by Duan and Wei (2009) who show that, for firms with greater proportions of systematic risk, the difference between implied and historical volatilities is larger.†

We test our predictions on credit default swaps (CDS) and option data for the approximately 130 firms that are part of the CDX index between 2003 and 2007. We find that the split of a firm's volatility into systematic and idiosyncratic components matters for both option and CDS prices even after controlling for total volatility. We first document that the results of Duan and Wei (2009) are robust to a broader and more recent sample of stock option data; that is, the proportion of systematic

---

*A similar point has been made in the context of collateralized debt obligations (CDO) by Coval, Jurek, and Stafford (2009).

†See Bollerslev, Tauchen, and Zhou (2009), Driessen, Maenhout, and Vilkov (2009), and Todorov (2010) for alternative explanations of the volatility (risk) premium.

risk in a firm's total volatility increases the difference between implied and historical volatilities, interpreted as a risk premium. These results are economically significant. A firm with 100 percent systematic volatility will, on average, have a volatility risk premium that is 3% greater, on an annual basis, than a firm with only idiosyncratic risk. The proportion of the former, "systematic" firm's implied volatility that is due to the premium is around 28% greater than for the latter, "idiosyncratic" firm.

For CDS contracts, the issue is more subtle. We first show that the proportion of systematic volatility is important also for explaining CDS spreads. However, surprisingly, CDS spreads are robustly, negatively related to the systematic risk proportion for individual firms. We show that this intriguing result can at least partially be explained by differences in systematic volatility proportions across rating categories; differences which are in fact consistent with what the augmented Merton model would predict. Although higher leverage in lower rating categories yields higher betas on average, the proportion of systematic volatility is highest for the most highly rated firms. At the same time, these firms tend to have lower CDS spreads. In essence, understanding the impact of systematic volatility on the price of default protection requires correcting for the various firm characteristics that impact the level of the physical default probability. Candidates for such characteristics can be gleaned by studying risk premia in the augmented Merton and Geske models.

As noted, corporate bond and credit default swap spreads can be separated into two fundamentally different components: expected loss and risk premium components. Elton, Gruber, Agrawal, and Mann (2001) were the first to measure the size of the expected loss component. They found that expected losses explain a negligible amount of short term spreads and between 10-30% for 10 year bonds. Elkamhi and Ericsson (2008) find that the fraction due to expected losses is highly time varying. In addition, risk premium and expected loss components are less than perfectly correlated, depending in nonlinear and different ways on volatility and leverage. These, in turn vary both cross-sectionally and over time. Thus it may not be all that surprising that it is difficult to link spreads with a risk premium proxy, as these nonlinearities cloud the link.

We then consider an alternative dependent variable - the spread corrected for the expected loss component. When we repeat our empirical tests on this risk premium component, we find a robust and positive relationship, confirming our initial intuition that default protection is more expensive, all else equal, for firms with more systematic risk.

In summary, we have documented that equity and credit derivatives prices contain risk premia in accordance with the predictions of simple contingent claims pricing models augmented with the CAPM. This is of course interesting by itself. In addition it has, as mentioned above, implications for asset allocation and portfolio risk measurement. Moreover, a growing literature has begun to use information in stock option markets as explanatory variables for default swap spreads, and our findings adds to this literature by discussing the relationship between pricing in these two derivatives markets.

One of the first papers to relate risk premia in the two markets is Cremers, Driessen, and Maenhout (2008). They consider whether jump risk premia from individual stock option prices can explain the high observed level of credit spreads. They rely on a structural credit risk model extended to incorporate jump-diffusion firm value dynamics. They calibrate this model to firm and index options as well as historical default and equity premium data. Option-implied jump risk premia allows the calibrated model to produce spreads much nearer to observed levels.

Another recent example is Cao, Yu, and Zhong (2010), who find that individual stock options' implied volatility alone can explain about half of the variation in CDS spreads. They also document that the link between the CDS market and the options market is stronger among firms with lower credit ratings. Further, they argue that the volatility risk premium is a significant determinant of CDS spreads, even in the presence of future predicted volatility. They argue that the success of implied volatility based variables is at least in part due to an embedded volatility risk premium.* Our purpose here is, although related, quite different. We seek to understand whether the cross-section of option and default swap prices can be explained by systematic risk. Given the evidence in favour provided so far, and that the empirical predictions for the two markets are derived using two internally

*Wang, Zhou, and Zhou (2010) measure the volatility risk premium differently – relying on high frequency data for expected volatility and model-free implied volatilities to compute option risk premia. They find that a strong predictive power for the variance risk premium as regards spreads, a strength that increases as the credit quality of CDS deteriorates.

consistent models, we claim that CDS and options tend to be simultaneously more expensive for firms with greater systematic risk.

The paper is organized as follows. Section 4.2 discusses risk premia in the augmented Merton and Geske models with a view to defining our empirical predictions. In Section 4.3, we then test these predictions on stock option data (§4.3.2), confirming the results of Duan and Wei (2009), and on CDS spread data (§4.3.3). Finally, Section 4.4 concludes.

4.2 Risk Premia in Credit and Stock Option Markets

In this section, we revisit Merton's and Geske's models. Once augmented with the CAPM under the objective measure, these models provide important empirical implications with respect to risk premia in credit and stock option markets, implications that we here discuss. Our purpose in doing so is to understand the differences predicted under the respective frameworks for prices of credit instruments and stock options with and without risk-adjustment.

4.2.1 Merton (1974) and Geske (1979)

The Merton (1974) and Black and Scholes (1973) models are intrinsically related. The latter allows one to compute the value of an equity option given the current value of the underlying stock and its volatility. In the Merton (1974) model, one simply considers the equity and debt of a firm as claims written on the fundamental value of the firm. The equity can be seen as a call option, the debt as a risk-free bond and a (short) put option.

Both models allow for risk-neutral pricing, which is often convenient since estimating the proper discount factor under the physical measure can prove to be difficult. Indeed, provided that we know the expected growth rate of the firm, μ, so that

$$dV_t = \mu V_t dt + \sigma V_t dW \,, \tag{4.2.1}$$

simple calculus would allow us to compute a "physical" value for the equity

$$\begin{aligned} E^{\mathbb{P}} &= e^{-rT_2}\,\mathrm{E}_0^{\mathbb{P}}\left[(V_{T_2}-F)\cdot I_{V_{T_2}>F}\right] \\ &= e^{-rT_2}\,\mathrm{E}_0^{\mathbb{P}}\left[V_{T_2}\cdot I_{V_{T_2}>F}\right] - Fe^{-rT_2}\,\mathrm{E}_0^{\mathbb{P}}\left[I_{V_{T_2}>F}\right] \\ &= V_0 e^{(\mu-r)T_2}N(d_1^{\mathbb{P}}) - Fe^{-rT_2}N(d_2^{\mathbb{P}}), \end{aligned} \tag{4.2.2}$$

where r is the risk-free rate, F the face value, T_2 the maturity of the firm's debt,

$$d_1^{\mathbb{P}} = \frac{\log V_0/F + (\mu+\sigma^2/2)T}{\sigma\sqrt{T}}, \quad \text{and} \quad d_2^{\mathbb{P}} = d_1^{\mathbb{P}} - \sigma\sqrt{T}. \tag{4.2.3}$$

For the debt

$$\begin{aligned} D^{\mathbb{P}} &= e^{-rT_2}\,\mathrm{E}_0^{\mathbb{P}}\left[F\cdot I_{V_{T_2}>F}\right] + e^{-rT_2}\,\mathrm{E}_0^{\mathbb{P}}\left[V_{T_2}(1-I_{V_{T_2}>F})\right] \\ &= Fe^{-rT_2}\,\mathrm{E}_0^{\mathbb{P}}\left[I_{V_{T_2}>F}\right] + e^{-rT_2}\,\mathrm{E}_0^{\mathbb{P}}\left[V_{T_2}\right] - e^{-rT_2}\,\mathrm{E}_0^{\mathbb{P}}\left[V_{T_2}\cdot I_{V_{T_2}>F}\right] \\ &= Fe^{-rT_2}N(d_2^{\mathbb{P}}) + V_0e^{(\mu-r)T_2} - V_0e^{(\mu-r)T_2}N(d_1^{\mathbb{P}}). \end{aligned} \tag{4.2.4}$$

Note that discounting takes place at the risk-free rate as we wish to compute prices without any compensation for systematic risk. Of course, the growth rate for the assets, μ, matters for the physical probabilities of survival, but implies no risk premium. These formulas are in all regards similar to those obtained when working under the risk-adjusted measure $\mathbb{Q}$ apart from the role played by the firm's growth rate, i.e. the classical Merton formulas obtain when one replaces μ with r in Equation (4.2.3). Assuming that $\mu > r$, would the claims on the firm priced under the physical measure, the value of the equity would be higher than it actually is. The stock price is thus discounted to reflect the risk inherent to the investment in the firm, and this discount is proportional to $\mu - r$.

The same rationale applies to the physical and risk-neutral value of the debt. Given the actual value of debt $D^{\mathbb{Q}}$, bond spreads, $s^{\mathbb{Q}}$, can be obtained by solving

$$D^{\mathbb{Q}} = e^{-(r+s^{\mathbb{Q}})T_2}F. \tag{4.2.5}$$

Note that in what follows we derive comparative statics for bond spreads rather than CDS spreads using the Merton model. This is done purely for expositional purposes and all the qualitative implications are identical for CDS spreads computed in the same model.* Letting $s^{\mathbb{P}}$ be the spread that would be obtained under the physical

*Default swap pricing in the context of structural models is discussed in detail in the appendix.

measure (expected loss), the risk premium component of the spread would simply be $s^{\mathbb{Q}} - s^{\mathbb{P}}$. For two otherwise equal firms with growth rates $\mu_1 > \mu_2 > r$, the Merton (1974) model predicts $s_1^{\mathbb{Q}} = s_2^{\mathbb{Q}}$. It is however obvious from Equations (4.2.4) and (4.2.5) that $s_1^{\mathbb{P}} < s_2^{\mathbb{P}}$. That is, the expected loss of the firm with the greater expected growth is lower than that of the firm with the smaller expected growth.

By predicting that $s_1^{\mathbb{Q}} = s_2^{\mathbb{Q}}$, the Merton (1974) model implies that, while the expected loss of the first firm is lower, the risk premium required to hold the bond of that same firm is higher and offsets the lower expected loss. In other words, while $s_1^{\mathbb{Q}} = s_2^{\mathbb{Q}}$, the proportion of the spread that is due to the risk premium, $(s_i^{\mathbb{Q}} - s_i^{\mathbb{P}})/s_i^{\mathbb{Q}}$, is greater for the firm with the highest expected growth rate, μ_1. While intriguing at first sight, this result actually makes sense in a CAPM world. Indeed, if

$$\mu_i = r + \beta_i \, \mathrm{E}\left[R_{\mathrm{M}} - r \right], \tag{4.2.6}$$

$\mu_1 > \mu_2$ then implies that the risk faced by firm 1 has a larger systematic component, thus the larger risk premium.

The left panel of Figure 4.1 plots the risk premium ratio $(s_i^{\mathbb{Q}} - s_i^{\mathbb{P}})/s_i^{\mathbb{Q}}$ as a function of the systematic risk proportion,

$$\mathrm{SRP}_i = \frac{\beta_i^2 \sigma_{\mathrm{M}}^2}{\sigma_i^2}, \tag{4.2.7}$$

the proportion of the firm's volatility that is explained by the firm's exposure to systematic risk. All else equal, the relationship between the firm's SRP and the proportion of its spread that is due to a risk premium is clearly monotonic and positive.

Stock Options

Focusing on the stock as the underlying, the Black and Scholes (1973) model gives the price of a stock option assuming that the stock price follows a geometric Brownian motion just like that in Equation (4.2.1). Geske (1979) sheds further light on stock options by allowing the instantaneous equity volatility vary with the market leverage of a firm. Geske's model considers stock options as compound options on the fundamental value of the firm described in Equation (4.2.1). This model has rich implications and, in particular, allows us to better understand the links between

Figure 4.1: Risk Premia and Systematic Risk

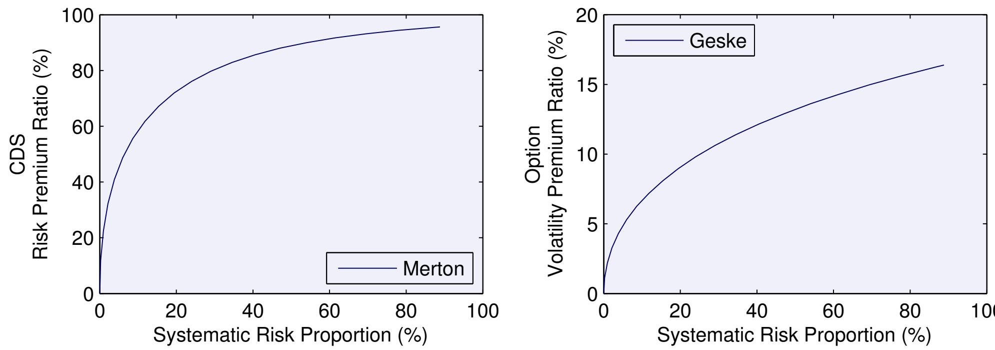

We consider firms with initial value 100, 30% volatility, and paying no dividends. The firms have a debt-to-firm value ratio of 50% and the debt has a maturity of 5 years. The risk-free rate is set at 2.88%, the equity risk premium $\mathrm{E}\left[R_{\mathrm{M}}-r\right]$ *is set at 8.32% and the volatility of the market is set at 12.13%; these numbers were obtained using 3-month Treasury Bill rates and S&P 500 data between January 2003 and June 2007. The total spread,* $s^{\mathbb{Q}}$*, implied by the Merton model is 1.48% and the Black-Scholes implied volatility of the call as priced under Geske's model is* 51.43%*. The left panel of the figure reports the Merton-implied CDS Risk Premium Ratios,* $\mathrm{RPR}_i=(s_i^{\mathbb{Q}}-s_i^{\mathbb{P}})/s_i^{\mathbb{Q}}$*, as a function of the firms' Systematic Risk Proportion,* $\mathrm{SRP}_i=\beta_i^2\sigma_{\mathrm{M}}^2/\sigma_i^2$*. The right panel of the figure reports the Volatility Premium,* $\mathrm{VPR}_i=(\mathrm{IV}_i^{\mathbb{Q}}-\mathrm{IV}_i^{\mathbb{P}})/\,\mathrm{IV}_i^{\mathbb{Q}}$*, as a function of* SRP_i*.*

credit and stock option markets. Derivatives in both markets ultimately depend on the same fundamental firm value.

Consider a call option on the stock that expires at $T_1<T_2$ and has strike K. The option's value under the Geske model is

$$\begin{aligned}
C &= e^{-rT_1}\,\mathrm{E}^{\mathbb{Q}}\left[\Big(E^{\mathbb{Q}}(V_{T_1})-K\Big)\cdot I_{V_{T_1}\geq V^*}\right]\\
&= e^{-rT_1}\,\mathrm{E}^{\mathbb{Q}}\left[V_{T_2}\cdot I_{V_{T_2}>F;V_{T_1}>V^*}\right]\\
&\qquad -Fe^{-rT_2}\,\mathrm{E}^{\mathbb{Q}}\left[I_{V_{T_2}>F;V_{T_1}>v^*}\right]-Ke^{-rT_1}\,\mathrm{E}^{\mathbb{Q}}\left[I_{V_{T_1}>V^*}\right]\\
&= V_0M(a_1,b_1,\rho)-Fe^{-rT_2}M(a_2,b_2,\rho)-Ke^{-rT_1}N(a_2),
\end{aligned}\qquad(4.2.8)$$

where the a_i and b_i are similar to Merton's d_i. Applying a reasoning similar to that which led to Equation (4.2.2), one can obtain a physical value for the call under Geske's framework, with $a_i^{\mathbb{P}}$ and $b_i^{\mathbb{P}}$ replacing a_i and b_i, while accounting for the objective growth rate of the firm value, μ.

The right panel of Figure 4.1 plots the volatility premium ratio $(\mathrm{IV}_i^{\mathbb{Q}}-\mathrm{IV}_i^{\mathbb{P}})/\,\mathrm{IV}_i^{\mathbb{Q}}$ as a function of the systematic risk proportion, SRP_i. Once again, all else equal, the relationship between the firm's SRP and the proportion of its implied volatility that is due to a risk premium is clearly monotonic and positive.

4.2.2 Empirical Implications

Figure 4.1 is obtained in a highly controlled environment. The initial firm value, the firm volatility, the debt-to-firm value ratio are kept constant through the cross section of firms. To better understand the empirical implications of the CAPM-augmented Merton and Geske models with respect to risk premia, we recast the analysis behind Figure 4.1 in settings that are closer to what is observed in our data.

In the empirical section of this paper, we study CDS and stock options written on roughly 130 firms that were part of the CDX index between January 2003 and June 2007. Data are aggregated monthly for a total of 3862 firm/month observations. In the two upper panels of Figure 4.2, we grouped these observations according to the firm's rating on a given month. In the upper-left panel, we computed the ratingwise average of firms' equity volatility; the solid line reports the average and the grey-shaded area form a two-standard error band around the average. As expected, volatility increases as ratings deteriorate.

The upper-right panel of Figure 4.2 repeats the analysis of the upper-left panel, but with the systematic risk proportion (SRP). While the volatility increases with decreasing ratings, the SRP displays the opposite pattern. That is, the volatility of highly rated firms tend to be explained in a larger proportion by market volatility.

We repeat the exercise of Figure 4.1, but considering seven hypothetical firms, one per rating, with volatilities and SRP that mimic those observed in the upper panels of Figure 4.2. Since volatility now varies across the ratings, the total spread implied by the Merton model also varies, increasing with deteriorating ratings as depicted in the mid-left panel of Figure 4.2. Nonetheless, as depicted by the mid-right panel of Figure 4.2, the intuition from Figure 4.1's left panel remains: the higher the proportion of systematic risk, the greater the ratio of the risk premium to total spread. However, the exercise points at potential difficulties in using total spreads as a proxy for the risk premium component, an issue that will become apparent in the empirical section of this paper. The stock option implications of the Geske model are analogous. The lower-left panel of Figure 4.2 illustrates that the predicted implied

Figure 4.2: Another Look at Risk Premia

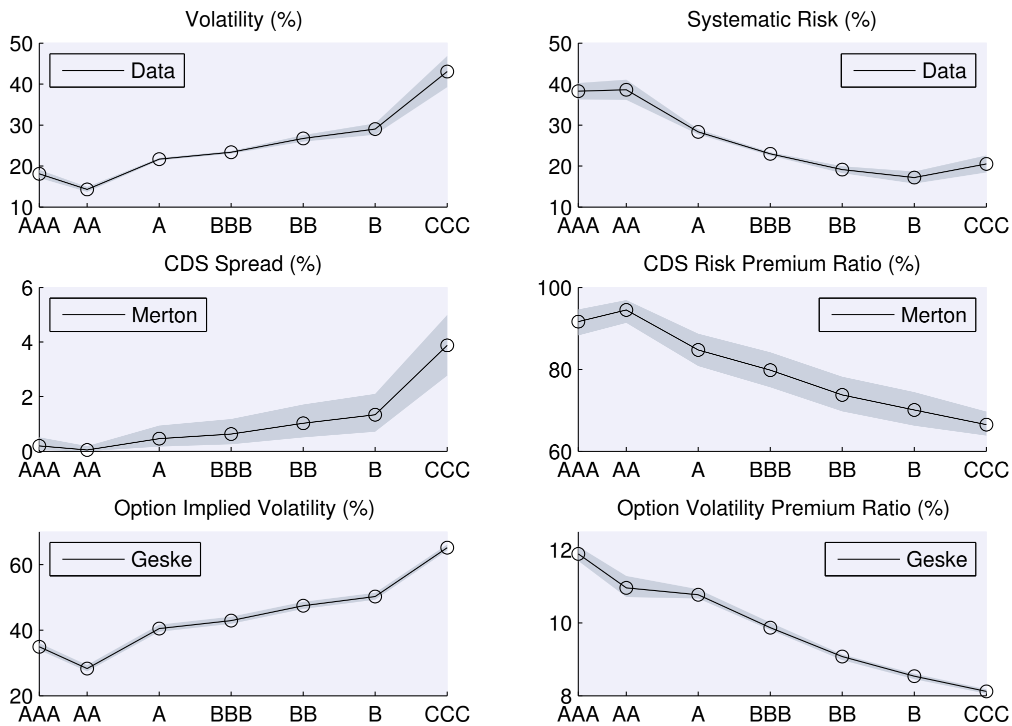

*We grouped the 3862 firm/month observations according to the firm's rating on a given month. In the upper-left panel, the solid line reports the ratingwise average of the firms' equity volatility and the grey shaded area form a two-standard error band around the average. The upper-right similarly reports the ratingwise average of the firms' systematic risk proportion (*SRP*). The mid-left panel reports, for each rating, the Merton spread of a firm with volatility and systematic risk proportion set to those reported in the upper panels; representative firms on the solid line have a debt-to-firm value ratio (d) of 50% and the grey area reports the ranges obtained by letting d vary from 40% to 60%. The mid-right panel reports the corresponding risk premium proportions. Finally, the lower-left panel reports, for each rating, the implied volatility from the Geske price of an at-the-money call with a grey band illustrating the sensitivity of this implied volatility to the call's moneyness (ranging from 85% to 115%). The lower-right panel reports the proportion of this implied volatility that is due to the volatility premium.*

volatility of an at-the-money call option increases with decreasing ratings, which is consistent with the increasing volatility and the positive vega of stock options. Yet, the proportion of implied volatility that is explained by the volatility premium is increasing with the systematic risk proportion, and thus decreasing with ratings.

These "calibrated comparative statics" for stock options are analogous to what we find for credit spreads. We thus find analogous implications for option and credit markets: greater proportional systematic volatilities are related to greater risk premia in the cross-section. Next, we seek to validate these implications in our sample of stock option and CDS data for the constituent firms of the CDX index.

4.3 Empirical Results

In this paper, we study CDS and stock options written on roughly 130 firms that were part of the CDX index between January 2003 and June 2007. Data is aggregated monthly for a total of 3862 firm/month observations.

4.3.1 Firm Characteristics

Figure 4.3 provides a visual summary of the firms in our dataset. The top two panels remind us that although overall risk, as measured by equity volatility, is greater for the lower rating categories, the better rated firms are proportionally more exposed to systematic risk. Levered betas appear stable across rating categories, with the exception of the lowest rated firms for which greater leverage increases the beta. There is no clear pattern in asset volatility across rating categories, nor for book and market leverage ratios.

4.3.2 Stock Option Empirics

Duan and Wei (2009) demonstrate that systematic risk has an impact on stock option prices. They find that, after controlling for the total volatility of the option's underlying, a higher proportion of systematic variance in the stock's total variance leads to higher implied volatility levels. As seen above in Section 4.2, this result emerges as a natural implication of the Geske (1979) model once the model is augmented with the CAPM.

We start our empirical analysis by considering whether Duan and Wei's results hold in our broader and more recent data set. Figure 4.4 summarizes the variable of interest in this regard. The upper-left panel of the figure reproduces the systematic risk pattern across ratings that we observed in Figures 4.2 and 4.3. The upper-right panel of Figure 4.4 reports the volatility premium, $\mathrm{IV}^{\mathbb{Q}}_{t,j} - \sigma_{t,j}$, averaged through time and over firms in a given rating. The positive relationship between the systematic

Figure 4.3: Firm Characteristics

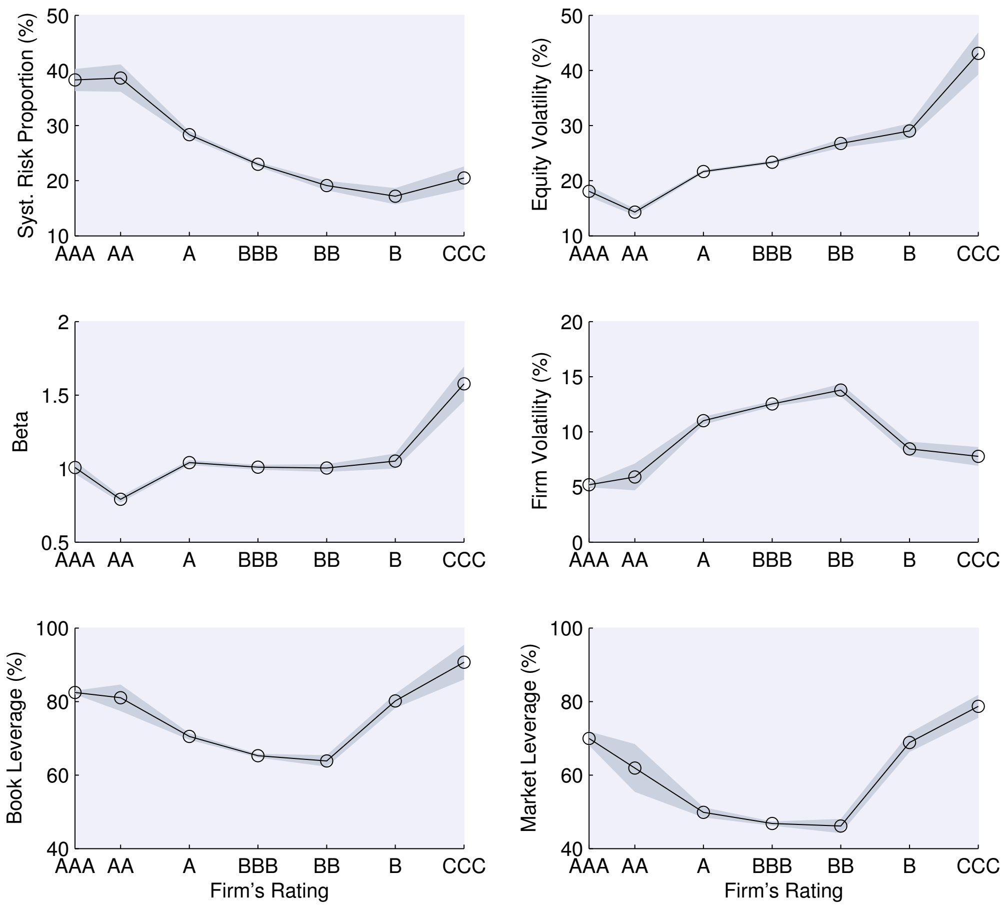

We grouped the 3862 firm/month observations according to the firm's rating on a given month. Each panel in this figure reports the ratingwise average of a given variable for the firms with the given rating. A two-standard error band around the average is also displayed.

risk proportion (SRP) and $\mathrm{IV}_{t,j}^{\mathbb{Q}} - \sigma_{t,j}$ is far from obvious by only considering this figure.

Yet, Table 4.1 reports that the result which Duan and Wei obtained using options written on 30 firms between January 1991 and December 1995 actually holds on a broader and more recent data set: out-of-the-money (OTM) calls and puts quoted on the stocks of 129 firms between January 2003 and June 2007. We reproduce exactly Duan and Wei's methodology here. For each month in our data set, we compute the average implied volatility of OTM calls and puts in a given moneyness/maturity bucket. This procedure is repeated for each firm for which we have data in the month

Figure 4.4: Option Data

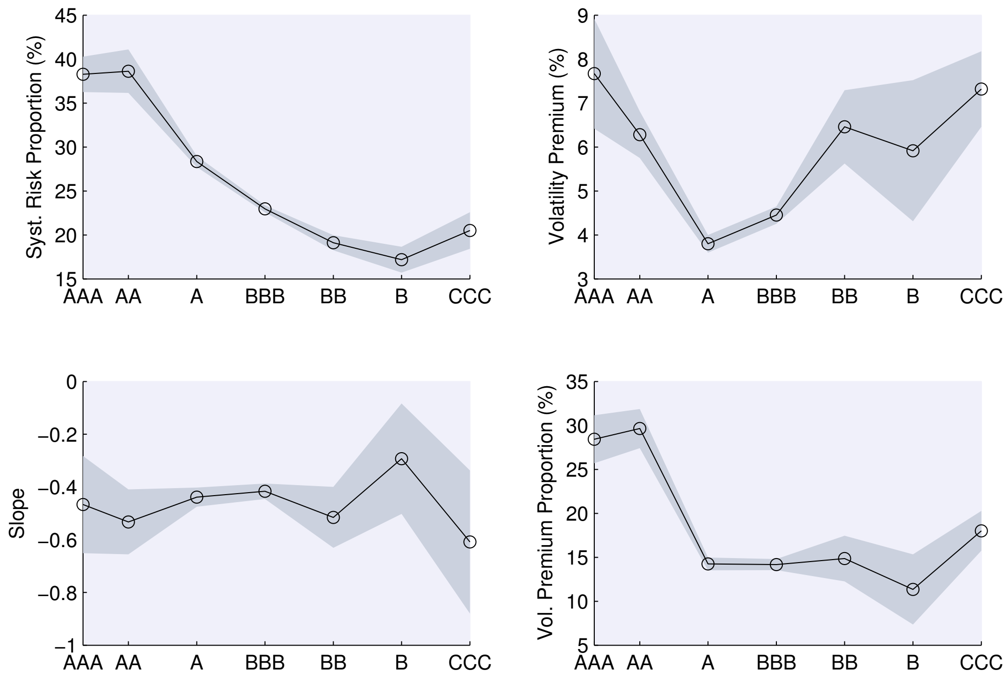

We grouped the 3862 firm/month observations according to the firm's rating on a given month. Each panel in this figure reports the ratingwise average of a given variable for the firms with the given rating. A two-standard error band around the average is also displayed.

under consideration. We also compute the historical volatility and systematic risk proportion of each of these firms using a rolling window of 12 months of observations.

Each month t, this gives us a cross-section of implied volatility, volatility and SRP measures. We then regress the volatility premium on systematic risk proportion, i.e.

$$IV_{t,j} - \sigma_{t,j} = \eta_{0,t} + \eta_{\mathrm{SRP},t}\,\mathrm{SRP}_{t,j} + \varepsilon_{t,j}\,. \tag{4.3.9}$$

We perform such a cross-sectional regression each month, *à la* Fama and MacBeth (1973), and report in the first two columns of Table 4.1 the average coefficient on systematic risk and its Newey and West (1987) standard error, correcting for the high autocorrelation in the coefficient time series. While they are not as striking, our findings do confirm Duan and Wei's results.

Consider two types of firms, the first with no exposure to systematic risk (SRP = 0), which we will refer to as "idiosyncratic" firms, and the second with no idiosyncratic risk (SRP = 1), which we will refer to as a "systematic" firms. The average value of

$\eta_{\mathrm{SRP},t}$ in the first column of Table 4.1 implies that options on a systematic firm would embed a volatility premium from 0.9% to 5.8% than those written on the stock of an idiosyncratic firm. Statistical significance does not obtain for all moneyness/maturity buckets however. The proportion of SRP coefficients that are greater than zero in Duan and Wei's analysis varies from 71.4% and 100% in the settings of our Table 4.1; we obtain proportions that range from 50% to 74.1%. The R^2 in their analysis ranges from 4.2% to 23.1%; in ours, from 2.3% to 8.3%.

Table 4.1: Fama-MacBeth Regressions of Volatility Premia on Systematic Risk Proportions

		Duan and Wei				Vol. Premium Ratios			
		Avg. Coeff	**t-stat**	**Coeff > 0**	$\mathbf{R^2}$	**Avg. Coeff**	**t-stat**	**Coeff > 0**	$\mathbf{R^2}$
Moneyness	**All maturities**	0.0348	2.098	68.5%	3.5%	0.2436	4.523	83.3%	5.7%
(K/S)	**Short-term**	0.0375	2.024	70.4%	3.8%	0.2553	4.491	83.3%	6.1%
0.90-0.95	**Medium-term**	0.0239	1.170	60.0%	7.3%	0.2345	3.277	75.6%	8.4%
	Long-term	0.0579	3.273	70.2%	8.3%	0.3204	4.595	85.1%	10.6%
Moneyness	**All maturities**	0.0412	3.004	74.1%	3.0%	0.2447	4.112	83.3%	4.1%
(K/S)	**Short-term**	0.0386	2.457	70.4%	3.1%	0.2431	3.782	79.6%	3.9%
0.95-1.00	**Medium-term**	0.0532	3.059	67.4%	5.9%	0.2669	3.386	73.9%	6.6%
	Long-term	0.0449	2.944	61.7%	7.5%	0.2783	3.405	76.6%	9.8%
Moneyness	**All maturities**	0.0193	1.660	66.7%	2.3%	0.1276	2.447	68.5%	2.6%
(K/S)	**Short-term**	0.0190	1.437	66.7%	2.4%	0.1269	2.220	66.7%	2.8%
1.00-1.05	**Medium-term**	0.0087	0.548	50.0%	5.4%	0.0828	1.547	57.7%	5.9%
	Long-term	0.0219	1.910	68.5%	5.1%	0.1639	2.117	70.4%	5.7%
Moneyness	**All maturities**	0.0234	1.952	61.1%	2.5%	0.1416	2.471	64.8%	3.1%
(K/S)	**Short-term**	0.0208	1.501	63.0%	2.7%	0.1342	2.280	66.7%	3.3%
1.05-1.10	**Medium-term**	0.0367	1.905	60.0%	6.6%	0.2452	1.725	60.0%	6.5%
	Long-term	0.0233	2.375	57.4%	4.9%	0.1229	3.106	57.4%	5.6%

On each month t, for each subset of options quoted on firm j, we compute the average implied volatility, $IV_{t,j}$. Likewise, we obtain the level of historical volatility $\sigma_{t,j}$ for firm j during month t. Then, in a Fama-MacBeth fashion, we perform the following regression

$$IV_{t,j} - \sigma_{t,j} = \eta_{0,t} + \eta_{\mathrm{SRP},t}\,\mathrm{SRP}_{t,j} + \varepsilon_{t,j}$$

each month. A given cross-section (month/option subset) is neglected if less than 10 observations are available. We report the average of each coefficient $\eta_{\cdot,t}$ time series, and t-statistics that are computed using Newey-West standard errors with lag 3 in order to account for the likely autocorrelation and heteroskedasticity in the series of coefficients. We also report the proportion of $\eta_{\mathrm{SRP},t}$ coefficients that are positive (Coeff > 0) as well as the average of the monthly R^2 statistics.

There are many reasons why this was to be expected. First and foremost, the data set we consider is by far more heterogeneous than considered by Duan and

Wei. Their data set comprises of the S&P 100 and its 30 largest component stocks. The firms' average annualized volatility ranges from 6.4% to 16.2%. Our data set here comprises of 129 firms with ratings varying from AAA to CCC, with average volatility levels that range from 15.3% (General Mills Inc.) to 51.6% (Visteon Corp.) annually.*

The Geske model implies that the higher the systematic risk proportion of a firm, the greater the *proportion* of its implied volatility that will be explained by the volatility premium.† Given that the firms in our data set have such a wide range of volatility levels, we revisit Duan and Wei's regressions and use volatility premium ratios (VPR) rather than absolute premia, i.e.

$$\frac{IV_{t,j} - \sigma_{t,j}}{IV_{t,j}} = \eta_{0,t} + \eta_{\mathrm{SRP},t}\,\mathrm{SRP}_{t,j} + \varepsilon_{t,j} \qquad (4.3.10)$$

The lower-right panel of Figure 4.4 reports the ratingwise average proportion of implied volatilities that is explained by the volatility premium, i.e. the average of $\mathrm{VPR}_{t,j} = \left(\mathrm{IV}_{t,j} - \sigma_{t,j}\right)/\,\mathrm{IV}_{t,j}$ for firms with a given rating. The low-rating averages are scaled down (compared to those in the upper-right panel), reflecting the fact that these firms display higher implied volatility levels on average. The fifth to eighth columns of Table 4.1 report the results of this new regression. The relationship between volatility premium ratios and systematic risk is positive for all moneyness/-maturity buckets, highly statistically significant for OTM puts and slightly less so for OTM calls. Proportions of positive $\eta_{\mathrm{SRP},t}$ coefficients range from 73.9% to 85.1% of OTM puts and from 57.4% a 68.5% for OTM calls. As for economic significance, the proportion of the implied volatility explained by volatility premium would be 23.5% to 32.0% higher for OTM put options written on a systematic firm than for similar options written on an idiosyncratic firm. For OTM call options, the difference ranges between 8.3% and 24.5%. The relationship between systematic risk and the volatility premium seems to be of greater statistical and economic significance

*The heterogeneity of our data set as compared to that of Duan and Wei (2009) might also explain why, while we find their results with regards to the level effect to hold, we couldn't find a significant effect in our data set using the slope proxy suggested in Duan and Wei (2009).

†In a comparative static setting, this holds true also for the difference between volatilities $IV_{t,j} - \sigma_{t,j}$. However, recall Figure 4.3, where the implied-volatility ratio decreases as credit quality declines. This relationship when plotted on the basis of the difference rather than the ratio, yields a more or less flat relationship. This suggests that given the heterogeneity in the data, it should be easier to identify a relationship based on ratios.

on OTM puts than on OTM calls, which could be a consequence of the slope effect documented by Duan and Wei (2009). However, we couldn't find a significant slope effect in our data set using the slope proxy suggested in Duan and Wei (2009); as discussed above, the heterogeneity of our data set as compared to that of Duan and Wei (2009) might explain why we can't reproduce their results with regards to the slope effect.

In Table 4.2, we focus, on the OTM puts ($0.90 \leq K/S < 0.95$) with short or long maturities and assess the robustness of the positive relationship between systematic risk and the volatility premium by control for firm characteristics. The constant in Equation (4.3.10) is replaced by dummy variables controlling for the firm's rating. When only the firm's rating is controlled for, short-term options on the stock of a systematic firm will embed a 27.5% higher volatility premium proportion than those on an idiosyncratic firm; for long-term options, the difference rises to 28.8%. We further control for the annualized conditional equity volatility (as mesured by standard deviation of daily stock returns over the last 12 months), firm size (using the logarithm of market capitalization), and the firm's market leverage. A systematic firm would have 11.0% higher volatility premium ratios on its short-term options and 7.4% higher on its long-term options than would an otherwise equal idiosyncratic firm. Interestingly, after controlling for firm characteristics, the statistical and economic significance of the volatility premium in stock option prices is greater for short-term options. This is interesting and suggestive in light of the findings in Elton, Gruber, Agrawal, and Mann (2001) who show that expected losses in corporate bond spreads are negligible for short term bonds. This implies that risk premia and other sources of compensation in bond spreads need to be relatively larger, and thus likely easier to discern empirically.

In summary, we find that options on stocks with more systematic volatility are relatively more expensive, consistent with the findings of Duan and Wei (2009). In turn this is consistent with the Geske (1979) model augmented by the CAPM. We now turn our attention to the predictions obtained under the framework of Merton for the credit derivatives markets.

Table 4.2: Regressing the Stock Option Relative Volatility Premium on Systematic Risk

		Systematic Risk Only			Controls Only			Full		
Short-Term Options	**#Obs**	**Avg.Coeff**	**t-stat**	**Coeff > 0**	**Avg.Coeff**	**t-stat**	**Coeff > 0**	**Avg.Coeff**	**t-stat**	**Coeff > 0**
AAA	77	23.8650	13.6920		22.7950	12.9160		21.5950	12.5880	
AA	123	19.6220	13.6750		17.6330	11.5390		16.5520	11.1350	
A	1498	13.6440	12.8930		13.9100	18.2600		13.4000	17.0280	
BBB	2300	14.3480	14.9860		17.8410	24.6500		17.8700	24.3320	
BB	524	16.6760	11.8280		19.3350	11.0160		20.0800	12.4470	
B	167	10.9250	1.7407		18.0650	3.2719		18.8800	3.5023	
CCC	78	19.3420	3.2475		23.7290	16.5530		24.3010	12.8200	
Systematic Risk		0.2750	4.8937	92.5%				0.1096	3.1287	82.5%
Equity Volatility					-1.3789	-9.3866	0.0%	-1.3488	-9.5582	0.0%
Firm Size					-2.0022	-3.0511	17.5%	-1.9467	-3.0977	20.0%
Market Leverage					-0.0101	-0.4423	52.5%	-0.0165	-0.6811	55.0%
R^2				18.7%			52.1%			53.6%
Long-Term Options										
AAA	77	16.6500	4.2475		17.6130	5.1917		14.4740	4.0712	
AA	123	15.3680	8.0016		15.0510	7.4742		14.2150	7.6107	
A	1498	11.1100	13.0830		12.1160	11.3620		11.5820	10.5420	
BBB	2300	10.6550	8.8765		16.4230	12.7370		16.3660	12.2070	
BB	524	12.7940	3.1613		13.9290	3.1874		13.0680	3.1480	
B	167	15.4470	2.3814		19.5430	3.3216		21.4080	3.7335	
CCC	78	18.4620			28.3670			28.4670		
Systematic Risk		0.2877	4.1817	81.1%				0.0741	1.3386	59.5%
Equity Volatility					-1.3443	-11.5650	2.7%	-1.2932	-10.3820	2.7%
Firm Size					-2.4418	-3.3509	21.6%	-2.1788	-2.9109	24.3%
Market Leverage					-0.0072	-0.2680	54.1%	-0.0103	-0.3875	56.8%
R^2				34.3%			68.9%			72.3%

In a Fama-MacBeth fashion, we perform the following regression

$$(\mathrm{IV}^{\mathbb{Q}}_{t,j} - \sigma_{t,j})/\,\mathrm{IV}^{\mathbb{Q}}_{t,j} = \eta_{\mathrm{AAA},t}\mathrm{AAA}_{t,j} + \eta_{\mathrm{AA},t}\mathrm{AA}_{t,j} + \ldots + \eta_{\mathrm{CCC},t}\mathrm{CCC}_{t,j} + \eta_{\mathrm{SRP},t}\,\mathrm{SRP}_{t,j} + \eta_{c,t}\cdot \textsc{controls}_{t,j} + \varepsilon_{t,j}$$

each month t, for all firms j. A given cross-section (month/firms subset) is neglected if less than 10 observations are available. We use rating dummies in lieu of a constant. Dummy $\textsc{aaa}_{t,j}$ is equal to 1 if firm j's rating is AAA and 0 otherwise; the same applies for all ratings. All regressors but dummies are centered. We report the average of each coefficient $\eta_{\cdot,t}$ time series, and t-statistics that are computed using Newey-West standard errors with lag 3 in order to account for the likely autocorrelation and heteroskedasticity in the series of coefficients. We also report the proportion of $\eta_{\cdot,t}$ coefficients that are positive (Coeff > 0) as well as the average of the monthly R^2 statistics. Note that, in the regressions, volatility premium ratios, systematic risk ($\mathrm{SRP}_{t,j}$), equity volatility and market leverage regressors appear in percentage terms. We use the natural logarithm of market capitalization to control for firm size.

Figure 4.5: CDS Data

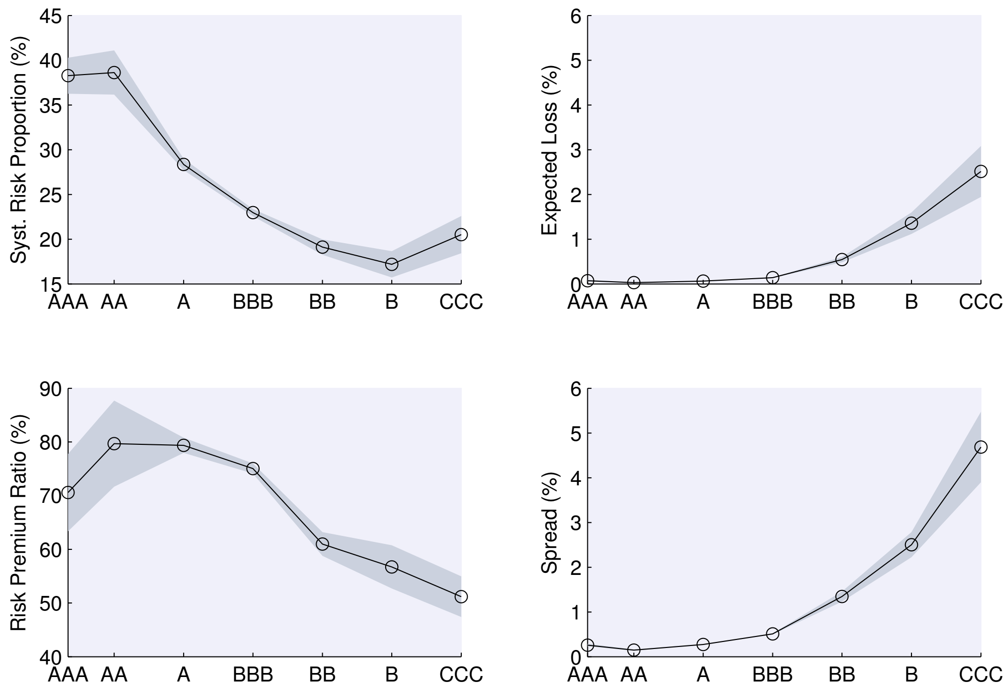

We grouped the 3862 firm/month observations according to the firm's rating on a given month. Each panel in this figure reports the ratingwise average of a given variable for the firms with the given rating. A two-standard error band around the average is also displayed.

4.3.3 CDS Empirics

The Merton model, which is internally consistent with the Geske model, provides an analogous intuition for the relationship between systematic risk and risk premia observed in credit markets. The upper-left panel of Figure 4.5 reproduces the systematic risk pattern across ratings that we observed in previous figures. The upper- and lower-right panels report on spreads and expected losses; firm's with better ratings have lower probability of default, thus requires a lower compensation for expected loss and, overall, have lower spreads. The lower-left panel however tells us that a larger proportion of the spreads of high-rated firms, which have larger systematic risk proportions, is due to the risk premium component of spreads. This is preliminary evidence that the empirical prediction we obtained under Merton's framework holds in the data. It seems that, across rating categories, firms with a greater proportion of systematic volatility also have higher proportions of risk premia in their spreads.

It is tempting to interpret the prediction of the Merton model as implying that for two otherwise equal firms, the one with greater exposure to systematic risk should require larger spreads. Table 4.3 clearly shows that this is not the case. In this table, we repeat the regression of Table 4.2, but with spreads as dependent variable. The relationship between spreads and systematic risk is robustly negative. Li (2008) also obtains this result and finds this negative relationship to be robust.* To explain her findings, Li suggests an explanation based on transaction costs. She argues that contracts on a firm with a lower systematic risk proportion are subject to higher transaction or hedging costs for the seller; the more closely a firm's value moves with the market, she says, the easier it is for the seller to find cheap hedging instruments. Thus, higher hedging costs would justify higher CDS prices.

We suggest an alternative explanation that appeals to the decomposition of a spread into its expected loss and its risk premium components. In recent work on structured credit products, Coval, Jurek, and Stafford (2009) argue that, for a given expected loss, a fixed income product that is more likely to default in bad states of the economy should require a relatively larger risk premium than an otherwise identical product more likely to default in good states.

Rephrasing this insight and assuming that the value of the market portfolio proxies for the state of the economy, firms with higher systematic risk proportions should require a higher risk premium and thus higher spreads. However, there is an endogeneity issue here; in light of Figure 4.5, firms with larger systematic risk proportions are also likely to be well established firms with lower expected losses, and thus lower spreads. In addition, risk premia depend in a nonlinear fashion on leverage and volatility. To deal with these issues, we propose decomposing spreads into their expected loss and risk premium components before running our regression analysis. The appendix provides a summary of our methodology. In essence we rely on physical default probabilities constructed using a structural model to compute "physical" bond prices as in Section 4.2. We then use the ratio of the spread risk premium component to total spread as dependent variable.

Table 4.4 reports the results of the same regressions as Table 4.3, but with risk premium ratios as the dependent variable. The proportion of 5-year spreads that

*Li's dataset was CDS spreads provided by GFI for the period of January 2000 to December 2004.

Table 4.3: Regressing 5-Year and 10-Year CDS Spreads on Systematic Risk

		Systematic Risk Only			Controls Only			Full		
5-years Spreads	**#Obs**	**Avg.Coeff**	**t-stat**	**Coeff > 0**	**Avg.Coeff**	**t-stat**	**Coeff > 0**	**Avg.Coeff**	**t-stat**	**Coeff > 0**
AAA	69	32.6950	8.6892		27.9290	5.2499		32.3990	5.5330	
AA	63	24.3880	6.9863		36.0490	6.1520		42.5830	6.1496	
A	989	31.9510	8.2692		33.5100	11.4930		36.7670	19.8690	
BBB	2088	49.6850	18.7460		50.6870	18.0540		51.3000	21.3830	
BB	419	129.3900	15.4520		121.1500	15.4970		119.8400	17.2430	
B	156	231.5500	7.0241		204.5200	7.0774		200.1700	7.2476	
CCC	78	484.9500	6.7666		382.9000	7.6724		380.0600	7.8278	
Systematic Risk		-0.6549	-3.8208	11.9%				-0.5582	-4.4141	9.5%
Equity Volatility					2.5825	10.1950	100.0%	2.4848	11.2960	100.0%
Firm Size					-3.4872	-1.0997	45.2%	-3.1994	-1.0537	45.2%
Market Leverage					0.7382	5.7391	100.0%	0.7650	5.7849	100.0%
R^2				70.1%			77.4%			78.0%
10-years Spreads										
AAA	69	45.6810	9.4950		44.2790	7.5673		48.8850	7.6226	
AA	63	35.8060	8.8044		51.9610	7.4482		58.7040	7.4034	
A	989	48.3120	13.1690		50.6270	13.1480		54.0300	20.4910	
BBB	2088	74.0690	37.1720		74.7010	22.0010		75.3220	27.6180	
BB	419	166.0600	12.9890		156.4600	13.3440		155.1500	14.3080	
B	156	271.6400	7.8111		245.0500	7.8835		240.7300	8.0681	
CCC	78	508.0000	7.5336		407.8400	8.3551		405.0400	8.5459	
Systematic Risk		-0.6990	-3.9635	9.5%				-0.5672	-4.6092	7.1%
Equity Volatility					2.6098	11.0650	100.0%	2.5076	12.9580	100.0%
Firm Size					-5.2354	-1.3076	45.2%	-4.9840	-1.2825	45.2%
Market Leverage					0.6899	5.2255	95.2%	0.7164	5.3713	97.6%
R^2				70.8%			78.1%			78.6%

On each month t, using 5-year (upper panel) and 10-year (lower panel) CDS spreads quoted on firm j, we compute the average spread, $s_{t,j}$. Then, in a Fama-MacBeth fashion, we perform the following regression

$$s_{t,j} = \eta_{\mathrm{AAA},t}\mathrm{AAA}_{t,j} + \eta_{\mathrm{AA},t}\mathrm{AA}_{t,j} + \ldots + \eta_{\mathrm{CCC},t}\mathrm{CCC}_{t,j} + \eta_{\mathrm{SRP},t}\,\mathrm{SRP}_{t,j} + \eta_{c,t} \cdot \mathrm{CONTROLS}_{t,j} + \varepsilon_{t,j}$$

each month. A given cross-section (month/firms subset) is neglected if less than 10 observations are available. We use rating dummies in lieu of a constant. Dummy $\mathrm{AAA}_{t,j}$ is equal to 1 if firm j's rating is AAA and 0 otherwise; the same applies for all ratings. All regressors but dummies are centered. We report the average of each coefficient $\eta_{\cdot,t}$ time series, and t-statistics that are computed using Newey-West standard errors with lag 3 in order to account for the likely autocorrelation and heteroskedasticity in the series of coefficients. We also report the proportion of $\eta_{\cdot,t}$ coefficients that are positive (Coeff > 0) as well as the average of the monthly R^2 statistics. Note that, in the regressions, systematic risk ($\mathrm{SRP}_{t,j}$), equity volatility and market leverage regressors appear in percentage terms, while spreads appear in basis points. We use the natural logarithm of market capitalization to control for firm size.

Table 4.4: Regressing 5-Year and 10-Year CDS Risk Premium Ratios on Systematic Risk

		Systematic Risk Only			Controls Only			Full		
5-years RP Ratios	**#Obs**	**Avg.Coeff**	**t-stat**	**Coeff > 0**	**Avg.Coeff**	**t-stat**	**Coeff > 0**	**Avg.Coeff**	**t-stat**	**Coeff > 0**
AAA	69	71.1750	16.3440		69.8300	18.8470		68.8640	18.5950	
AA	63	78.0250	22.9170		82.8880	31.3480		79.0140	19.7230	
A	989	78.4400	43.3030		78.6900	49.1720		78.7740	54.6200	
BBB	2088	74.2470	35.5170		74.9310	37.8690		74.9930	35.3220	
BB	419	63.8030	20.2570		61.2960	23.9250		63.7530	18.3640	
B	156	61.5670	31.2800		51.0300	22.6460		55.7820	31.2550	
CCC	78	59.7720	18.6710		36.1340	7.6636		40.5720	9.2579	
Systematic Risk		0.3464	2.0813	64.3%				0.3113	1.8289	61.9%
Equity Volatility					0.2180	1.3647	64.3%	0.2044	1.1770	61.9%
Firm Size					0.4720	0.2948	59.5%	-0.0490	-0.0312	57.1%
Market Leverage					0.2408	9.8879	97.6%	0.2107	10.0770	97.6%
R^2				18.9%			23.8%			28.7%
10-years RP Ratios										
AAA	69	96.1690	124.2100		95.5430	55.7070		94.1540	52.2700	
AA	63	95.4690	65.2390		90.6350	62.1990		87.9720	44.2770	
A	989	90.6090	141.4000		88.7710	137.4000		88.4770	145.1200	
BBB	2088	88.3700	141.0400		88.5600	135.6600		88.6800	142.2700	
BB	419	74.8310	70.3470		76.1400	92.7870		77.3420	87.1730	
B	156	54.3720	18.6370		55.2010	15.7290		57.0930	20.0880	
CCC	78	17.4930	2.4051		22.2310	3.7676		23.8970	4.0119	
Systematic Risk		0.2152	2.8335	81.8%				0.2003	3.0878	87.9%
Equity Volatility					0.0536	0.7614	69.7%	0.0832	1.1195	72.7%
Firm Size					4.2562	5.8341	97.0%	4.0540	6.3379	97.0%
Market Leverage					-0.0493	-0.7671	51.5%	-0.0633	-1.0021	51.5%
R^2				56.1%			61.9%			63.2%

*On each month t, using the risk premium ratios (*RPR*) of 5-year (upper panel) and 10-year (lower panel) CDS spreads quoted on firm j, we compute the average* RPR, $s_{t,j}$. *Then, in a Fama-MacBeth fashion, we perform the following regression*

$$\mathrm{RPR}_{t,j} = \eta_{\mathrm{AAA},t}\mathrm{AAA}_{t,j} + \eta_{\mathrm{AA},t}\mathrm{AA}_{t,j} + \ldots + \eta_{\mathrm{CCC},t}\mathrm{CCC}_{t,j} + \eta_{\mathrm{SRP},t}\,\mathrm{SRP}_{t,j} + \eta_{c,t}\cdot \mathrm{CONTROLS}_{t,j} + \varepsilon_{t,j}$$

each month. A given cross-section (month/firms subset) is neglected if less than 10 observations are available. We use rating dummies in lieu of a constant. Dummy $\mathrm{AAA}_{t,j}$ *is equal to 1 if firm j's rating is AAA and 0 otherwise; the same applies for all ratings. All regressors but dummies are centered. We report the average of each coefficient $\eta_{\cdot,t}$ time series, and t-statistics that are computed using Newey-West standard errors with lag 3 in order to account for the likely autocorrelation and heteroskedasticity in the series of coefficients. We also report the proportion of $\eta_{\cdot,t}$ coefficients that are positive (Coeff > 0) as well as the average of the monthly R^2 statistics. Note that, in the regressions,* RPR, *systematic risk* ($\mathrm{SRP}_{t,j}$)*, equity volatility and market leverage regressors appear in percentage terms. We use the natural logarithm of market capitalization to control for firm size.*

is due to the risk premium ranges from 51.2% on average for CCC-rated firms to 79.7% for AA-rated firms. This corroborates the findings of Elton, Gruber, Agrawal, and Mann (2001). The dummies here control for the rating of the different firms. Omitting other controls, a systematic firm would have 34.6% higher risk premium ratio than an idiosyncratic firm on its 5-year spreads, and 21.5% higher on its 10-year spreads. Even after controlling for the level of equity volatility, firm size, and market leverage, the impact of systematic risk remains of economic significance. A systematic firm would have 21.5% higher risk premium ratio than an idiosyncratic firm on its 5-year spreads, and 20.0% higher on its 10-year spreads.

Interestingly firm size is not significant for the risk premium ratios on 5-year spreads, but is significant for RPRs on 10-year spreads. The opposite is true for market leverage. Equity volatility is never significant in explaining risk premium ratios.

4.4 Conclusion

We show that two simple and internally consistent models for option prices and credit spreads provide an intuitive prediction for both markets. Options and credit default swaps on firms with more systematic risk exposure should be more expensive, all else equal. In other words, purchasing insurance on a firm with put options or credit default swaps should be costlier, after controlling for among other things total risk, leverage and risk free interest rates, the greater the systematic risk. This is a simple insight, yet options and credit derivatives are often viewed through the prism of risk-neutral (relative) pricing and, as a result, the wedge between risk-adjusted and physical probability return distributions tends to remain out of view.

Recently, Duan and Wei (2009) have shifted the focus onto the cross-sectional pricing of risk in stock option markets, relating prices to the proportions of systematic risk in equity volatilities. We document that their findings are robust to a more recent and broader dataset, and then ask whether (as our comparative statics suggest) these findings have an analogue in the credit derivative markets. We find this to be the case: the proportion of the price of default insurance due to a risk premium is greater the larger the systematic risk exposure a firm has.

These findings have important implications both in asset allocation and risk management applications. It becomes clear that risk metrics that abstract from any systematic risk adjustment (such as credit ratings) may lead to incentive issues in portfolio selection. If an institution faces a ratings based investment constraint, then it is possible it may favour higher promised yields per rating unit, which in turn will lead to systematic risk concentration. Similarly, an institution facing capital charges on the basis of ratings-defined risk categories, may seek to enhance yield in much the same way.

Appendix — Extracting the Expected Loss Component of Spreads

The Merton (1974) model, the simplest structural model, is rich enough to generate the empirical implications of Section 4.2.2. For expositional purposes, there was thus no reason to consider a more complex model in Section 4.2. When it comes to estimating default probabilities, the Merton model however lacks the flexibility needed to capture the rich structure of the data. We thus use the Leland and Toft (1996) model in the estimation of default probabilities and resulting expected losses.

Here, we briefly describe the Leland and Toft (1996) model; the reader is invited to consult the original paper for details. Since the model does not treat the valuation of credit default swaps (CDS), we also present the building blocks needed to value a CDS. We calibrate our structural model to both equity and CDS spreads to get precise default probabilities.

4.A.1 Overview of the Leland-Toft Model

The fundamental variable in the Leland-Toft model is the value of the firm's assets, which is assumed to evolve as a geometric Brownian motion under the risk-adjusted measure,

$$dv_t = (r - q)v_t dt + \sigma v_t dW_t, \tag{4.A.1}$$

where v is the value of the firm's assets, r is the risk-free interest rate, q is the payout ratio, σ is the volatility of the asset value, and W_t is a Brownian motion under the risk-adjusted measure.

Default is triggered by the shareholders' endogenous decision to stop servicing debt. The exact asset value at which this occurs is determined by several parameters of the model and is denoted by L. The value of the firm differs from the value of the assets by the values of the tax shield and the expected bankruptcy costs. Coupon payments are tax deductible at a rate τ and the realized costs of financial distress amount to a fraction α of the value of the assets in default (i.e. L). In this setting, the value of the firm (F) is equal to the value of assets plus the tax shield (TS) less the costs of financial distress (BK). The value of the firm is, of course, split between equity (E) and debt (D), and thus

$$\begin{aligned} F(v_t) &= v_t + TS(v_t) - BK(v_t) \\ &= E(v_t) + D(v_t). \end{aligned} \tag{4.A.2}$$

4.A.2 Building Blocks for CDS valuation

First, define default as the first time ($\mathcal{T}$) at which the asset value hits the default boundary L from above, so that $\ln\left(\frac{v_{\mathcal{T}}}{L_{\mathcal{T}}}\right)$=0. Then, define $G(v_t, t)$ as the value of a claim paying off \$1 in default,

$$G(v_t, t) = \mathrm{E}^B\left[e^{-r(\mathcal{T}-t)} \cdot 1\right]. \tag{4.A.3}$$

We let E^B denote expectations under the standard pricing measure. The value of G is given by

$$G(v_t, t) = \left(\frac{v_t}{L_t}\right)^{-\theta} \tag{4.A.4}$$

with the constant given by

$$\theta = \frac{\sqrt{(h^B)^2 + 2r} + h^B}{\sigma}, \quad \text{where} \quad h^B = \frac{r - q - \alpha - 0.5\sigma^2}{\sigma}.$$

Define the dollar in default with maturity T, $G(v_t, t; T)$, as the value of a claim paying off \$1 in default if it occurs before T

$$G(v_t, t; T) = \mathrm{E}^B\left[e^{-r(\mathcal{T}-t)} \cdot 1 \cdot (1 - I_{\mathcal{T} \not< T})\right] \tag{4.A.5}$$

and define the binary option $H(v_t, t; T)$ as the value of a claim paying off \$1 at T if default has not occurred before that date

$$H(v_t, t; T) = \mathrm{E}^B\left[e^{-r(\mathcal{T}-t)} \cdot 1 \cdot I_{\mathcal{T} \not< T}\right]. \tag{4.A.6}$$

$I_{\mathcal{T} \not< T}$ is the indicator function for the survival event, i.e. the event in which the asset value (v_T) did not hit the barrier prior to maturity ($\mathcal{T} \not< T$).

The price formulae for the last two building blocks are given below. They contain a term that expresses survival probabilities, the probabilities (under different measures) of the survival event $\mathcal{T} \not< T$. The probabilities of the survival event at t under the probability measures $Q^m : m = \{B, G\}$ are

$$Q^m(\mathcal{T} \not< T) = \mathrm{N}\left(k^m\left(\frac{v_t}{L_t}\right)\right) - \left(\frac{v_t}{L_t}\right)^{-\frac{2}{\sigma}h^m} \mathrm{N}\left(k^m\left(\frac{L_t}{v_t}\right)\right), \tag{4.A.7}$$

where

$$\begin{aligned} k^m(x) &= \frac{\ln x}{\sigma\sqrt{T-t}} + h^m\sqrt{T-t}, \\ h^G &= h^B - \theta\cdot\sigma = -\sqrt{(h^B)^2 + 2\tau}, \end{aligned}$$

and N(k) denotes the cumulative standard normal distribution with integration limit k.

The probability measure Q^G is the measure having $G(v_t, t)$ as numeraire (the Girsanov kernel for going to this measure from the standard pricing measure is $\theta\cdot\sigma$). Using this lemma we obtain the pricing formulae for the building blocks in a convenient form. The price of a down-and-out binary option is

$$H(v_t, t; T) = e^{-r(T-t)}\cdot Q^B(\mathcal{T} \nless T).$$

The price of a dollar-in-default claim with maturity T is

$$G(v_t, t; T) = G(v_t, t)\cdot(1 - Q^G(\mathcal{T} \nless T)).$$

To understand this second formula, note that the value of receiving a dollar if default occurs prior to T must be equal to receiving a dollar-in-default claim with infinite maturity, less a claim where you receive a dollar in default conditional on it not occuring prior to T:

$$G(v_t, t; T) = G(v_t, t) - e^{-r(T-t)}E^B\left[G(v_T, T)\cdot I_{\mathcal{T}\nless T}\right].$$

Using a change of probability measure, we can separate the variables within the expectation brackets,

$$\begin{aligned} G(v_t, t; T) &= G(v_t, t) - e^{-r(T-t)}E^B\left[G(v_T, T)\cdot I_{\mathcal{T}\nless T}\right] \\ &= G(v_t, t)\cdot(1 - Q^G(\mathcal{T} \nless T). \end{aligned}$$

CDS Valuation

A CDS provides insurance for a specified coporate bond termed the reference obligation. The firm issuing this bond is designated as the reference entity. The seller of insurance, the protection seller, promises, should a default event occur, to buy the reference obligation from the protection buyer at par.

The valuation of a CDS involves two parts, the premium paid by the protection buyer to the protection seller and the potential buy-back by the protection seller. Let T denote the maturity of a CDS contract and Q the fee. The value of the premium at time t is

$$\mathrm{E}\left[\int_t^T e^{-r(s-t)} \cdot Q \cdot I_{\mathcal{T} \not< s} \cdot ds\right], \tag{4.A.8}$$

where we let $I_{\mathcal{T} \not< s}$ be the indicator function for nondefault before s and E denotes the expectation under the standard pricing measure. The maturity of the credit default swap is typically shorter than the maturity of the reference obligation (T). In fact, the most common maturity for CDS in practice is $T^* = 5$ years.

Assume that the bondholder in the event of a default recovers a fraction R of the par value P. The second part of the value of a CDS contract is the expected value of receiving the difference between the par value P and the amount recovered,

$$\mathrm{E}\left[e^{-r(T-t)} \cdot (P - R \cdot P) \cdot I_{\mathcal{T} \not< T^*}\right]. \tag{4.A.9}$$

The expectation is conditional on default occurring before maturity of the CDS. Using the previously outlined building blocks, we can formalize the value of a CDS as below.

Assume a CDS involves receiving an amount $P - R \cdot P$ if $\mathcal{T} \not< T$, and paying a continuous premium Q until $\min(T^*, \mathcal{T})$. The value of the CDS is

$$CDS(v_t, t) = (P - R \cdot P) \cdot G(v_t, t; T^*) - \frac{Q}{r}(1 - H(v_t, t; T^*) - G(v_t, t; T^*)).$$

The first term of the CDS formula captures the value of receiving the bond's face value in case of default. The second term captures the cost of paying the premium as a risk-free, infinite stream $\left(\frac{Q}{r}\right)$ less two terms: the first (H) reflecting the discount attributable to the finite maturity of the swap, and the second (G) repoducing the discount due to disrupted payments when and if default occurs.

Typically, the fee is chosen so that the credit default swap upon initiation ($t = 0$) has zero value:

$$Q = \frac{r \cdot (P - R \cdot P) \cdot G(v_t, t; T^*)}{(1 - H(v_t, t; T^*) - G(v_t, t; T^*))}.$$

Intuitively, holding a CDS together with the reference obligation is close to holding the corresponding risk-free bond only. The positions are not identical, however, since the CDS typically has a different maturity and assures its holder the nominal amount (P), rather than value of the risk free bond (B), upon default. Yet, it is often convenient to think of the default swap premium of a just initiated swap as akin to spread on the underlying corporate bond.

4.A.3 Estimation methodology

In the previous section we laid down the pricing formulae for credit default swaps. In this section we discuss issues related to the practical implementation as well as extracting risk premia from credit default swaps. In addition to benchmark term structures, the following inputs are needed to price equity as well as CDS. The recovery rate of the bond R,the total nominal amount of debt, N, coupon C and maturity T. The costs of financial distress α. The tax rate τ. The rate at which earnings are generated by the assets q. Finally the current value, v, and volatility of assets σ.

Since, the recovery rate of the bond in financial distress is not readily observed we set it equal to 40%, roughly consistent with average defaulted debt recovery rate estimates for US entities between 1985-2003. The nominal amount of debt is measured by the total liabilities as reported in COMPUSTAT. For simplicity, we assume that the average coupon paid out to all the firm's debtholders equals the risk-free rate: $c = r \cdot N$. We set the maturity of newly issued debt equal to 6.76 years, consistent with empirical evidence reported in Mauer and Stohs (1994).

Finally, we assume that 15% of the firm's assets are lost in financial distress before being paid out to debtholders and fix the tax rate at 20%. The choice of 15% distress costs lies within the range estimated by Andrade and Kaplan (1998). The choice of 20% is intentionally lower than the corporate tax rate to reflect personal tax benefits to equity returns, thus reducing the tax advantage of debt.

We compute q as the weighted average of net of tax interest expenses (relative to total liabilities (TL)) and the equity dividend yield (DY):

$$q = \frac{IE}{TL} \times lev \times (1 - TR) + DY \times (1 - lev) \tag{4.A.10}$$

where

$$lev = \frac{TL}{TL + MC}$$

where MC denotes the firm's equity market capitalization and TR is the effective tax rate.

Next, we require estimates of asset value and volatility. Since, we have two unknowns (asset value and volatility), we require two equations to solve for our two unknowns. The usual approach is to match the observed equity value to the model-implied equity value and observed equity volatility to model-implied equity volatility. In particular, this approach involves solving backwards from equity value and volatility to get implied asset value and asset volatility, where equity is treated as a call option on the firm's assets. We need to value the underlying CDS contracts precisely. The approach of backing out asset value and volatility from equity value and volatility may not provide precise valuation of a CDS contract for each data point. So, instead of using equity value and volatility to back out asset value and volatility, we use equity value and CDS spread to back out the two unknowns. The advantage of using this approach is that it allows us to match the observed CDS spread and hence, allows to extract the default probabilities under risk-neutral measure precisely.

4.A.4 Estimating default probabilities

Above, we describe our methodology for valuation of CDS and calibration of the structural model. We also need to extract default probabilities under both risk-neutral and physical measure. Below, we describe the methodology to extract default probabilities from the calibrated structural model.

We set the loss rate l equals to 60%, roughly consistent with average defaulted debt recovery rate estimates for US entities between 1985-2003. Previous studies on default risk premium Berndt, Douglas, Duffie, Ferguson, and Schranz (2008), Saita (2006) and Berndt, Lookman, and Obreja (2006) used Expected Default frequencies (EDFs) provided by Moody's KMV as their estimate of the historical default probabilities. In this paper, we estimate company specific default probabilities using the Leland and Toft (1996) model. This methodology yields estimates conceptually similar to EDFs.

Under the physical measure, the survival probabilities $P_t(\tau > T_i)$ are given in closed form by

$$P_t(\tau > T_i) = \mathrm{N}\left(d^P_{T_i}\left(\frac{v_t}{L}\right)\right) - \left(\frac{v_t}{L}\right)^{-2\frac{\mu_v - 0.5\sigma^2}{\sigma^2}} \mathrm{N}\left(d^P_{T_i}\left(\frac{L}{v_t}\right)\right), \qquad (4.A.11)$$

with

$$d^P_{T_i}\left(\frac{v_t}{L}\right) = \frac{\ln\left(\frac{v_t}{L}\right) + \left(\mu_v - 0.5\sigma^2\right)(T_i - t)}{\sigma\sqrt{T_i - t}}$$

$$\text{and} \quad d^B_{T_i}\left(\frac{L}{v_t}\right) = \frac{\ln\left(\frac{L}{v_t}\right) + \left(\mu_v - 0.5\sigma^2\right)(T_i - t)}{\sigma\sqrt{T_i - t}},$$

and where μ_v, the expected return of the asset value under the objective measure, is the only parameter that still needs to be estimated at this point.

Previous studies such as Leland (2004) and Huang and Huang (2002) have used the CAPM beta of the firm multiplied with an average market risk premium figure to provide an estimate of the expected asset return. In contrast, we use the same methodology that we applied above to link the bond risk premium and the equity risk premium. Equity is a contingent claim on the asset value and we can write

$$\mu_v - r = (R_v(t) - r) = \Delta_E \cdot (R_E(t) - r),$$

where

$$\Delta_E = \left(\frac{\frac{\partial E(v_t,t)}{\partial v_t} v}{E(v_t,t)}\right)^{-1}$$

and where $\frac{\partial E(v_t,t)}{\partial v_t}$ is computed using the Leland-Toft model, and $(R_E(t) - r)$ is the estimated equity risk premium.

5 Conclusion

This thesis analyzes three closely interrelated themes of financial derivatives pricing: model risk, macroeconomic risk, and systematic risk. Whichever the source of risk under consideration, the risk being faced is unavoidably quantified using some model. Model risk is thus a very broad concept of which we study only a very specific aspect.

Nonetheless, the study in Chapter 2 is of importance as it sheds light on the relative performance of eight different models divided along three major modeling axes. We demonstrate that the modeling choices being made impact on the economic properties of the different option-valuation models. For example, we find that constraining variance processes to an affine specification puts an undue burden on some of the model's parameters that end up playing multiple economical roles. We also find that omitting to properly model for the long-memory-like properties of volatility significantly impairs the model's ability to capture the variations in options prices, through time, and across maturities and moneyness levels.

Models with two volatility components allow for a richer structure of variance than single-component models. For example, given the current volatility level, a component model's variance forecast can significantly differ depending on whether this current level is above or below the long-run level of variance given by the second component. In Chapter 3, I consider a model where this second component is designed to capture time-varying expectations and relies on observables rather than being latent. In particular, I analyze the impact of changes in business conditions on option prices. The model I introduce outperforms the benchmark in fitting asset returns and in pricing options, especially in the 1990-1991 and 2001 recessions. By letting macroeconomic risk play a significant role in its volatility forecasts and, consequently, in its option price predictions, the model removes the typically observed counter-cyclical patterns in pricing errors.

In Chapter 4, rather than correcting the weaknesses of a given model, we simply revisit fairly simple models, the Merton (1974) and Geske (1979) models, and highlight some their implications for risk premia that have been partly neglected so far. Interestingly, these simple models, augmented with the CAPM, straightforwardly link the risk premia in credit and stock option markets, a link that is due to the systematic risk faced by the firm from which these contingent claims derive their value. However imperfect the Merton and Geske models may be, we find that their predictions regarding systematic risk holds in the data. The greater the systematic risk faced by a firm, the greater the importance of the risk premium in observed spreads and option prices.

Macroeconomic risk is an important determinant of market volatility and inherently contributes to systematic risk, which in turn is a fundamental determinant of asset prices. From a portfolio management perspective, acknowledging and understanding the different sources of risk to which an investment is exposed is a first step of paramount importance in properly managing the risk inherent to the investment. This thesis contributes in deepening, although modestly, our understanding of these risks and of their determinants.

References

Adrian, T. and J. Rosenberg (2008). Stock returns and volatility: Pricing the short-run and long-run components of market risk. *Journal of Finance 63*(6), 2997–3030.

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.

Andrade, G. and S. N. Kaplan (1998). How Costly is Financial (not Economic) Distress? Evidence from Highly Leveraged Transactions that Became Distressed. *Journal of Finance 53*, 1443–1493.

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.

Arora, N., P. Gandhi, and F. A. Longstaff (2009, July). Counterparty Credit Risk and the Credit Default Swap Market. Working Paper.

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.

Bakshi, G., P. Carr, and L. Wu (2008). Stochastic Risk Premiums, Stochastic Skewness in Currency Options, and Stochastic Discount Factors in International Economies. *Journal of Financial Economics 87*(1), 132–156.

Bakshi, G., N. Kapadia, and D. Madan (2003). Stock Returns Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options. *Review of Financial Studies 16*, 101–143.

Barndorff-Nielsen, O. (1997). Processes of Normal Inverse Gaussian Type. *Finance and stochastics 2*(1), 41–68.

Bates, D. (1996). Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options. *Review of Financial Studies 9*, 69–107.

Bates, D. (2000). Post-87 Crash Fears in S&P500 Futures Options. *Journal of Econometrics 94*, 181–238.

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.

Berndt, A., R. Douglas, D. Duffie, M. Ferguson, and D. Schranz (2008). Measuring Default Risk Premia from Default Swap Rates and EDFs. Working paper.

Berndt, A., R. Jarrow, and C. Kang (2007). Restructuring Risk in Credit Default Swaps: An Empirical Analysis. *Stochastic Processes and their Applications 117*, 1724–1749.

Berndt, A., A. A. Lookman, and I. Obreja (2006). Default Risk Premia and Asset Returns. Working Paper, CMU.

Bhamra, H., L.-A. Kuehn, and I. A. Strebulaev (2009). The Aggregate Dynamics of Capital Structrure and Macroeconomic Risk. *Review of Financial Studies*, Forthcoming.

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.

Black, F. and J. C. Cox (1976). Valuing Corporate Securities: Some Effects of Bond Indenture Provisions. *Journal of Finance 31*, 351–367.

Black, F. and M. S. Scholes (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy 7*, 637–654.

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

Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics 31*, 307–327.

Bollerslev, T., M. S. Gibson, and H. Zhou (2009). Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities. *Journal of Econometrics*, Forthcoming.

Bollerslev, T. and H. O. Mikkelsen (1999). Long-Term Equity AnticiPation Securities and Stock Market Volatility Dynamics. *Journal of Econometrics 92*, 75–99.

Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected Stock Returns and Variance Risk Premia. *The Review of Financial Studies 22*, 4463–4492.

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.

Cao, C., F. Yu, and Z. Zhong (2010). The Information Content of Option-Implied Volatility for Credit Default Swap Valuation. *Journal of Financial Markets*, Forthcoming.

Carr, P., H. Geman, D. Madan, and M. Yor (2002). The Fine Structure of Asset Returns: An Empirical Investigation. *The Journal of Business 75*(2), 305–332.

Carr, P. and L. Wu (2003a). The Finite Moment Log Stable Process and Option Pricing. *Journal of Finance 58*(2), 753–777.

Carr, P. and L. Wu (2003b). What Type of Process Underlies Options? A Simple Robust Test. *Journal of Finance 58*(6), 2581–2610.

Carr, P. and L. Wu (2004). Time-Changed Lévy Processes and Option Pricing. *Journal of Financial Economics 17*, 113–141.

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, L., D. A. Lesmond, and J. Wei (2007). Corporate Yield Spreads and Bond Liquidity. *Journal of Finance 62*, 119–149.

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

Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen (2003). Alternative Models for Stock Price Dynamics. *Journal of Econometrics 116*, 225–257.

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 (2008). Volatility Components, Affine Restrictions and Non-Normal Innovations. *Journal of Business and Economic Statistics*, Forthcoming.

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., S. Heston, and K. Jacobs (2009). The Shape and Term Structure of the Index Option Smirk: Why Multifactor Stochastic Volatility Models Work so Well. Working Paper.

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.

Collin-Dufresne, P. and R. Goldstein (2001). Do Credit Spreads Reflect Stationary Leverage Ratios? *Journal of Finance 56*, 1929–1957.

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

Coval, J., J. Jurek, and E. Stafford (2009). Economic Catastrophe Bonds. *American Economic Review*, Forthcoming.

Cox, J., J. Ingersoll, and S. Ross (1985). A Theory of the Term Structure of Interest Rates. *Econometrica 53*, 385–407.

Cremers, M., J. Driessen, and P. Maenhout (2008). Explaining the Level of Credit Spreads: Option-Implied Jump Risk Premia in a Firm Value Model. *Review of Financial Studies 21*(5), 2209–2242.

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

Driessen, J., P. Maenhout, and G. Vilkov (2009). Option-implied Correlations and the Price of Correlation Risk. *Journal of Finance*, Forthcoming.

Duan, J., P. Ritchken, and Z. Sun (2006). Approximating GARCH-jump models, jump-diffusion processes, and option pricing. *Mathematical Finance 16*(1), 21–52.

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

Duan, J.-C. (1999). Conditionally Fat-Tailed Distributions and the Volatility Smile in Options. Manuscript, Hong Kong University of Science and Technology.

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

Duan, J.-C. and J. Wei (2009). Systematic Risk and the Price Structure of Individual Equity Options. *Review of Financial Studies 22*(5), 1981–2206.

Elkamhi, R. and J. Ericsson (2008). Time Varying Risk Premia in Corporate Bond Markets. Working Paper.

Elton, E. J., M. J. Gruber, D. Agrawal, and C. Mann (2001). Explaining the Rate Spread on Corporate Bonds. *Journal of Finance 56*, 247–277.

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

Engle, R. F. (1982). Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation. *Econometrica 50*, 987–1008.

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 V. K. Ng (1993, December). Measuring and Testing the Impact of News on Volatility. *Journal of Finance 48*, 1749–78.

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. (2004). Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance 59*, 1367–1403.

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

Ericsson, J. and O. Renault (2006). Liquidity and Credit Risk. *Journal of Finance 61*, 2219–2250.

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

Fama, E. and J. MacBeth (1973). Risk, Return, and Equilibrium: Empirical Tests. *Journal of Political Economy 81*(3), 607–636.

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.

Gauthier, P. and D. Possamaï (2009). Efficient simulation of the double Heston model. Working Paper.

Geske, R. (1979). The Valuation of Compound Options. *Journal of Financial Economics 7*, 63–81.

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 (1994). *Time Series Analysis*. Princeton University Press.

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.

Huang, J.-Z. and M. Huang (2002). How Much of the Corporate-Treasury Yield Spread is due to Credit Risk? A New Calibration Approach. Working Paper, Pennsylvania State University.

Huang, J.-Z. and L. Wu (2004). Specification Analysis of Option Pricing Models Based on Time-Changed Levy Processes. *Journal of Finance 59*, 1405—1439.

Jones, C. (2003). The Dynamics of Stochastic Volatility: Evidence from Underlying and Options Markets. *Journal of Econometrics 116*, 181–224.

Kim, G. H., H. Li, and W. Zhang (2009, November). The CDS/Bond Basis and the Cross Section of Corporate Bond Returns. Working Paper.

Lehnert, T. (2003). Explaining Smiles. *The Journal of Derivatives 10*(3), 27–39.

Leland, H. (2004). Predictions of Expected Default Frequencies in Structural Models of Debt. *Journal of Investment Management 2*, 5–20.

Leland, H. E. (1994, September). Corporate Debt Value, Bond Covenants, and Optimal Capital Structure. *Journal of Finance 49*, 1213–1252.

Leland, H. E. and K. B. Toft (1996). Optimal Capital Structure, Endogenous Bankruptcy and the Term Structure of Credit Spreads. *Journal of Finance 51*, 987–1019.

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

Li, Y. (2008). Credit Default Swap Spreads and the Composition of the Total Risk. Working Paper.

Lintner, J. (1965, February). The Valuation of Risk Assets and the Selection of Risky Investments in Stock Portfolios and Capital Budgets. *The Review of Economics and Statistics 47*(1), 13–37.

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

Longstaff, F. A. and E. S. Schwartz (1995). Valuing Credit Derivatives. *Journal of Fixed Income June*, 6–12.

Maheu, J. (2005). Can GARCH Models Capture Long-Range Dependence? *Studies in Nonlinear Dynamics & Econometrics 9*(4), 1.

Mauer, D. C. and M. H. Stohs (1994). The Determinants of Corporate Debt Maturity Structure. Working Paper, University of Wisconsin.

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

Merton, R. C. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance 29*, 449–470.

Mossin, J. (1966, October). Equilibrium in a Capital Asset Market. *Econometrica 34*(4), 768–783.

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.

Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica 59*, 347–370.

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.

Ritchken, P. and R. Trevor (1999). Pricing options under generalized GARCH and stochastic volatility processes. *Journal of Finance*, 377–402.

Roll, R. (1977). A Critique of the Asset Pricing Theory's Tests Part I: On Past and Potential Testability of the Theory. *Journal of Financial Economics 4*, 129–176.

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.

Saita, L. (2006). The Puzzling Price of Corporate Default Risk. Working Paper, Stanford University.

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

Sharpe, W. F. (1964, September). Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *Journal of Finance 19*(3), 425–442.

Stentoft, L. (2005). Pricing American options when the underlying asset follows GARCH processes. *Journal of Empirical Finance 12*(4), 576–611.

Tang, D. Y. and H. Yan (2007, August). Liquidity and Credit Default Swap Spreads. Working Paper.

Taylor, S. and X. Xu (1994). The Term Structure of Volatility Implied by Foreign Exchange Options. *Journal of Financial and Quantitative Analysis 29*, 57–74.

Theodossiou, P. P. (2001). Skewed Generalized Error Distribution of Financial Assets and Option Pricing. Working Paper, School of Business, Rutgers University.

Todorov, V. (2010). Variance Risk-Premium Dynamics: The Role of Jumps. *Review of Financial Studies 23*, 345–383.

Wang, H., H. Zhou, and Y. Zhou (2010). Credit Default Swap Spreads and Variance Risk Premia. Working Paper.