

GCOE Discussion Paper Series

Global COE Program

Human Behavior and Socioeconomic Dynamics

Discussion Paper No.242

Option-Implied Risk Aversion Anomalies: Evidence from Japanese Market

Nattapol TAKKABUTR

January 2012; revised September 2012

GCOE Secretariat
Graduate School of Economics
OSAKA UNIVERSITY

1-7 Machikaneyama, Toyonaka, Osaka, 560-0043, Japan

Option-Implied Risk Aversion Anomalies: Evidence from Japanese Market

Nattapol TAKKABUTR

Graduate School of Economics, Osaka University, Japan

Email: jgm042tn@mail2.econ.osaka-u.ac.jp

September 2012

Abstract

This paper empirically studied the relative risk aversion (RRA) implied from the options and historical returns of the Nikkei 225 index around the 2007-2008 subprime loan crisis. The extended use of Japanese option data and an estimation method of physical density are innovations introduced in this study. The RRA are typically downward sloping across the options' moneyness but show a clear U-shape and become negative around the at-the-money (ATM) level. Also, the RRA level decreases substantially during the crisis. Previous studies have explained these anomalies as the result of a change in the investor mix or a mispricing of options.

JEL Classification Code: G01, G130, G140

Key words: Crisis, Financial Crisis, Option, Option-implied, Risk Aversion

Acknowledgments

I gratefully acknowledge the comments and opinions provided by participants—in particular, Teruo Nakatsuma—at the 35th meeting of the Japanese Association of Financial Econometrics and Engineering at Keio University in October 2011. Any remaining errors are of course solely my responsibility.

1. Introduction

This paper chiefly aims to empirically study the relative risk aversion (RRA) of a representative stock market investor in Japan to determine how the subprime loan crisis affected risk aversion. In particular, the study examined the shape of the risk-neutral density derived from a cross section of Nikkei 225 options data from 2006 to 2010 and its relationship to a real-world density, derived from a time series of historical returns. The relationship between risk-neutral and real-world densities, as proposed by Ait-Sahalia and Lo (2000) and Jackwerth (2000), is then used as a basis to estimate the implied risk aversion measure. I found that the average RRA level rose in response to major financial events and when the crisis began to build up. However, it fell below the pre-crisis level when the Lehman shock hit the Japanese stock market. In the post-crisis period, the RRA rose again but not as high as the pre-crisis level.

Previous studies that used option prices to derive empirical risk aversion measures can be divided into two broad groups. The first group assumes a representative agent's utility function in the form of power and exponential functions that have one parameter as a constant relative risk aversion measure. Studies in this group usually begin with an estimation of risk-neutral density from option prices, and then use risk aversion coefficients as risk preference adjustments to derive real-world densities, for example, Liu et al. (2007), Bliss and Panigirtzoglou (2004), and Anagnou et al. (2002). The risk aversion measures can then be estimated conveniently using a closed-form solution based on the assumed utility function. However, the estimated risk aversion measures are either time-invariant or constant across price levels of the underlying assets because of the pre-assumed utility function.

The second group of studies does not specify a representative agent's utility functional form, but rather derives a risk aversion measure directly from the relationship between risk-neutral and real-world densities, such as Ait-Sahalia and Lo (2000) and Jackwerth (2000). The risk aversion measure estimated in this group is flexible and time-variant, which may reflect the current aggregate risk aversion better than the measures estimated from the first group. However, the estimated risk aversion measures generally form a U-shape curve across the values of the underlying assets and turn negative for certain asset prices.

This paper can be categorized into the latter group. The risk aversion measures of the Japanese stock market shows a downward sloping trend across a range of underlying indexes; however, they also show a U-shape pattern and partially fall into the negative region. This ill-behaved risk aversion measure contradicts the economic theory that a representative agent is risk averse and has a concave utility

function. Jackwerth (2000) offers a credible explanation for this behavior: the options used to derive risk-neutral densities are mispriced. The results of Ait-Sahalia et al. (2001) and Bondarenko (2003) as well support this mispricing explanation.

This paper further aims to propose a simple approach to solving a dilemma in the estimation of real-world density between the validity of a stationary assumption and the reliability of estimated parameters. Real-world density estimation usually requires a long historical return series for the identification of reliable distribution parameters. However, risk-neutral density parameters are estimated from a cross section of option data at a single point in time. Ideally, data of the same period should be used to estimate both risk-neutral and real-world densities to derive an exact implied risk aversion measure. Moreover, the use of a long past sampling period implies that the return series is stationary over the entire period, a questionable assumption in the real world. In addition, a non-overlapping series of return should be used in density estimation to avoid the autocorrelation problem. With a 30-day target horizon, one year data will provide only 12 non-overlapping monthly return observations. This leads to a dilemma of whether to use a non-overlapping-return series observed over a long period or employ overlapping returns from the shortest possible sampling period but subject to the autocorrelation problem.

To address this dilemma, I apply a parametric rather than a non-parametric method to reduce the amount of data required for reliable parameter estimation. Moreover, instead of using a target multi-period return series, a single-period model is used to forecast multi-period distribution parameters. In particular, I estimate real-world density using GARCH models both with and without leverage effects. The GARCH(1,1) model is used for non-leverage and the Glosten–Jagannathan–Runkle [GJR(1,1)] model (Glosten et al., 1993) for leverage situations. Daily Nikkei 225 returns were used for parameter estimation of a daily GARCH(1,1) or GJR(1,1) model, and then the 30-day-ahead conditional mean and variance of the distribution were forecasted using the estimation model. Therefore, with only past-one-year data, the number of returns increased to around 250 observations. The multi-period conditional mean and variance are estimated based on the simple law of expectations that does not alter the properties of the assumed model.

According to the likelihood ratio test proposed by Berkowitz (2001), in which the same GARCH model specification with both normal and Student's *t* innovation distributions are compared, models with normally distributed innovations are better ex-ante forecasts of realized returns density. This contradicts

the common belief that in comparison with the normal distribution the Student's t distribution is preferred in modeling equity returns because the t distribution has a fat-tailed property

For risk-neutral density estimation, I also chose a parametric rather than a non-parametric method. Ait-Sahalia and Lo (2000) and Ait-Sahalia et al. (2001) used pooled cross sections of option data for the observation period to non-parametrically estimate risk-neutral density. However, as Bliss and Panigirtzoglou (2004) pointed out, pooling cross sections of options data over the observation period implies that risk-neutral density is stationary across periods, which is a questionable assumption because the Nikkei 225 varied substantially during the sampling period and the number of available strikes used in estimation changes every observation day. Even if the stationary assumption is satisfied, Rosenberg and Engle (2002) suggested that the obtained risk-neutral density is an average of risk-neutral density over the observation period, not the risk-neutral density that reflects the current preference of the representative agent.

I used three different models to estimate risk-neutral density: the Heston (1993) density, the generalized beta density of the second kind (GB2), and a mixture of two lognormal densities (MLN). Though these three models produce similar results, the MLN density is prone to producing erratic RRA functions, although it is a direct extension of the basic Black-Scholes option pricing models. The Heston density is a representative of the stochastic volatility models in option pricing but requires complex numerical integration. The GB2 density has a flexible function form, but this, as well, sometimes produces erratic results. However, this paper does not aim to determine which risk-neutral density model is the best.

Section 2 discusses the estimation method in detail. Section 3 then describes the data and section 4 presents an empirical estimation of the results. Finally, I conclude in section 5.

2. Methodology

2.1. Risk Aversion Measures

This study builds on the assumptions that there exist a representative investor who is rational and risk averse or has a concave utility function, and other general characteristics such as a complete and frictionless market. With these assumptions, Ait-Sahalia and Lo (2000) showed that market-wide or

aggregate risk-neutral density $q_t(S_T)$ and the real-world density $p_t(S_T)$ are related through the representative investor's utility function $U(S_T)$ as follows:

$$\frac{q_t(S_T)}{p_t(S_T)} = \lambda \frac{U'(S_T)}{U'(S_t)} = \xi_t(S_T)$$

where S_t is the current price and S_T the end-of-period price of the underlying asset; $\xi_t(S_T)$ is a pricing kernel function, and λ a constant independent of the price level of the underlying asset, S_T . Jackwerth (2000) and Ait-Sahalia and Lo (2000) further illustrated that instead of using a utility function to derive a risk-aversion measure, which requires knowledge of the λ value, the pricing kernel can be written as

$$\xi_t'(S_T) = \lambda \frac{U''(S_T)}{U'(S_t)}.$$

Then, the Arrow-Pratt relative risk aversion measure can be derived from

$$RRA_t(S_T) = -S_T \frac{U''(S_T)}{U'(S_T)} = -S_T \frac{\xi_t'(S_T)}{\xi_t(S_T)} = S_T \left[\frac{p_t'(S_T)}{p_t(S_T)} - \frac{q_t'(S_T)}{q_t(S_T)} \right]. \quad (1)$$

To obtain the Arrow-Pratt relative risk aversion measure $RRA_t(S_T)$ defined above, the risk-neutral density $q_t(S_T)$ and the real-world density $p_t(S_T)$ need to be estimated. This risk aversion measure is also called option-implied risk aversion because the estimation of risk-neutral density is usually based on observed option data. Note that it is also possible for a researcher to predefine the $RRA_t(S_T)$ function and one of the two densities, and then derive the remaining density. In other words, knowledge of any two of the elements $RRA_t(S_T)$, $q_t(S_T)$, and $p_t(S_T)$ enables the researcher to estimate the remaining element. For example, to derive the risk aversion function $RRA_t(S_T)$, Jackwerth (2000) fitted option-implied volatility to the volatility from his model to derive risk-neutral density $q_t(S_T)$ and used a Gaussian kernel to estimate the real-world density $p_t(S_T)$. Ait-Sahalia and Lo (2000) used a non-parametric model with a Gaussian kernel to estimate both the risk-neutral density $q_t(S_T)$ and the real-world density $p_t(S_T)$. On the other hand, Bliss and Panigirtzoglou (2004) aimed to derive the implied real-world density $p_t(S_T)$ based on predefined models of the risk neutral density $q_t(S_T)$ and the risk aversion function.

2.2. Risk-Neutral Density Estimation

I used three different parametric models to estimate risk-neutral density: the Heston (1993) density, the GB2, and a mixture of two lognormal densities (MLN).

With regard to the Heston (1993) density, I followed the method applied by Gatheral (2006). Under the Heston density, stock prices S_t follow the process:

$$\begin{aligned} dS_t &= \mu_t S_t dt + \sqrt{v_t} S_t dZ_1, \\ dv_t &= -\lambda (v_t - \bar{v}) dt + \eta \sqrt{v_t} dZ_2, \\ \langle dZ_1 dZ_2 \rangle &= \rho dt \end{aligned}$$

where v_t is the conditional variance at time t and λ is the speed of reversion of v_t to its long-term mean \bar{v} .

The return density under Heston can be written as

$$f_{Hest,t}(X_i | \theta_{Hest,t}) = \frac{1}{2\pi} \int_{-\infty}^{\infty} d\varphi \cdot \exp\{\Omega(\varphi, \tau) \bar{v} + \Psi(\varphi, \tau) v - i\varphi \log(X_i/S_t)\} \quad (2)$$

where X_i is the exercise price of the option, φ is an integral variable, and $\Omega(\varphi, \tau)$ and $\Psi(\varphi, \tau)$ are functions of the model parameters $\theta_{Hest,t}$. For details of the parameters and the derivation of the model, refer to Gatheral (2006) or to the original work, Heston (1993). It should be noted that since risk neutral densities are estimated for each observation date, the parameter set $\theta_{Hest,t}$ implies that the parameters are dependent on time t . However, for simplicity, the subscript t is omitted in the distribution parameters.

With regard to the GB2 introduced by Bookstaber and MacDonald (1987), I follow the estimation method used by Liu et al. (2007). The GB2 density is defined as

$$f_{GB2,t}(X_i | \theta_{GB2,t}) = \frac{a X_i^{ap-1}}{b^{ap} B(p, q) [1 + (X_i/b)^a]^{p+q}}, \quad X_i > 0 \quad (3)$$

where $\theta_{GB2,t} = (a, b, p, q)$ are the parameters of the distribution and $B(p, q)$ is a beta function with parameters p and q . Risk neutrality imposes a condition that the forward price of the underlying asset should be equal to

$$F_{t,\tau} = \frac{bB(p + 1/a, q - 1/a)}{B(p, q)}.$$

The two-lognormal-density mixture, as well, follows the Liu et al. (2007) method, defined as

$$f_{MLN,t}(X_i | \theta_{MLN,t}) = w f_{LN1,t}(X_i | \theta_{LN1,t}) + (1 - w) f_{LN2,t}(X_i | \theta_{LN2,t}) \quad (4)$$

where $\theta_{MLN,t} = (F_1, F_2, \sigma_1, \sigma_2, w)$ is a set of distribution parameters and $\theta_{LNj,t}$, $j = 1$ to 2 is a set of parameters associated with the individual lognormal density $f_{LNj,t}$ defined as

$$f_{LNj,t}(X_i|\theta_{LNj,t}) = \frac{1}{X_i \sqrt{2\pi\sigma_j^2}} \exp\left(-\frac{[\ln X_i - \ln F_j + 0.5\sigma_j^2\tau]^2}{2\sigma_j^2\tau}\right).$$

A European call option under each assumed risk-neutral density is computed as follows:

$$C_{t,i}^{\theta_{m,t}}(X_i, \tau) = \exp(r_{f,t}\tau) \int_0^\infty (S_T - X_i)^+ f_{m,t}(S_T|\theta_{m,t}) dS_T \quad (5)$$

where $C_{t,i}^{\theta_{m,t}}(X_i, \tau)$ is a theoretical European call option price at strike X_i with time to maturity $\tau = T - t$ computed under the assumed density models $f_{m,t}(S_T|\theta_{m,t})$ with m equal to *Hest* in (2), *GB2* in (3), and *MLN* in (4), given the set of parameters $\theta_{m,t}$. $r_{f,t}$ is the risk-free rate and S_T is the price of the underlying asset at maturity date T . All assumed densities $f_{m,t}(S_T|\theta_{m,t})$ in (5) satisfied the relationship found by Ross (1976), Breeden and Litzenberger (1978), and Banz and Miller (1978):

$$f_{m,t}(S_T|\theta_{m,t}) = \frac{\partial^2 C_{t,i}^{\theta_{m,t}}(X_i, \tau)}{\partial X_i^2} \Bigg|_{X_i=S_T}.$$

A closed-form solution for the theoretical option prices in (5) under the Heston density in (2) is based on Gatheral (2006), whereas the GB2 density in (3) and the MLN density in (4) are based on Liu et al. (2007).

The least-squares method is used to estimate the parameters of densities in (2), (3), and (4) by minimizing the sum of the squared differences between the theoretical price $C_{t,i}^{\theta_{m,t}}(X_i, \tau)$ and the observed market price $C_{t,i}^{mkt}(X_i, \tau)$ at all strike prices X_i according to the equation

$$lsq_t(\theta_{m,t}^*, \tau) = \min_{\theta} \sum_i^N [C_{t,i}^{\theta_{m,t}}(X_i, \tau) - C_{t,i}^{mkt}(X_i, \tau)]^2.$$

Previous studies that estimated risk-neutral density using Nikkei 225 option data include Nakamura and Shiratsuka (1999) and Shiratsuka (2001) which used the finite difference method documented in Breeden and Litzenberger (1978) and Neuhaus (1995).

Ait-Sahalia and Lo (2000) applied a non-parametric method to estimate risk-neutral density, but their method is subject to the following assumption: the representative investor's preference is constant across the observation periods, and the derived risk-neutral density reflects the average density over

multiple periods rather than the current period density, as highlighted by Bliss and Panigirtzoglou (2004) and Rosenberg and Engle (2002).

A parametric approach offers many candidate models to choose from. Previous studies seem to be in favor of a variation of the spline functions estimation, a mixture of the lognormal density, the generalized beta density, and the Heston (1993) density. Examples of relevant studies are Shackleton et al. (2010), Liu et al. (2007), Moodley (2005), Anagnou et al. (2002), and Melick and Thomas (1997).

For comparability with previous studies, I chose Heston as a representative risk-neutral density based on the stochastic volatility model. A mixture of two lognormal densities should be more flexible than one lognormal density, while a mixture of three lognormal densities requires a larger cross section of option data for reliable parameter estimation. The GB2 is relatively quick to estimate compared with the other two models, with Heston being the most time-consuming.

2.3. Real-World Density Estimation

Ideally, estimation of real-world density should use the shortest possible observation period of past returns because risk-neutral density can be estimated with a cross section of option data for the current observation period. However, estimation of real-world density generally requires the use of a long period of observed past returns, which in turn may not reflect the true distribution of assets at the current observation date. Consequently, risk aversion measures derived from risk-neutral and real-world densities based on different observation periods may be inaccurate.

Any researcher can increase the sample of observed returns by using an overlapping-returns series. However, the overlapping-returns series would be plagued by autocorrelation problems that would, in turn, result in unreliable parameter estimation. To estimate reliable real-world density parameters using the shortest possible observation period, I use a daily parametric model to forecast multi-period distribution parameters. Therefore, a daily returns series can be used in model parameter estimation that, in turn, results in an instant increase in observation data with a minimum autocorrelation problem. Specifically, I use a three-step estimation process summarized as follows.

1. Estimate GARCH(1,1) and GJR (1, 1) model parameters based on a daily returns series over the past one year from the current observation date.
2. Select the model with the lowest AIC statistic.

3. Use estimated parameters of the selected model to forecast multi-period distribution parameters matching the target maturity horizon of risk-neutral density.

In the first step, daily stock returns are assumed to follow either a symmetric GARCH(1,1) model or an asymmetric GARCH model, proposed by Glosten et al. (1993), known as GJR (1, 1). Under the GJR (1, 1) model, the daily returns of the Nikkei 225 index follow the process:

$$r_t = \mu_t + h_t^{1/2} z_t, \quad z_t \sim iid(0,1), \quad (6)$$

$$h_t = \omega + (\alpha_1 + \alpha_2 d_{t-1}) \varepsilon_{t-1}^2 + \beta h_{t-1} \text{ with } d_{t-1} = 1 \text{ if } r_{t-1} < \mu, \text{ otherwise } d_{t-1} = 0$$

where r_t is the daily return $\ln(S_t/S_{t-1})$ at time t , h_t is the conditional variance at time t , and d_{t-1} is an asymmetric dummy variable. The innovation process z_t is assumed to follow either a normal distribution or a standardized Student's t distribution. Under step 1, parameters ω , α_1 , α_2 , and β and the degree of freedom ν in the case of a t distribution were estimated. The GARCH(1,1) model is a restricted version of (6) with the coefficient of leverage term α_2 equal to zero. μ_t is assumed to be constant, or it follows the first-order autoregressive [AR (1)] process that is automatically selected according to the AIC statistic criterion in step 2.

Under step 1, a total of eight models are estimated. Four models are estimated under each assumed distribution of z_t ; GJR(1,1) with AR(1), GJR(1,1) without AR1, GARCH(1,1) with AR(1) and GARCH(1,1) without AR(1). An AIC statistic is then used to select the most appropriate model under each assumed distribution.

Next, the estimated parameters of the selected model are used to forecast the next single-period conditional mean $\mu_{t+1}^{(1)}$ and variance $h_{t+1}^{(1)}$. The single-period conditional mean and variance of period $t+2$ are then forecasted based on $\mu_{t+1}^{(1)}$ and $h_{t+1}^{(1)}$. The process is repeated until the forecast period reaches the target horizon, in this case, 30 days ahead. All of the 30-single-period-forecasts are then summed to obtain the multi-period conditional mean $\mu_{T-t}^{(30)} = \sum_{i=1}^{29} \mu_{t+i}^{(1)}$ and the conditional variance $h_{T-t}^{(30)} = \sum_{i=1}^{29} h_{t+i}^{(1)}$. The relationship between single-period and multi-period parameters is based on the properties of expectation.

3. Data

This study applies the Nikkei 225 from December 2005 to December 2010 as it represents the broad index of the Japanese stock market. Option data on the Nikkei 225 index, obtained from NEEDS-FinancialQUEST2.0, are used to estimate risk-neutral density. Nikkei 225 options are European-style options that can be exercised on the second Friday of the expiration month. The settlement price is based on special quotation (SQ) calculated from the total opening prices of each component stock of the Nikkei 225 on the business day following the last trading day. Strike prices are multiples of ¥500 intervals based on the Nikkei 225, but for the nearest three expiration months the strike prices are multiples of ¥250 intervals. On any trading day, strike prices were set such that 17 strike prices were available for any maturity month, eight below and eight above the at-the-money (ATM) strike price. Of a total of 7,155 call and put options at the end of each month during the sampling period, only 4,223 options, or around 60%, passed the screening process.

The screening process began by eliminating options with maturities less than seven days. Then, for each day, the pair of ATM call and put options with the least difference in the closing transaction prices was identified for each maturity. If a pair could not be identified, the options in that maturity were entirely disregarded. During this step, an implied forward index level, $F_{t,\tau}$, was calculated based on the closing transaction prices of ATM call and put options from the relationship described in Ait-Sahalia and Lo (2000),

$$F_{t,\tau} = \exp(r_{f,t}\tau)(C_{t,ATM} - P_{t,ATM}) + X_{ATM} \quad , \quad (7)$$

where $r_{f,t}$ is the risk-free rate, which is an average of both buying and selling rates of new issues of three-month certificate of deposits (CDs). τ is the remaining time-to-maturity of the options, and $C_{t,ATM}$ and $P_{t,ATM}$ are the prices of ATM call and put options, respectively. X_{ATM} is the ATM strike price.

For maturities that ATM options were identified, only options with existing closing bid and ask prices and ask prices within two times the bid prices were retained. Next, mid-prices were calculated for all options. The risk-neutral densities estimated were based on out-of-the-money options that have more liquidity than in-the-money options and, hence, are less subject to pricing error. Out-of-the-money puts were converted to calls using the put-call parity. Only options data that implied volatility could be estimated and only values less than 200% per year were retained. Table 1 shows summary statistics of the

screened options for 61 months, from the end of December 2005 to the end of December 2010. Only 4,223 options, or around 60% of the total, passed the screening process; of these, 41% were out-of-the-money call and 59% were out-of-the-money put data. An ATM option is defined as a call option with strike price directly below the implied forward index level estimated from (7). By definition, this option is regarded as an in-the-money call option. Consequently, the call price implied from out-of-the-money put price was used instead of the real in-the-money call price at ATM strike.

[Table 1 around here]

Real-world densities were also estimated at the end of every month from December 2005 to December 2010. The estimation window was one year prior to each end-of-month date. For example, a real-world density estimated on the last trading day of December 2005 requires one-year daily Nikkei 225 data from December 2004. Non-overlapping one-month empirical RRA measures are then derived using the estimated real-world and risk-neutral densities.

4. Empirical Results

4.1. Real-World Density Estimation

At the end of each month from December 2005 to December 2010, two real-world densities with normal and t-distributed innovations were estimated based on the selected models either GARCH(1,1) or GJR (1, 1) using (6). Parameters were estimated from the daily returns series from one year before the observation date. The estimation used AIC statistic as a basis to automatically select among four candidate models, either a constant mean or AR(1) conditional mean GJR(1,1) or GARCH(1,1), separately for assumed normal and t-distributed innovations. Table 2 shows average AIC statistics of the four candidate models as well as how many times each model was selected during an observation period. Models with the lowest AIC statistic were selected separately under each assumed distribution of innovations, with Panel (a) showing the results for Student's t distribution and Panel (b) for the normal distribution.

[Table 2 around here]

The observation periods were divided into three sub-periods, the pre-crisis period from December 2005 to June 2007, the crisis period from July 2007 to March 2009, and the post-crisis period from April

2009 to December 2010. I divided the periods based on the observed movement of the Nikkei 225 index. Although the Lehman shock actually hit the Japanese stock market in October 2008, the Nikkei 225 continuously declined from July 2007 and did not pick up until April 2009.

It should be noted that the numbers presented in Table 2 are the average AIC statistics of an observation period. The statistics for any observation month can be quite different from the numbers shown in Table 2. To provide a meaningful conclusion, Table 2 shows, in parentheses, how many times each model achieved the lowest AIC statistic and was selected in an observation period. Regardless of the assumed distribution and observation periods, symmetric GARCH(1,1) models dominated GJR(1,1) models with leverage effects, which were not selected at all.

Under the Student's t-distributed innovation assumption in Panel (a), the GARCH(1,1) model with a constant mean but without a leverage term was the dominant model, and was selected 56 times out of 61 observations. The GARCH(1,1) model with conditional mean AR(1) was selected three times in the pre-crisis period and two times in the post-crisis period.

Under the normal innovation assumption, the symmetric GARCH(1,1) model with a constant mean was selected 59 times out of 61 observations. The GARCH(1,1) with conditional mean AR(1) was selected once during the pre-crisis period, in September 2009, and once during the post-crisis period, in November 2010.

Table 3 shows averages of the estimated parameters under the GARCH(1,1) models by observation period. Numbers in parenthesis are average robust standard errors of each coefficient. The numbers in the AR(1) coefficient ϕ column are averages of applicable sampling dates only. N/A indicates that models with the coefficients were not selected at all during the observation period. Again, it should be noted that average estimated coefficients and standard errors can be used only as reference values for the average magnitude of each observation period. An inference about the statistical significance of coefficients should not be made because the estimated coefficients on any observation date may be considerably different from the number shown in Table 3.

[Table 3 around here]

While practitioners often use AIC statistics as a tool for model selection, it cannot be used rigorously to compare models across different distribution assumptions of innovation series. To compare

between the assumptions under normal and Student's t-distributed innovation, I used the likelihood ratio (LR) statistic in Berkowitz (2001), which provides a reliable test of the entire distribution even with a small sample size of forecasted distributions. I could compare the Berkowitz LR statistic across different estimated densities because it is robust against the functional form of densities.

Nevertheless, it should be noted here that the Berkowitz LR statistic is conceptually different from the AIC statistic. The AIC statistic indicates how well a model fits a given data set, but the Berkowitz LR statistic shows how well a model predicts future realization. It can be seen from Table 2 that the GARCH(1,1) with a constant mean is the dominant model under both the assumed distributions. Consequently, to estimate the Berkowitz LR statistic, I only took samples from an observation date when the constant mean GARCH(1,1) model was simultaneously selected under both normal and t-distributed innovations. The final observation inputs for the Berkowitz LR statistic is reduced to only 55 observations under this setting.

According to Berkowitz (2001), an ex-ante forecasted density function is considered appropriate if the following transformed variable, z_t , is iid $N(0, 1)$:

$$z_t = \Phi^{-1} \left[\int_{-\infty}^{y_t} f(u) du \right] \quad (8)$$

where $f(u)$ is an ex-ante forecasted density function of the realization y_t . In this study, y_t is the realization of the 30-days-ahead Nikkei 225 index level, which is the actual index level on observation day $t + 30$. On every end-of-month date t , the z_t value is calculated and then the series z_1, z_2, \dots, z_N over the entire observation period is used to estimate the following AR (1) model with a constant mean and variance:

$$z_t - \mu = \rho(z_{t-1} - \mu) + \epsilon_t \quad (9)$$

The maximum-likelihood estimated $\hat{\mu}$, $\hat{\sigma}^2$, and $\hat{\rho}$ are used to calculate the log-likelihood value of (9). The likelihood ratio test statistic can now be given with a restricted likelihood value obtained by setting $\mu = 0, \sigma^2 = 1$, and $\rho = 0$. Specifically, the Berkowitz LR statistic is in the form

$$LR = -2[L(0,1,0) - L(\hat{\mu}, \hat{\sigma}^2, \hat{\rho})] \quad (10)$$

Table 4 shows the Berkowitz LR statistic and log-likelihood values obtained from the real-world density under the constant mean GARCH(1,1) model with the assumption of normal and Student's t-distributed innovations. The LR statistics show that a real-world density under normally distributed innovations is appropriate to forecast a 30-days-ahead density of Nikkei 225 returns. The p-value of a Student's t-distributed innovations model implies that the model was rejected at a 1% significance level. This finding contradicts the general assumption that the Student's t distribution is a more appropriate choice than the normal distribution to model equity returns because of its fat-tail property.

[Table 4 about here]

4.2. Risk-Neutral Density Estimation

At the end of each observation month, a cross section of the options data that passed the screening process described in section 3 was used to estimate the 30-days-ahead risk-neutral density. An option series with a maturity nearest to ± 15 days from 30 days to maturity is used as a proxy for 30-day maturity options. However, if two option series were within ± 15 days from 30 days to maturity, risk-neutral densities were estimated from both near-term and next-term option series and then interpolated in the same manner as the Chicago Board Options Exchange (CBOE) does for the VIX index. Before interpolation, each density is normalized by the respective implied forward level, $F_{t,T}$.

The number of available strikes must meet the minimum strikes requirement of each density model to ensure the estimation of density parameters. Because Heston's (1993) density has five parameters, I set the required minimum number of strikes equal to eight. The GB2 has four parameters with one parameter restricted to be equal to the forward index level under risk neutrality. The minimum number of strikes for the GB2 density was set to seven. Because the mixture of two lognormal densities has five parameters, the minimum number of strikes required was set to 10. However, if the minimum number of strikes for the Heston and GB2 densities were not attained, the cross section option data of the previous trading day were used instead. The mixture of two lognormal densities reduced to a lognormal density with a minimum strike requirement of five. If any of the models did not pass the minimum strike requirement, all models used the same previous trading day data.

4.3. Risk-Aversion Measure Estimation

With both the risk-neutral density $q_t(S_T)$ and the real-world density $p_t(S_T)$, the relative risk aversion measure $RRA_t(S_T)$ on each observation month can be estimated using (1). The following results include only RRA measures that are calculated based on real-world density under the normally distributed innovations assumption as suggested by Berkowitz's LR statistics. All RRA measures based on three different models of risk-neutral densities specified in section 2.2 are presented.

The RRA measures defined by (1) were estimated only over the observable strike ranges on each month end. Consequently, RRA measures are defined over a different domain of underlying assets. To compare the RRA measures across observation periods, I normalized the domain by using moneyness defined as a ratio of strike $X_i = S_T$ to the forward index level $F_{t,\tau}$ instead of the gross level of the underlying index.

Figure 1 shows the commonly observed 30-days-ahead option-implied risk-neutral densities and the matching horizon real-world density estimated from the GARCH(1,1) model with normal innovation. The three models of risk-neutral density include the Heston (1993) density, the GB2, and a mixture of two lognormal densities. The upper rows of Panels (a), (b), and (c) show three typical shapes of densities found during the pre-crisis period, the crisis period, and the post-crisis period, respectively. The lower rows indicate the associated relative risk aversion functions $RRA_t(X_i)$ estimated from each pair of real-world and risk-neutral densities using (1).

[Figure 1 around here]

During the pre-crisis period, from December 2005 to June 2007, the mode of real-world densities (labeled RWD-N) is located to the right of the mode of risk-neutral densities (labeled Hest, GB2, and MLN). A well-behaved downward sloping RRA function in line with economic theory can be observed when the mode of the real-world density is located further to the right and higher than the mode of the risk-neutral density [Panel (a), left]. A U-shaped pattern can be observed around the ATM level as found in previous studies, for example, Jackwerth (2000) and Ait-Sahalia and Lo (2000). The RRA function exhibits a sharp U-shaped curve when the mode of the real-world density falls below and approaches the mode of the risk-neutral density [Panel (a), middle]. When news about the subprime mortgage crisis in

the U.S. began to emerge in 2007, the mode of the real-world density moved toward the left of the mode of risk-neutral density, resulting in steep downward sloping RRA functions with sharp U-shape curves [Panel (a), right]. The RRA functions also partially become negative around ATM levels.

In the crisis period, risk-neutral densities become more left skewed, showing expectations of further decline in asset prices in the future. The relationships between real-world and risk-neutral densities are the opposite of those in the pre-crisis period. All the modes of real-world densities across the observation months are located to the left of the risk-neutral densities' modes. The same downward sloping RRA function and U-shaped curve around the ATM level can be observed, but the RRA function turned more negative than the pre-crisis level. The steep downward sloping RRA function can be observed when the modes of both real-world and risk-neutral densities are close to each other [Panel (b), left]. When the mode of the real-world density fell below those of risk-neutral densities, sharp U-shaped curves are observed again [Panel (b), middle]. When the Lehman shock hit the Japanese market in October 2008, an obvious upward sloping risk aversion function implied a perverse risk preference of the representative investor [Panel (b), right]. In other words, the representative investor became risk seeking instead of risk averse.

During the post-crisis period, the mode of real-world density began to move right toward the mode of risk-neutral density. A downward sloping risk aversion function is observed, and the RRA function becomes less negative around the ATM level. In late 2009, the relationship between real-world density and risk-neutral densities was similar to the crisis period [Panel (c), left]. In the beginning of 2010, the RRA function around the ATM level began to recover and shift into the positive region when the mode of real-world density moved toward the right of the mode of risk-neutral densities [Panel (c), middle]. However, the RRA function around the ATM level fell back a little into the negative region and remained around -5 till the end of December 2010 [Panel (c), right].

All of the sub-periods show that the RRA functions, on average, are downward sloping in line with economic theory. However, the functions sometimes show a sharp U-shaped curve, partially become negative, and then increase with the level of expected asset prices in the future. These three anomalies suggest that the standard assumptions in economic theory may be wrong and representative investors may not be rational or may have a convex utility function. However, as pointed out in Jackwerth (2000), the standard assumptions may hold true, but the investors' expectation about future asset prices may not be

accurate. The wrong expectation of future returns results in mispricing of options, which in turn distort the estimation of risk-neutral densities.

Bliss and Panigirtzoglou (2004) explained that the risk preference of the representative investor might actually change during a period of high volatility of underlying assets. The mix of market participants changed because investors with greater risk aversion left the market during the high volatile period. However, I found that Jackwerth's (2000) option mispricing interpretation might be a more plausible explanation. Figure 2 shows the effective strike range introduced by Andersen et al. (2011) to measure the coverage of options used to compute the volatility index (VIX). The effective range on any observation date t is defined as

$$ER_t = \left[\frac{\ln(X_{1,t}/F_{t,\tau})}{\hat{\sigma}_{BS,t}\sqrt{T}}, \frac{\ln(X_{n,t}/F_{t,\tau})}{\hat{\sigma}_{BS,t}\sqrt{T}} \right], \quad (11)$$

where $\hat{\sigma}_{BS,t}$ is the Black-Scholes implied volatility of ATM options and $F_{t,\tau}$ is the forward price level of the underlying asset. $X_{1,t}$ and $X_{n,t}$ are, respectively, the lowest and highest strikes observed. The effective range can be considered as the range of highest and lowest returns at the future date, T .

[Figure 2 around here]

The square shaded area in Figure 2 indicates the crisis period. The minus region represents the coverage of out-of-the-money put data as $X_1 < F_{t,\tau}$. The positive region shows the coverage of out-of-the-money call data. Notice how the effective range of out-of-the-money put reduced to nearly zero when the Lehman shock hit the Japanese market in October 2008, during which daily returns varied from +13% to -13%. Options are subject to a high degree of mispricing during such a volatile period. The cross section of options data used in the risk-neutral density estimation of October 2008 observations was from the previous two trading days, and not the precise last trading day of the month. The estimation failed to obtain valid parameters from the data of the previous day and the last trading day because of a lack of valid options data.

Figure 3 shows the average relative risk aversion measures $RRA(X_i)$ at each level of moneyness, defined by X_i/F_t . The range of moneyness was limited to 0.90–1.05 levels, which are available ranges shared by all sub-sampling periods. Results for the pre-crisis, crisis, and post-crisis periods are shown in

Panels (a), (b), and (c), respectively. Across all observation periods, the RRA function shows a U-shaped pattern around the ATM level, as in Jackwerth (2000). Figure 3 also confirms the findings in Figure 1 that on average the RRA functions were positive during the pre-crisis period [Panel (a)], become negative during the crisis period [Panel (b)], and recover a little in the post-crisis period [Panel (c)].

[Figure 3 around here]

The U-shape pattern found in Figures 1 and 3 suggest that, after the prices of the underlying asset increase to a certain level, the representative investor's risk aversion level again increases. Instead of welcoming the higher returns from rising asset prices, the representative investor prefers lower returns. As Jackwerth (2000) suggested, the U-shape pattern and negative risk aversion may be the result of option mispricing. Bondarenko (2003) studied the anomaly of the so-called overpriced puts puzzle of S&P 500 put options and explained that excessive weighting by investors of the probability of negative S&P 500 returns results in biased subjective future returns density and, hence, a biased risk-neutral distribution. Jackwerth (2000) and Ait-Sahalia et al. (2001) showed that an option trading strategy can earn positive returns, which supports the mispricing explanation.

4.4. *Relative Risk Aversion Over Time*

To study the structure of risk aversion over time, the definition of relative risk aversion in (1) is modified. On each observation month t , I calculate the sample average of $RRA_t(X_i)$ across levels of strike prices X_i to obtain a representative relative risk aversion level \overline{RRA}_t . Table 5 shows the average \overline{RRA}_t over various observation periods. The average levels of \overline{RRA}_t , regardless of the risk-neutral density model, decreased during the crisis period and recovered in the post-crisis period.

[Table 5 around here]

Figure 4 shows the movement of relative risk aversion, \overline{RRA}_t , over the observation periods together with Nikkei 225 index levels and returns, as well as effective strike ranges. The shaded area represents the crisis period. Figure 4 suggests that the RRA function fluctuates over time. During the pre-crisis period, the \overline{RRA}_t moves around 5 to 20 points. During the crisis period, it varies substantially from

a peak of approximately 28 points in September 2007 to approximately -2 points in September 2008. The September 2007 peak may be the result of bankruptcies of distressed financial institutions in the U.S. and Europe. In the post-crisis period, \overline{RRA}_t rose again, but not until 2010, when it began to recover to the pre-crisis level.

One finding from Figure 4 is in contrast to the general belief that the risk aversion level will skyrocket when an external shock hits the stock market the hardest. The risk aversion level bottomed out in September 2008 when Lehman Brothers filed for bankruptcy protection on September 15, and the Nikkei 225 closed 5% lower than the previous day. Even in October 2008, when the Daiwa Life Assurance bankruptcy sent the Nikkei 225 index down 10% in one day, the RRA measure was still well below the non-crisis level. The change in investor mix explanation provided by Bliss and Panigirtzoglou (2004) can also be applied to this anomaly. Only investors with a high degree of risk tolerance remained in the market, and those with a higher risk aversion left the market during the highly volatile period. Although their explanation is rational and perhaps accurate, they did not formally test this hypothesis. Nevertheless, option mispricing is a more plausible explanation, as evident from the lower panel of Figure 4, which shows the effective range of out-of-the-money put options contracted substantially during September and October 2008. With highly volatile underlying assets, options are subject to a high degree of mispricing.

[Figure 4 around here]

5. Conclusion

This study extends previous studies on option-implied risk-neutral density in the Japanese market by making use of additional information from a real-world density data to derive a relative risk aversion measure.

A simple approach that uses a daily model to estimate multi-period distribution parameters worked well in deriving real-world density, as evidenced by the likelihood ratio test of Berkowitz (2001). However, contrary to the assumption that the Student's t distribution models stock returns better, the Berkowitz LR test suggested a model with normally distributed innovations. With regard to risk-neutral density estimations, three different models produced similar results. However, the GB2 and MLN density models, though flexible and easy to estimate, are prone to producing erratic RRA functions, while the

Heston model involves complex estimation but is less likely to produce an erratic result. I suggest the use of a model that fits the available data on hand. An objective selection from these three candidate models is beyond the scope of this paper.

In general, the relative risk aversion in the Japanese market shows downward sloping characteristics consistent with the economic theory that the representative agent is risk averse. However, the market also shares the ill-behaved risk aversion function as documented in previous S&P 500 studies. Through an examination of density shapes and the associated RRA function patterns, I showed that the desired relationship between the real-world and risk-neutral densities occurs when the mode of the real-world density $p_t(S_T)$ stays further to the right and higher than the mode of the risk-neutral density $q_t(S_T)$.

U-shape and partially negative patterns are also observed in the Japanese market when RRA measures at each price level of the underlying assets were averaged across observation periods, similar to Jackwerth (2000). The RRA curves stayed in the negative region at almost all levels of moneyness during the crisis period, but recovered in the post-crisis period.

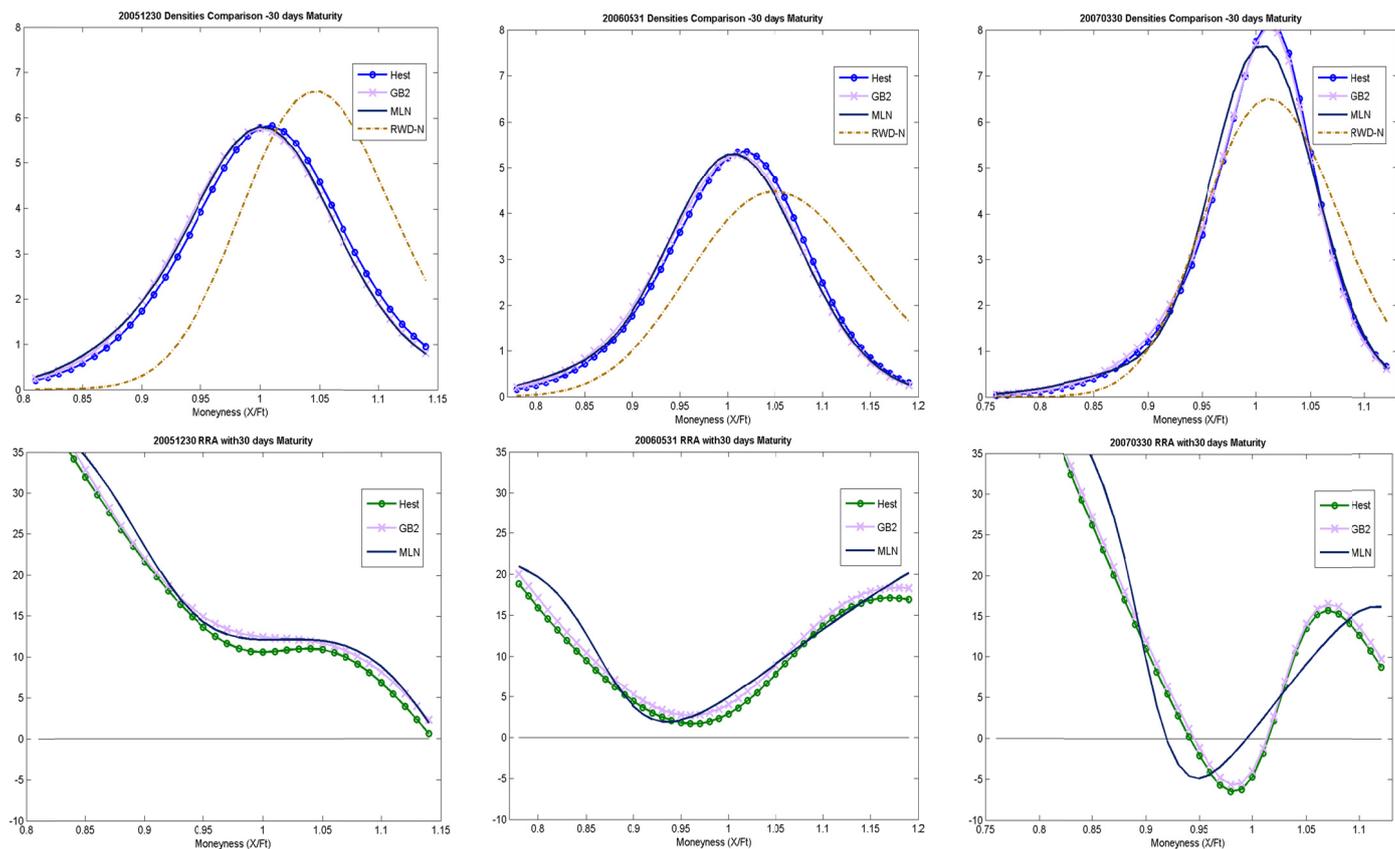
On each observation date, the market-wide level of relative risk aversion measure \overline{RRA}_t , which is an average of $RRA_t(X_i)$ across price levels of underlying assets, showed that risk aversion levels dropped below the normal level during a market turmoil characterized by high-return volatility. Previous literature preferred the option mispricing explanation for both the U shape and the sinking RRA level during the crisis period, and showed that a simple option trading strategy can result in abnormal positive returns. I agree with the mispricing explanation, particularly for the bottomed-out phenomenon, as evident from a very volatile effective strike range of a cross section of option data. However, whether some arbitragers could exploit a constant mispricing of options and correct the prices is arguable. Further research on the mispricing explanation is needed in the Japanese market.

References

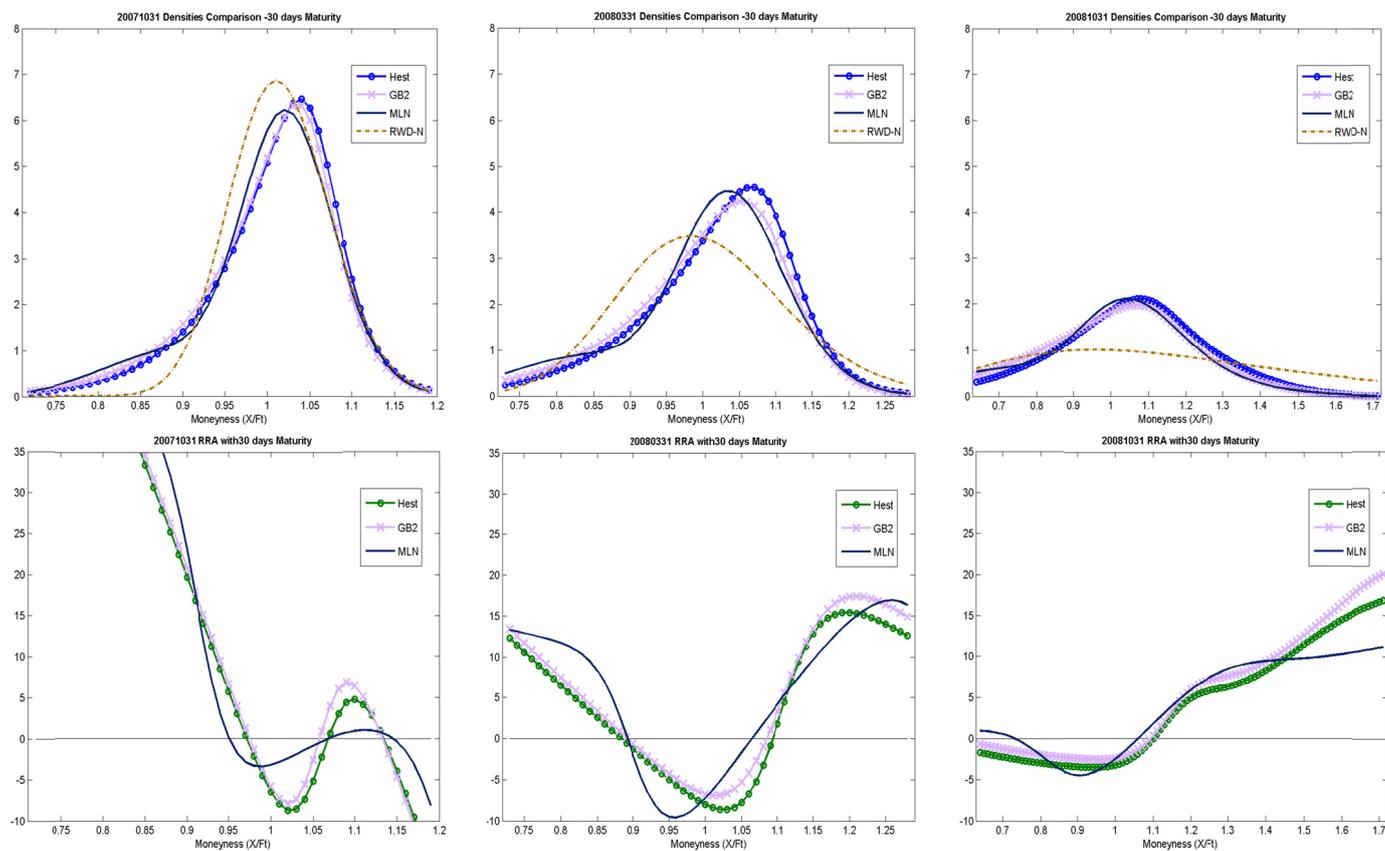
- Ait-Sahalia, Y. and Lo, A.W. (2000), "Nonparametric Risk Management and Implied Risk Aversion," *Journal of Econometrics* 94, pp. 9–51.
- Ait-Sahalia, Y., Wang, Y. and Yared, F. (2001), "Do Option Markets Correctly Price the Probabilities of Movement of the Underlying Asset?" *Journal of Econometrics* 102, pp.67–110.
- Anagnou, I., Bedendo, M., Hodges, S. and Tompkins, R. (2002, March), "*The Relation between Implied and Realized Probability Density Functions*," Paper presented at the EFA 2002 Berlin Meetings, Berlin.
- Andersen, T. G., Bondarenko, O. and Gonzalez-Perez, M. T. (2011, January), "*A Corridor Fix for VIX: Developing a Coherent Model-free Option-implied Volatility Measure*," Paper presented at the 2nd Osaka Conference on "High-Frequency Data Analysis in Financial Markets," Osaka.
- Banz, R. and Miller, M. (1978), "Prices for State-Contingent Claims: Some Estimates and Applications," *Journal of Business* 51, pp.653–672.
- Berkowitz, J. (2001), "Testing Density Forecasts, with Applications to Risk Management," *Journal of Business & Economic Statistics* 19, pp.465–474.
- Bliss, R. and Panigirtzoglou, N. (2004), "Option-implied Risk Aversion Estimates," *Journal of Finance* 59, pp.407–446.
- Bondarenko, O. (2003, November), "*Why are Put Options so Expensive?*," Paper presented at the 64th Annual Meeting of the American Finance Association, California.
- Bookstaber, R. and MacDonald, J. (1987), "A General Distribution for Describing Security Price Returns," *Journal of Business* 60, pp.401–424.
- Breeden, D. and Litzenberger, R. (1978), "Prices of State-contingent Claims Implicit in Options Prices," *Journal of Business* 51, pp.621–651.
- Gatheral, J. (2006), *The Volatility Surface: A Practitioner's Guide*, New York: Wiley Finance.
- Glosten, L.R., Jagannathan, R. and Runkle, D. (1993), "On the Relationship between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance* 48, pp.1779–1801.
- Heston, S. (1993), "A Closed-form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies* 6, pp.327–343.

- Jackwerth, J. (2000), “Recovering Risk Aversion from Option Prices and Realized Returns,” *Review of Financial Studies* 13, pp.433–451.
- Liu, X., Shackleton, M.B., Taylor, S.J. and Xu, X. (2007), “Closed-form Transformations from Risk-Neutral to Real-world Distributions,” *Journal of Banking & Finance* 31, pp.1501–1520.
- Melick, W. R. and Thomas, C. P. (1997), “Recovering an Asset’s Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis,” *Journal of Financial and Quantitative Analysis* 32, pp.91–115.
- Moodley, N. (2005), “The Heston Model: a Practical Approach with Matlab Code,” University of Witwatersrand.
- Nakamura, H. and Shiratsuka, S. (1999), “Extracting Market Expectations from Option Prices: Case Studies in Japanese Option Markets,” *Monetary and Economic Studies* 17, pp.1–43.
- Neuhaus, H. (1995), “The Information Content of Derivatives for Monetary Policy,” Discussion Paper 3/95, Economic Research Group of the Deutsche Bundesbank.
- Rosenberg, J. and Engle, R. F. (2002), “Empirical Pricing Kernels,” *Journal of Financial Economics* 64, pp.341–372.
- Ross, .S. (1976), “Options and Efficiency,” *Quarterly Journal of Economics* 90, pp.75–89.
- Shackleton, M.B., Taylor, S.J. and Yu, P. (2010), “A Multi-horizon Comparison of Density Forecasts for the S&P 500 using Index Returns and Option Prices,” *Journal of Banking & Finance* 34, pp.2678–2693.
- Shiratsuka, S. (2001), “Information Content of Implied Probability Distributions: Empirical Studies of Japanese Stock Price Index Options,” *Monetary and Economic Studies* 19, pp.143–170.

Panel (a): Typical shapes of densities and RRA functions during pre-crisis period from December 2005 to June 2007



Panel (b): Typical shapes of densities and RRA functions during crisis period from July 2007 to March 2009



Panel (c): Typical shapes of densities and RRA functions during post-crisis period from April 2009 to December 2010

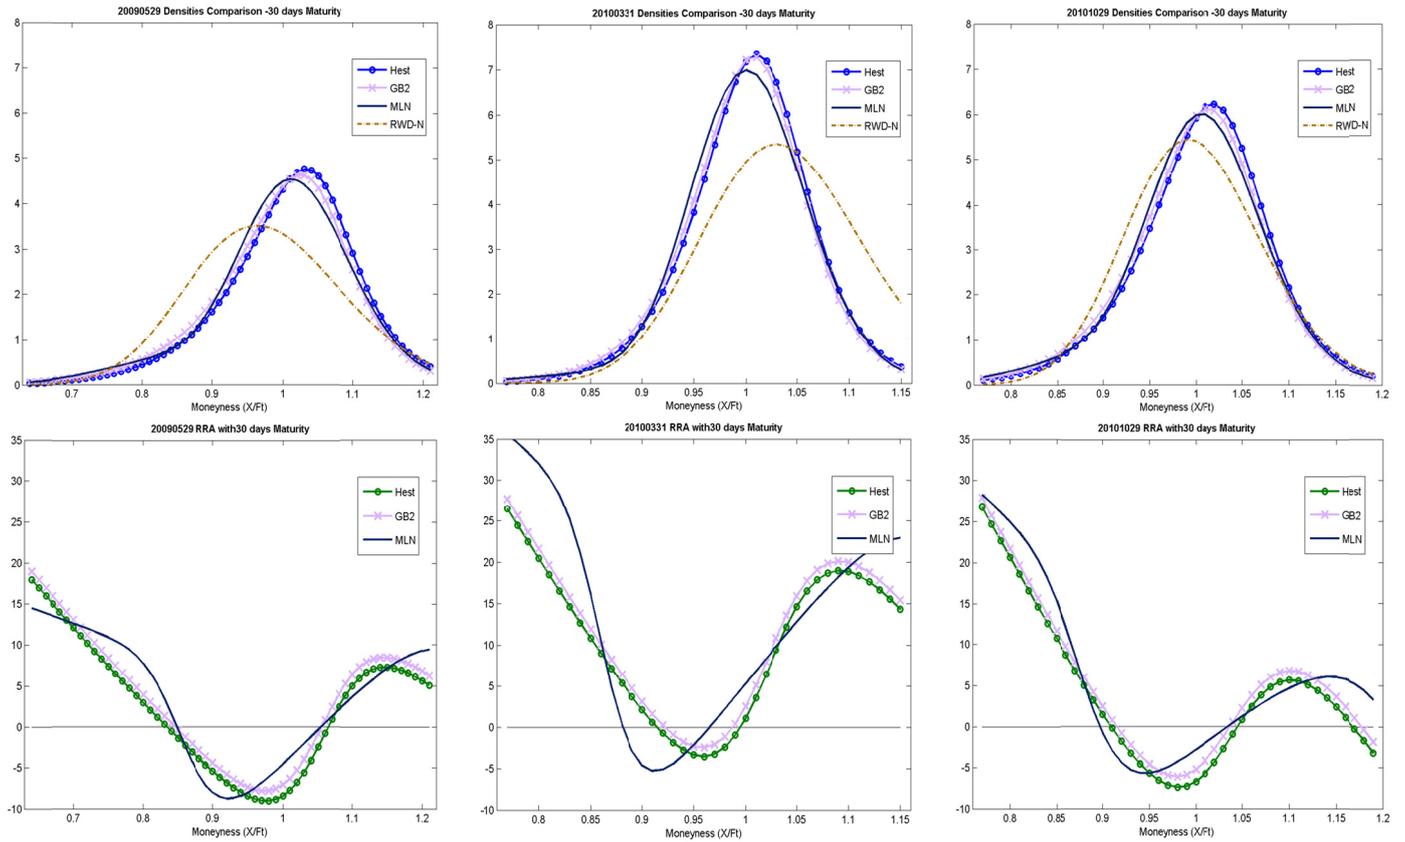


Figure 1: Density estimation and associated RRA functions

The upper row in each panel shows three typical shapes found from estimated real-world densities $p_t(S_T)$ under GARCH(1,1) with normal innovations and risk-neutral densities $q_t(S_T)$ under assumed models: Heston in (2), GB2 in (3), and MLN in (4). The lower rows show relative risk aversion functions $RRA_t(X_i)$ estimated from each pair of real-world and risk-neutral densities using (1).

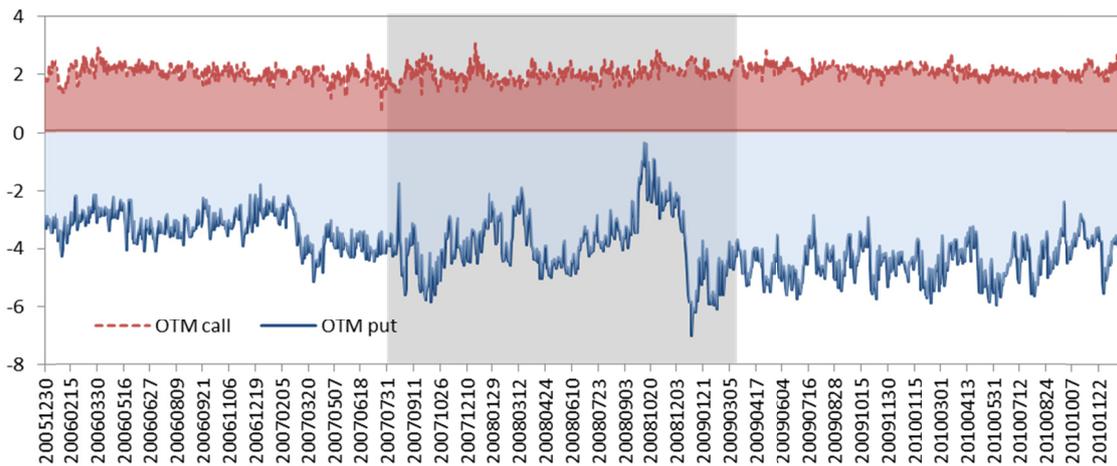


Figure 2: Effective strike range

The effective strike range on each observation date ER_t is defined in (11). Andersen et al. (2011) introduced the effective range to measure the coverage of options used in the volatility index (VIX) calculation.

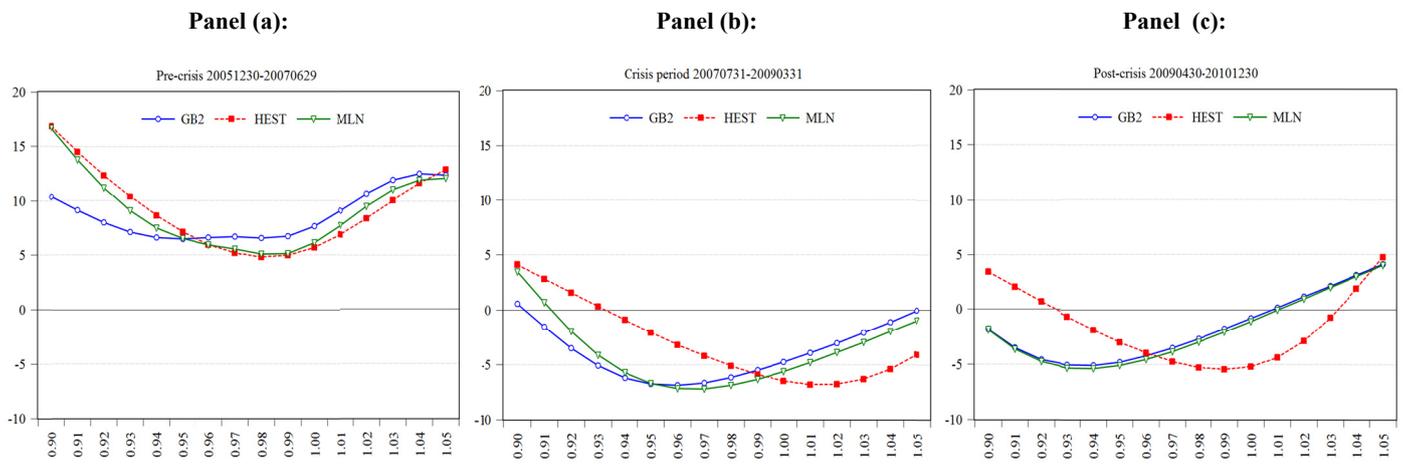
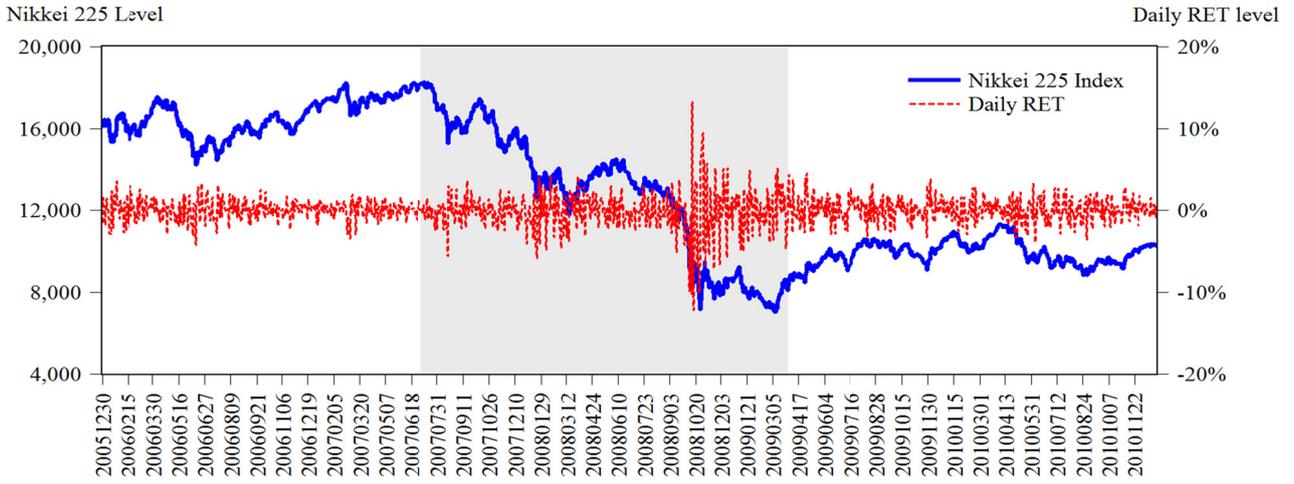


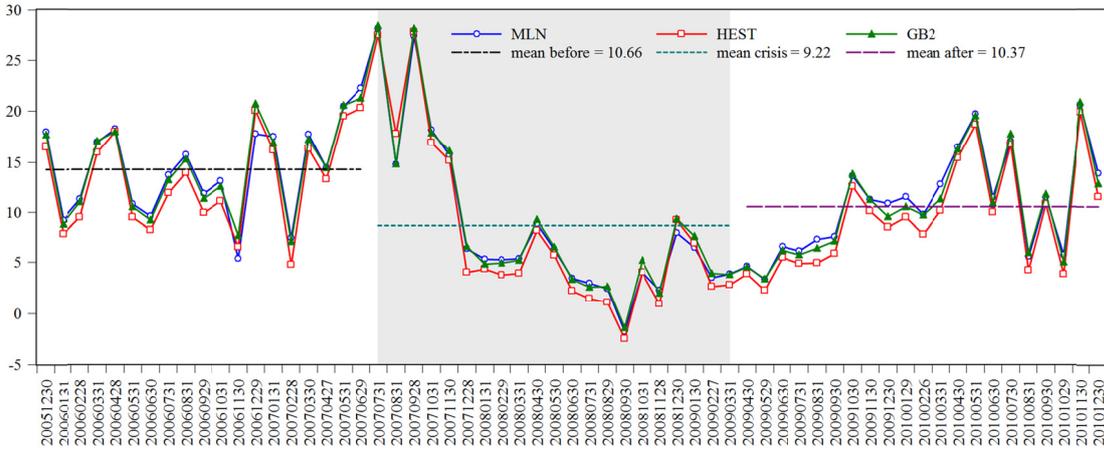
Figure 3: Average RRA at each price level of the underlying asset.

Curves labeled GB2, Heston, and MLN are average relative risk aversion measures $RRA_t(X_t)$ at each level of moneyyness defined, by $X_t/F_{t,\tau}$ calculated from risk-neutral densities under assumed models and real-world density under GARCH(1,1) with normal innovations.

Nikkei 225 index and daily returns



Relative risk aversion measure $\overline{RR\bar{A}}_t$ based on real-world density under normal innovation models and three risk-neutral densities



Effective ranges of option strikes

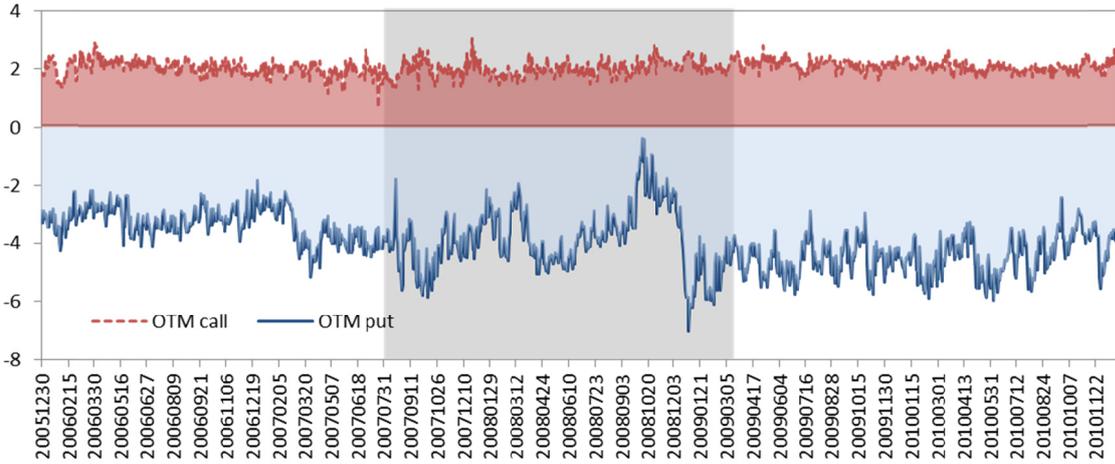


Figure 4: Average relative risk aversion measure $\overline{RR\bar{A}}_t$ and the Nikkei 225 index

The middle chart shows $\overline{RR\bar{A}}_t$ across observation periods. $\overline{RR\bar{A}}_t$ is a sample average of relative risk aversion $RR\bar{A}_t(X_i)$ across levels of moneyness defined as $X_i/F_{t,\tau}$ for each observation month t .

Table 1: Summary statistic of Nikkei 225 options

	Options before screening	Options after screening	Average before screening	Average after screening	Minimum	Maximum
OTM call	3,731	2,512	61	41	23	77
OTM put	3,424	1,711	56	28	13	63
Total	7,155	4,223	117	69		

Out-of-the-money is abbreviated as OTM. Options before and after screening refer to the total number of options before and after the screening process at the end of the month from December 2005 to December 2010 (61 months). Averages before and after screening refer to the average numbers of options before and after the screening process on each observation. Minimum and maximum are the minimum and maximum number of options that passed the screening process on each observation.

Table 2: Average AIC statistics for candidate models of each sampling period

Panel (a): Average AIC statistics for Student's t distribution models		With leverage		Without leverage	
		GJR(1,1)		GARCH(1,1)	
		AR mean	Constant mean	AR mean	Constant mean
Full	20051230 : 20101230	1,432.49 (0)	1,429.58 (0)	1,422.39 (5)	1,419.42 (56)
Before	20051230 : 20070629	1,569.16 (0)	1,566.74 (0)	1,557.54 (3)	1,555.37 (16)
Crisis	20070731 : 20090331	1,409.25 (0)	1,406.35 (0)	1,397.42 (0)	1,394.31 (21)
After	20090430 : 20101230	1,332.08 (0)	1,328.70 (0)	1,325.08 (2)	1,321.52 (19)

Panel (b): Average AIC statistics for normal distribution models		With leverage		Without leverage	
		GJR(1,1)		GARCH(1,1)	
		AR mean	Constant mean	AR mean	Constant mean
Full	20051230 : 20101230	1,429.57 (0)	1,426.62 (0)	1,417.72 (2)	1,412.66 (59)
Before	20051230 : 20070629	1,566.46 (0)	1,563.96 (0)	1,553.94 (0)	1,544.86 (19)
Crisis	20070731 : 20090331	1,405.14 (0)	1,402.23 (0)	1,389.64 (1)	1,387.22 (20)
After	20090430 : 20101230	1,330.16 (0)	1,326.75 (0)	1,322.54 (1)	1,318.48 (20)

For each observation period, a model with the lowest AIC statistic is used to estimate parameters of real-world density. Enclosed in parentheses is the number of times each model achieved the lowest AIC statistic and was selected during an observation period.

Table 3: Average of estimated parameters for each sampling period

Panel (a): GARCH(1,1) - Student's t distributed innovations

	Observation period		μ_t	ϕ^*	ω	α_1	β	ν
	Start	End						
Full	20051230	: 20101230	0.0003 <i>(1.99e-05)</i>	-0.0210 <i>(0.0004)</i>	1.05e-05 <i>(2.68e-07)</i>	0.0980 <i>(2.11e-03)</i>	0.8627 <i>(0.0188)</i>	75.9110 <i>(0.3468)</i>
Before	20051230	: 20070629	0.0012 <i>(1.99e-05)</i>	-0.0223 <i>(0.0002)</i>	3.25e-06 <i>(5.27e-08)</i>	0.0822 <i>(7.91e-04)</i>	0.8974 <i>(0.0131)</i>	83.6669 <i>(0.1137)</i>
Crisis	20070731	: 20090331	-0.0004 <i>(2.48e-05)</i>	N/A	8.03e-06 <i>(2.10e-07)</i>	0.1052 <i>(2.61e-03)</i>	0.8745 <i>(0.0181)</i>	30.2719 <i>(0.1836)</i>
After	20090430	: 20101230	0.0002 <i>(1.49e-05)</i>	-0.0189 <i>(0.0007)</i>	1.96e-05 <i>(5.20e-07)</i>	0.1050 <i>(2.80e-03)</i>	0.8196 <i>(0.0246)</i>	114.5328 <i>(0.7208)</i>

Panel (b): GARCH(1,1) – normally distributed innovations

	Observation period		μ_t	ϕ^*	ω	α_1	β
	Start	End					
Full	20051230	: 20101230	0.0002 <i>(1.33e-05)</i>	-0.0352 <i>(0.0007)</i>	1.14e-05 <i>(1.91e-07)</i>	0.1096 <i>(1.96e-03)</i>	0.8448 <i>(0.0134)</i>
Before	20051230	: 20070629	0.0012 <i>(1.49e-05)</i>	N/A	4.07e-06 <i>(3.60e-08)</i>	0.0904 <i>(1.60e-03)</i>	0.8821 <i>(0.0103)</i>
Crisis	20070731	: 20090331	-0.0005 <i>(1.64e-05)</i>	-0.0648 <i>(0.0013)</i>	9.67e-06 <i>(2.21e-07)</i>	0.1276 <i>(3.01e-03)</i>	0.8405 <i>(0.0206)</i>
After	20090430	: 20101230	0.0002 <i>(8.71e-06)</i>	-0.0056 <i>(0.0001)</i>	1.98e-05 <i>(3.01e-07)</i>	0.1089 <i>(1.23e-03)</i>	0.8153 <i>(0.0089)</i>

* Average of applicable observations only.

The GARCH(1,1) model is estimated using (6). The mean parameter is μ_t , which is automatically selected to be a constant mean or an AR(1) mean based on AIC statistics. If the AR(1) model is selected $\mu_t = c + \phi r_{t-1}$, where c is a constant, ϕ is a first-order autoregressive coefficient and r_{t-1} is the lag of the current return. The constant mean model implies $\phi = 0$. ω is a constant of the conditional variance equation, β is the coefficient of conditional variance σ_{t-1}^2 and α_1 is the coefficient of squared innovation of the previous period ε_{t-1}^2 . In the case of Student's t-distributed innovations, the degree of freedom ν is an additional parameter. Numbers in parentheses are average standard errors of coefficients.

Table 4: Berkowitz likelihood ratio test of real-world density models

Berkowitz likelihood ratio statistics of GARCH(1,1) with constant mean for full observation period		
	t-distributed innovations	Normally-distributed innovations
Berkowitz LR statistics	11.4686	4.3561
p-value	(0.0094)	(0.2255)
LL unrestricted	-61.5825	-65.5777
LL restricted	-55.8482	-63.3996

Berkowitz LR statistics are calculated using (10) and distributed as chi-squared with three degrees of freedom with a null hypothesis that the series of transformed variable z_t in (8) are independent with zero mean and unit variance. LL unrestricted is a maximized log-likelihood value from model (9) without restriction and LL restricted is a maximized log-likelihood value under model (9) with the restrictions $\mu = 0, \sigma^2 = 1$, and $\rho = 0$.

Table 5: Average \overline{RRA}_t based on real-world density with normally distributed innovations

Observation period				Risk-neutral density models		
	Start	End		Heston	GB2	MLN
Full	20051230	: 20101230	Ave	10.0338	11.0699	11.1104
			Min	-2.4926	-1.4471	-1.7866
			Max	27.8457	28.4419	28.1556
Before	20051230	: 20070629	Ave	13.1501	14.2775	14.3218
			Min	4.8355	7.1520	5.4612
			Max	20.3776	21.3791	22.3255
Crisis	20070731	: 20090331	Ave	7.8271	8.6873	8.4533
			Min	-2.4926	-1.4471	-1.7866
			Max	27.8457	28.4419	28.1556
After	20090430	: 20101230	Ave	9.4210	10.5503	10.8619
			Min	2.2752	3.4242	3.3852
			Max	19.8608	20.8870	20.6085

The numbers under each column labeled Heston, GB2, and MLN are average levels of \overline{RRA}_t across observation periods under the assumed risk-neutral density models. \overline{RRA}_t is a sample average of $RRA_t(X_t)$ across levels of moneyness defined as X_t/F_t for each observation month t .