Netspar Network for Studies on Pensions, Aging and Retirement

Bas Werker

Mark-Jan Boes

Feike Drost

Nonparametric risk-neutral joint return and volatility distributions

Discussion Paper 2005 - 011
July 2007 (replacing last version of May 2005)

Nonparametric risk-neutral joint return and volatility distributions

Mark-Jan Boes [*], Feike C. Drost [†] and Bas J.M. Werker [‡§]
VU University Amsterdam and *Tilburg University*

First version: May 2005
This version: July 2007

Abstract

We propose a nonparametric technique to estimate risk-neutral volatility distributions. Our method does not need to specify a parametric risk-neutral jump-diffusion for returns and volatilities, nor do we need observations on volatility derivatives. The method uses (daily) observations on plain vanilla options only. Using S&P-500 data, we confirm a negative volatility risk premium, but find four additional results. First, the nonparametric risk-neutral volatility distribution is right-skewed. This is consistent with a larger aversion towards unexpected positive volatility shocks. The right-skewness is most pronounced when overall volatility is high. Secondly, our method also identifies the joint risk-neutral return and volatility distribution and we find negative (still risk-neutral) correlation. Third, while we confirm overall negative risk-neutral return skewness, we also find that this skewness disappears in periods of decreasing volatility. Finally, our method allows to assess the fit of parametric models. For example, concerning the popular Heston (1993) model we show that it fails to describe simultaneously the risk-neutral return and volatility distributions.

Keywords: Return skewness, Volatility risk premium, Volatility skewness.
JEL Classification: G12, G13.

---

[*] Department Finance and Financial Sector Management, VU University Amsterdam, De Boelelaan 1105, 1081 HV, Amsterdam, The Netherlands. E-mail: mboes@feweb.vu.nl

[†] Econometrics and Finance Group, CentER, Tilburg University, P.O. Box 90153, 5000 LE, Tilburg, The Netherlands. E-mail: F.C.Drost@TilburgUniversity.nl

[‡] Econometrics and Finance Group, CentER, Tilburg University, P.O. Box 90153, 5000 LE, Tilburg, The Netherlands. E-mail: Werker@TilburgUniversity.nl

[§] The authors gratefully acknowledge support of ABN-Amro Asset Management in conducting this research. Moreover, we are grateful to Peter Bloomberg, René Garcia, Jens Jackwerth, Eric Renault, and Mingxin Xu for several useful remarks.

1 Introduction

During the past decades, a considerable amount of financial research has been devoted to the informational content of derivative prices. These prices depend on one or more underlying financial factors and, therefore, price changes give information about the (risk-neutral) stochastic evolution of these factors. This information is not only used in academic research but also in the everyday practice of risk management, investment strategies, and monetary policies. In this paper, we focus on the information revealed by option prices about risk-neutral return *and* volatility distributions. There exists an extensive literature on inference concerning the risk-neutral density of stock prices/returns. However, the present paper is the first to study the information contained in plain vanilla option prices concerning the joint density of volatility, jointly with returns, in a nonparametric way.

The risk-neutral density is also known as the state price density or the implied density. The name "state price density" derives from the insight that the risk-neutral density determines continuous Arrow-Debreu state prices. As option payoffs depend in a nonlinear way on the underlying states, their prices can be used to extract risk-neutral return distributions in a nonparametric way. One popular device to extract this risk-neutral return distribution from option prices is based on Breeden and Litzenberger (1978). That paper shows that the risk-neutral return density can be obtained as the second derivative of the call option pricing function with respect to the strike price. In practice, only a finite number of options is available. Since the work of Rubinstein (1994) and Jackwerth and Rubinstein (1996) numerous papers have appeared that provide different approaches to this problem. For an overview we refer to Coutant, Jondeau, and Rockinger (1998), Jackwerth (1999), and Panigirtzoglou and Skiadopoulos (2004). The before-mentioned methods aim to extract the implied state price density at a given single point in time for a given horizon. Panigirtzoglou and Skiadopoulos (2004) investigates the dynamics of the implied distributions and provides algorithms that make the results applicable to areas like option pricing and risk management.

Using the above-mentioned methods, the evolution of risk-neutral return distributions has been studied extensively. In particular, Jackwerth and Rubinstein (1996) shows that, before the 1987 stock market crash, both the risk-neutral return distribution and the objective distribution are close to log-normal (over a 1-month horizon). However, after the 1987 crash the objective distribution still appears log-normal but the shape of the implied distribution has changed considerably. Weinberg (2001) for US option data, Anagnou, Bedendo, Hodges, and Tompkins (2002) for the UK option market, and Bliss and Panigirtzoglou (2004) using US and UK option data, find a typical post-crash shape of the implied risk-neutral distribution. Bates (2000) gives three possible explanations for the change in the risk-neutral distribution after the 1987 crash: a change in view on the underlying process of the index, a change in aggregate risk aversion, and mispricing of (far) out-of-the-money put options. This second explanation is considered in a number of papers that estimate the risk aversion of a representative investor as implied by the risk-neutral return distribution and an estimated objective distribution. The resulting implied risk aversion function seems to be inconsistent with financial theory. For example, the risk aversion functions estimated in Jackwerth (2000) imply that investors are more risk averse at high and low levels of wealth. Aït-Sahalia and Lo (2000) find decreasing but non-monotonic implied risk aversion functions as wealth increases. Brown and Jackwerth (2001) calls the phenomenon that implied risk aversion functions do not fit the requirements of economic theory even "the pricing kernel puzzle".

The above methods all consider risk-neutral return distributions only. At the same time, they mostly recognize that risk-neutral volatility distributions are important as well to test financial theory and for risk management. So far, the papers that study risk-neutral volatility distributions use parametric methods, as for example in Pan (2002). Earlier work of, for instance, Canina and Figlewski (1993) and Lamoureux and Lastrapes (1993) investigates the informational content of Black-Scholes implied volatilities. Later studies of, among others, Britten-Jones and Neuberger (2000) and Jiang and Tian (2005) examine model-free measures of risk-neutral realized volatility. However, these studies mostly focus on the *expected* value of future *realized* volatility while the focus of our study is on the risk-neutral *distribution* of future *spot* volatility.

An obvious approach to nonparametric risk-neutral volatility distributions would be to extend the Breeden and Litzenberger (1978) result to volatility derivatives. However, such approach would be complicated by the fact that these volatility derivatives are not very liquidly traded and that they usually have a pay-off scheme in terms of averages of Black-Scholes implied volatilities and not (unobserved) spot volatility. In contrast, the present paper extracts information on the joint risk-neutral density of returns and spot volatility from plain vanilla option prices without making parametric assumptions. The main idea relies on the fact that *before* maturity plain vanilla options have a value which depends directly on spot volatility. More precisely, under the assumption that option prices depend on moneyness, spot volatility, and time-to-maturity, we provide a nonparametric (and thus model free) methodology that gives the empirical risk-neutral distribution of both asset returns and spot volatilities over a given horizon. While the risk-neutral volatility distributions are interesting per se for financial theory, they can also be used to test parametric stochastic volatility models and to obtain nonparametric estimates of prices of derivatives written on volatility, like variance swaps.

In particular, we document a number of interesting new facts about risk-neutral distributions. First, it is well-known that risk-neutral volatility distribution is shifted to the right, compared to the objective one. This is one consequence of negative volatility risk-premiums. Next to reestablishing this result, we also find that the risk-neutral volatility distribution is right-skewed. While a right-skewed risk-neutral volatility distribution is implied by several well-known parametric models, we obtain this result in a fully nonparametric way. Moreover, we establish that this skewnessis most pronounced in periods of high volatility. A possible explanation could be a more extreme aversion (of the representative agent) to positive volatility shocks when volatility is already high. Secondly, as our method identifies in fact the *joint* risk-neutral return and volatility distribution, it allows us to study the risk-neutral dependence of returns and volatility. Under the objective probability measure negative return/volatility correlation is well established. This is sometimes alluded to as the "leverage effect". We find that also the risk-neutral correlation is negative. Thirdly, we confirm overall negative skewness for the risk-neutral return distribution. However, we additionally establish that this skewness disappears when volatilities are decreasing. More precisely, the conditional risk-neutral distribution of returns given that future volatility is below average exhibits no significant skewness. Apparently, decreasing volatility reduces the risk premium for crashes.

Next to the above empirical findings, our methodology can be used to assess the fit provided by parametric (jump)diffusion models for asset prices and volatilities. A formal comparison of competing specification is outside the scope of this paper, but we do show that the popular Heston (1993) model is sufficiently flexible to describe the risk-neutral return distributions accurately. However, it does so by severely overestimating the volatility of volatility. As a result it fails in *jointly* describing both return and volatility distributions. Obvious parametric alternatives would be models that incorporate jumps, possibly both in asset prices and volatilities. Examples would be the models studied in Amin (1993), Scott (1997), Andersen and Andreasen (2000), Kou (2002), and Pan (2002). In general the quest for empirically validated parametric jump-diffusion models seems to be open. In this respect note that the joint risk-neutral return volatility distribution estimated by our method, Figure 10, indicates non-convex isodensity sets.

The rest of the paper is organized as follows. Section 2 presents our methodology for obtaining the joint risk-neutral distribution of returns and volatilities. Sections 3.1 and 3.2 show how our approach relates to existing methods that yield nonparametric estimates of risk-neutral return distributions. Moreover, Section 3.3 illustrates the scope of our method in an hypothesized Heston (1993) world. In Section 4 our method is applied to S&P-500 data, which leads to the before-mentioned new insights concerning risk-neutral volatility distributions. Section 5 concludes.

2 Nonparametric estimation of risk-neutral return/volatility distributions

We present our methodology for extracting the conditional (given the current volatility level) risk-neutral return and volatility distribution from plain vanilla option prices. Nowadays, nonparametric estimates of risk-neutral distributions are usually based on butterfly spreads or on the second derivative of an estimated relation between option prices and strike prices, see Breeden and Litzenberger (1978). Such approaches, by construction, only lead to risk-neutral *return* distributions. In case derivatives whose pay-off depends on volatility were traded, analogous approaches could be used to derive the risk-neutral volatility distribution. However, this is generally not feasible as such derivatives are not (liquidly) traded. Our approach is based on the straightforward observation that standard plain vanilla options have *before* maturity a value that does depend on both the asset's price and its volatility and, thus, can be considered an implicit derivative on spot volatility.

2.1 Identifying risk-neutral volatility distributions

Consider a financial market with an asset (index) whose price at time t is denoted by S_t. We assume that, next to the asset's price, there is a second underlying factor in the financial market that we identify with (spot) volatility and we denote it by σ_t.

Fix a horizon $\Delta > 0$. The return[1] of the asset over the interval $(t, t+\Delta]$ is written as $R_{t:t+\Delta} = \log S_{t+\Delta}/S_t$. For expository reasons, we abstract from possible dividends in this section. It is well-known that interest rates only marginally affect plain vanilla option prices and, therefore, we assume for that interest rates are constant at a fixed level r. The induced excess return on the asset S is denoted by $\tilde{R}_{t:t+\Delta} = R_{t:t+\Delta} - r\Delta$. Finally, the risk-neutral distribution (joint for the price process S and the volatility process σ) is denoted by $\mathbb{Q}$.

---

[1] All (excess) returns and interest rates are taken to be continuously compounded in this paper.

Let $\mathcal{F}_t$ denote all the information available in the financial market at time t. We are interested in the risk-neutral, i.e., under $\mathbb{Q}$, distribution of $(S_{t+\Delta}, \sigma_{t+\Delta})$ conditional on $\mathcal{F}_t$. Throughout the paper we assume that the risk-neutral process of prices and volatilities is Markov and homogeneous with respect to the initial price level. Formally, we impose the following.

Assumption 1 The conditional risk-neutral distribution of $\left(\tilde{R}_{t:t+\Delta}, \sigma_{t+\Delta}\right)$ given $\mathcal{F}_t$, is equal to the conditional distribution given σ_t, i.e., for all $\Delta > 0$,

$$\mathcal{L}_{\mathbb{Q}}\left(\begin{bmatrix} \tilde{R}_{t:t+\Delta} \\ \sigma_{t+\Delta} \end{bmatrix} |\mathcal{F}_t\right) = \mathcal{L}_{\mathbb{Q}}\left(\begin{bmatrix} \tilde{R}_{t:t+\Delta} \\ \sigma_{t+\Delta} \end{bmatrix} |\sigma_t\right). \tag{2.1}$$

□

Assumption 1 is, for instance, satisfied in a standard continuous time stochastic volatility model where, under the risk-neutral distribution,

$$\begin{aligned} \mathrm{d}S_t &= rS_t\mathrm{d}t + \sigma_t S_t \mathrm{d}W_{1t}, && (2.2) \\ \mathrm{d}\sigma_t &= \alpha(\sigma_t)\mathrm{d}t + \beta(\sigma_t)\mathrm{d}W_{2t}, && (2.3) \end{aligned}$$

for some appropriate functions α and β and with W_1 and W_2 two, possibly correlated, Brownian motions. Also, Assumption 1 does not rule out jumps: neither in the underlying stock price, nor in the volatility process. For instance, W_1 and W_2 in (2.2) and (2.3) may also be correlated Lévy processes.

Condition (2.1) stipulates that, for the risk-neutral distribution of returns and future volatilities, the current volatility level is the only relevant state variable. Other state variables can easily be added from a theoretical point of view, as long as the curse of dimensionality doesn't hamper the nonparametric estimation in Section 2.2. In particular, we assume that the current stock price level does not affect the conditional return distribution. Differently stated, the price process is homogeneous with respect to its level, a property studied in Renault (1997). In a model like (2.2)–(2.3) this amounts to saying that α and β do not depend on the stock price level S_t. It does not exclude leverage, i.e., W_1 and W_2 may be correlated.

Theorem 2.1 below allows us to identify the joint risk-neutral distribution of $(\tilde{R}_{t:t+\Delta}, \sigma_{t+\Delta})$, conditional on the current volatility level σ_t from plain vanilla option prices. This result, more precisely (2.8) below, is the main theoretical contribution of the present paper. Recall that the relations (2.6)–(2.7) have already been used by, for example, Aït-Sahalia and Lo (1998), to identify the marginal risk neutral return distribution, that is, $\int_v q_\Delta(z, v|\sigma)\mathrm{d}v$ in our notation. Theorem 2.1 is operationalized in the rest of this paper as follows: To obtain the conditional risk-neutral return/volatility density at horizon Δ, i.e., q_Δ in (2.6)–(2.8) below, we first estimate nonparametrically the option price c. Subsequently, we solve equations (2.6)–(2.8); see Section 2.2.

Theorem 2.1 *Under Assumption 1 the relative (with respect to the underlying value) option price c_H is such that the time t price $C_t(K,H)$ of a plain vanilla call with exercise price K, maturity $t+H$, and moneyness*

$$m_t = \frac{\exp(-rH)K}{S_t}, \tag{2.4}$$

satisfies

$$C_t(K, t+H) = S_t c_H(m_t|\sigma_t). \tag{2.5}$$

Moreover, the joint risk-neutral density $q_\Delta(z,v|\sigma)$ of $(\tilde{R}_{t:t+\Delta}, \sigma_{t+\Delta})$, conditional on the current volatility level $\sigma_t = \sigma$, satisfies

$$\int_z \int_{v>0} q_\Delta(z,v|\sigma)\,dz\,dv = 1 \text{ and } \int_z \int_{v>0} \exp(z) q_\Delta(z,v|\sigma)\,dz\,dv = 1, \tag{2.6}$$

$$\frac{\partial^2}{\partial m^2} c_\Delta(m|\sigma) = \frac{1}{m} \int_{v>0} q_\Delta(\log m, v|\sigma)\,dv, \tag{2.7}$$

and

$$c_H(m|\sigma) = \int_z \int_{v>0} \exp(z) c_{H-\Delta}(m\exp(-z)|v) q_\Delta(z,v|\sigma)\,dz\,dv, \tag{2.8}$$

for all $m > 0$, $H \geq \Delta$, and $\sigma > 0$.

Proof: The first claim of the theorem is well-known and follows from

$$\begin{aligned} C_t(K,t+H) &= \mathrm{E}^{\mathbb{Q}}\left\{\exp(-rH)\max\left\{S_{t+H} - K, 0\right\}|\mathcal{F}_t\right\} \\ &= S_t \mathrm{E}^{\mathbb{Q}}\left\{\max\left\{\exp(\tilde{R}_{t:t+H}) - \frac{\exp(-rH)K}{S_t}, 0\right\}\middle|\, \mathcal{F}_t\right\} \\ &= S_t \int_{z>\log m} (\exp(z) - m_t) \int_{v>0} q_H(z,v|\sigma_t)\mathrm{d}v\mathrm{d}z \qquad (2.9) \\ &= S_t c_H(m_t|\sigma_t), \end{aligned}$$

where the third equality is a direct consequence of Assumption 1. This result implies that option prices are homogeneous under our assumptions.

Relation (2.6) merely states that q is a risk-neutral density for the underlying stock (recall that the argument z refers to a continuously compounded excess return). Moreover, (2.7) is the well-known Breeden-Litzenberger result and follows from setting $H = \Delta$ and differentiating (2.9) twice with respect to moneyness.

To prove (2.8), observe that the time t price of an option with exercise price K and maturity T also satisfies, for $t \leq t+\Delta \leq t+H$,

$$C_t(K,t+H) = \mathrm{E}^{\mathbb{Q}}\left\{\exp(-r\Delta)C_{t+\Delta}(K, t+\Delta+(H-\Delta))|\,\mathcal{F}_t\right\}.$$

Using (2.5) to rewrite both $C_t(K,t+H)$ and $C_{t+\Delta}(K,t+\Delta+(H-\Delta))$, we find

$$S_t c_H(m_t|\sigma_t) = \mathrm{E}^{\mathbb{Q}}\left\{\exp(-r\Delta)S_{t+\Delta}c_{H-\Delta}(m_{t+\Delta}|\sigma_{t+\Delta})|\,\mathcal{F}_t\right\}.$$

Observing

$$m_{t+\Delta} = \frac{\exp(-r(H-\Delta))K}{S_{t+\Delta}} = m_t \exp(-\tilde{R}_{t:t+\Delta}),$$

implies

$$c_H(m_t|\sigma_t) = \mathrm{E}^{\mathbb{Q}}\left\{\exp(\tilde{R}_{t:t+\Delta})c_{H-\Delta}(m_t\exp(-\tilde{R}_{t:t+\Delta})|\sigma_{t+\Delta})\Big|\,\mathcal{F}_t\right\}. \tag{2.10}$$

Invoking the Markov property in Assumption 1 once more, we find that (2.10) can be rewritten, for all $H \geq \Delta$, $m > 0$, and $\sigma > 0$ as (2.8). $\square$

As discussed above, the relations (2.6)–(2.8) can be used to obtain a nonparametric estimate of the conditional risk-neutral return/volatility density. Observe that the key point of Theorem 2.1 is the use of options with a maturity H that is strictly larger than the horizon for the density of interest Δ. Indeed, Relation (2.8) is not informative for the risk-neutral *volatility* density in case $H = \Delta$ as $c_0(m|v) = \max\{1-m, 0\}$ does not depend on v.

2.2 Nonparametric estimation

In our setup, under Assumption 1, plain vanilla option prices (relative to the value of the underlying) depend on moneyness m_t, time-to-maturity $T-t$, and spot volatility σ_t. We follow Aït-Sahalia and Lo (1998) to come up with a nonparametric estimate of this relation, i.e., an estimate of the relative option price c_H. Once this estimate is available, we solve the equations in Theorem 2.1) for the joint risk-neutral return and volatility density q_Δ. Consistency of this density estimate is obtained when c_H is estimated consistently. Obviously, the price and volatility process must satisfy regularity conditions for this consistency to hold. In particular, returns and volatilities must be sufficiently stationary and mixing. We refer to Aït-Sahalia and Lo (1998) for details.

More precisely, given option prices $C_t(K,T)$, we rewrite (2.5) as

$$\frac{C_t(K, t+H)}{S_t} = c_H(m_t|\sigma_t) = BS_H\left(m_t, IV_H\left(m_t|\sigma_t\right)\right), \tag{2.11}$$

with

$$BS_H\left(m,\sigma\right) = \Phi\left(\frac{\ln(m)+\sigma^2 H/2}{\sigma\sqrt{H}}\right) - m\Phi\left(\frac{\ln(m)-\sigma^2 H/2}{\sigma\sqrt{H}}\right).$$

Here BS_H stands for the maturity-H Black-Scholes formula, normalized by the current stock price level, and IV denotes the option's implied volatility. In our setup, the implied volatility depends on moneyness m_t, time-to-maturity H, and current spot volatility σ_t. The function IV, and hence the function c_H, is estimated using standard nonparametric kernel regression following Aït-Sahalia and Lo (1998). As spot volatility levels σ_t are not observed, we rely on daily realized volatilities to approximate these. This leads to an estimate $\hat{c}_H$.

Subsequently, our estimate of the joint conditional (on the current volatility level $\sigma_t = \sigma$) risk-neutral return and volatility density q_Δ is obtained by solving the equations (2.6)–(2.8) for q_Δ using the estimated option price function $\hat{c}_H$. The equations (2.6)–(2.8) do not allow for an explicit analytical solution and thus we actually solve a discrete approximation. More precisely, we choose a grid $-\infty < z_0 < z_1 < z_2 < \ldots < z_N < \infty$ for possible asset excess returns and a grid $0 < v_0 < v_1 < \ldots < v_K < \infty$ for possible spot volatility levels.

Subsequently, (2.6) is discretized as

$$\begin{aligned}
\sum_{n=1}^{N}\sum_{k=1}^{K}\hat{q}_{\Delta}\left(\frac{z_{n-1}+z_n}{2},\frac{v_{k-1}+v_k}{2}\middle| v_i\right)dz_n dv_k &= 1, \qquad (2.12)\\
\sum_{n=1}^{N}\sum_{k=1}^{K}\exp\left(\frac{z_{n-1}+z_n}{2}\right)\hat{q}_{\Delta}\left(\frac{z_{n-1}+z_n}{2},\frac{v_{k-1}+v_k}{2}\middle| v_i\right)dz_n dv_k &= 1, \qquad (2.13)
\end{aligned}$$

with $dz_n = z_n - z_{n-1}$ and $dv_k = v_k - v_{k-1}$ and for $i = 0, \ldots, K$. Discretizing the log-moneyness $\log(m)$ on the same grid as returns, Equation (2.7) leads to

$$\begin{aligned}
&\frac{\hat{c}_{\Delta}\left(\exp\left(\frac{z_{l-2}+z_{l-1}}{2}\right)\middle| v_i\right) - 2\hat{c}_{\Delta}\left(\exp\left(\frac{z_{l-1}+z_l}{2}\right)\middle| v_i\right) + \hat{c}_{\Delta}\left(\exp\left(\frac{z_l+z_{l+1}}{2}\right)\middle| v_i\right))}{\left(\exp\left(\frac{z_{l-1}+z_l}{2}\right)-\exp\left(\frac{z_{l-2}+z_{l-1}}{2}\right)\right)\left(\exp\left(\frac{z_l+z_{l+1}}{2}\right)-\exp\left(\frac{z_{l-1}+z_l}{2}\right)\right)} \qquad (2.14)\\
&\qquad = \exp\left(-\frac{z_{l-1}+z_l}{2}\right)\sum_{k=1}^{K}\hat{q}_{\Delta}\left(\frac{z_{l-1}+z_l}{2},\frac{v_{k-1}+v_k}{2}\middle| v_i\right)dv_k,
\end{aligned}$$

for $l = 1, \ldots, N$ and $i = 0, \ldots, K$. We make sure that moneyness is discretized equally spaced. Finally, Equation (2.8) is discretized as

$$\begin{aligned}
&\hat{c}_H\left(\exp\left(\frac{z_{l-1}+z_l}{2}\right)\middle| v_i\right) \qquad (2.15)\\
&= \sum_{n=1}^{N}\sum_{k=1}^{K}\exp\left(\frac{z_{n-1}+z_n}{2}\right)\hat{c}_{H-\Delta}\left(\exp\left(\frac{z_{l-1}+z_l}{2}-\frac{z_{n-1}+z_n}{2}\right)\middle| \frac{v_{k-1}+v_k}{2}\right)\\
&\qquad \times \hat{q}_{\Delta}\left(\frac{z_{n-1}+z_n}{2},\frac{v_{k-1}+v_k}{2}\middle| v_i\right)dz_n dv_k,
\end{aligned}$$

for $l = 1, \ldots, N$, $i = 0, \ldots, K$, and a number of possible maturities $H = H_0, \ldots, H_P$ with $\Delta = H_0 < H_1 < \ldots < H_P$. Recall $\hat{c}_0(m|v) = \max\{1-m, 0\}$.

The equations (2.12)–(2.15) provide a linear system in the unknown quantities

$$\hat{q}_{\Delta}\left(\frac{z_{n-1}+z_n}{2},\frac{v_{k-1}+v_k}{2}\middle| v_i\right), \qquad (2.16)$$

for $n = 1, \ldots, N$, $k = 1, \ldots, K$, and $i = 0, \ldots, K$. This system is solved numerically by means of a standard least-squares algorithm. To avoid arbitrage possibilities, we impose that the probabilities in (2.16) are all positive. Consistency of the resulting estimate follows from consistent estimation of the relative option pricing function c_H under conditions that parallel those of, e.g., Aït-Sahalia and Lo (1998).

Since $\hat{c}_H$ is a nonparametric estimate, the numerical integral approximation (2.15) quite often leads to non-smooth densities. Therefore, (2.15) is complemented with a smoothness condition on the solution q_Δ in terms of the second derivative of q_Δ with respect to both returns and volatilities. Jackwerth and Rubinstein (1996) impose a similar condition but, of course, with respect to returns only. The approximations (2.12)–(2.14) are imposed as (linear) constraints in the penalized optimization.

3 Relation with existing methods

The literature mainly presents two approaches to get nonparametric estimates of risk-neutral *return* distributions: the method discussed in Jackwerth and Rubinstein (1996) and the kernel-estimation based approach of Aït-Sahalia and Lo (1998). The main difference between these two methods is that the first provides a different risk-neutral return density for each different cross-section of option prices while the second results in a single risk-neutral return distribution. With respect to the conditional information used to determine the risk-neutral return distribution, we argue below that our method is effectively in between these two approaches. More importantly, however, our method offers the additional advantage of estimating the risk-neutral *volatility* distribution, including the risk-neutral return-volatility dependence structure. For instance, our method applies to the popular Heston (1993) stochastic volatility model, but without any parametric assumptions. Also, our method does not impose nor exclude jumps in prices and/or volatilities. The present section discusses briefly the relation of our approach with both alternatives mentioned above and the performance of our approach in a simulated Heston (1993) world (Section 3.3).

3.1 Fully nonparametric methods

Jackwerth and Rubinstein (1996) use option prices observed at a given date to infer risk-neutral probabilities of returns for a given future date. Essentially the (discrete) risk neutral probability distribution of returns is determined such that all observed option prices today are within the bid-ask bounds. In the notation of Section 2, this method provides an estimate of $\mathcal{L}_{\mathbb{Q}}\left(R_{t:t+\Delta}|\mathcal{F}_t\right)$ without any further restrictions on the conditioning information set $\mathcal{F}_t$. In particular, as noted in Aït-Sahalia and Lo (1998), no time-consistency is imposed in this method. As a result, the estimates of risk-neutral return probabilities will vary over time. The method is thus fully nonparametric, but only few observations (i.e., only options traded on a given day with a particular maturity) can be used in the estimation. Our method builds on Assumption 1 which identifies the current spot volatility as the only relevant state variable for predicting risk-neutral return distributions and imposes time-consistency. As mentioned before, such assumption is common in most parametric stochastic volatility models.

3.2 Kernel-estimation based methods

Aït-Sahalia and Lo (1998) base their nonparametric risk-neutral return distributions on the Breeden and Litzenberger (1978) result that the risk-neutral return distribution is proportional to the second derivative of plain vanilla call prices with respect to the exercise price. The functional relation between the option prices and relevant explanatory variables is estimated using nonparametric kernel regression of Black-Scholes implied volatilities on the futures price associated with the underlying asset, the exercise price, and time-to-maturity. Aït-Sahalia and Lo (1998) consider other (vectors of) explanatory variables as well, but futures price, exercise price, and time-to-maturity come out as their preferred choice. Next to moneyness and time-to-maturity, we add the current spot volatility level σ_t as relevant state variable.

With respect to our method, Aït-Sahalia and Lo (1998) allows for time-varying volatility where the current *level* of the stock price induces a certain volatility. However, in cases of stochastic volatility as a separate state variable, the method does not lead to risk-neutral return distributions conditional on a certain volatility level, but to unconditional (with respect to volatility) return distributions. Since we consider spot volatility as a separate state variable, much in line with the Heston (1993) model, we impose that options' implied volatilities depend on futures prices and exercise prices through moneyness alone (see, also, Renault (1997)). Observe that in case of a stochastic volatility model, the Aït-Sahalia and Lo (1998) method will pick up of some of the stochastic volatility effects as stock prices and volatilities are (negatively) correlated. This will be discussed in more detail in the next section where we consider both our method and the Aït-Sahalia and Lo (1998) method in a theoretical Heston (1993) world. Simulation results for the Hull and White (1987) stochastic volatility model and the Stein and Stein (1991) stochastic volatility model are available upon request and comparable to the results reported below for the Heston (1993) model.

3.3 Simulation: Risk-neutral volatility distributions in the Heston model

Heston (1993) presents a parametric stochastic volatility model which is especially useful for calculating derivative prices due to the fact that the characteristic function of the risk-neutral return distribution is known in analytical form. This latter property is a demonstration of the fact that the Heston (1993) model belongs to the class of affine jump-diffusions (in the spirit of Duffie and Kan (1996)). The Heston (1993) model is given by the dynamics, under the risk-neutral probability measure and with $\lambda_t = (\mu - r)/\sigma_t$,

$$\begin{aligned} \mathrm{d}S_t &= (\mu - \lambda_t \sigma_t) S_t \mathrm{d}t + \sigma_t S_t \mathrm{d}W_t^S, \\ \mathrm{d}\sigma_t^2 &= (\kappa + \eta) \left(\frac{\kappa}{\kappa + \eta} \sigma^2 - \sigma_t^2 \right) \mathrm{d}t + \sigma_\sigma \sigma_t \mathrm{d}W_t^\sigma, \\ \mathrm{Cov}\left\{ \mathrm{d}W_t^S, \mathrm{d}W_t^\sigma \right\} &= \rho \mathrm{d}t. \end{aligned} \tag{3.1}$$

Under the objective probability measure, the dynamics are obtained by setting both the stock risk premium λ_t and the volatility risk premium η to zero. For given parameters, the risk-neutral return distribution is known in closed form as the inverse of its Fourier transform, see Heston (1993). The risk-neutral conditional distribution of spot volatility $\sigma_{t+\Delta}$, given σ_t, is also known in analytical form, see Cox, Ingersoll, and Ross (1985). Moreover, the Heston (1993) model satisfies the Markov condition in Assumption 1.

In order to study the performance of our proposed method, we simulate five years (1260 trading days) of daily S&P-500 data using the Pan (2002) parameters, i.e., in the notation of (3.1): $\kappa = 6.4$, $\sigma = 12.4\%$, $\sigma_\sigma = 30.0\%$, $\rho = -53\%$, and $\eta = -3.1$. Observe that the Pan (2002) estimates refer to a period of low overall volatilities. The interest rate is fixed at a constant annual level of 4.0%, the initial volatility level σ_0 is set equal to the unconditional mean $\sigma = 12.4\%$, and the expected instantaneous return is fixed at $\mu = 10.0\%$. All parameters are annualized. The diffusion (3.1) is simulated on a 2 minute frequency using an Euler discretization. Using the simulated price levels and volatilities, we calculate analytical option prices with the Heston (1993) formula. For each day we consider options with on average four different maturities and, for each maturity, forty different strike prices are considered. The actual number of options available each day varies as we have implemented a stylized option introduction scheme, in line with the methodology used by most exchanges. We apply the method as described in Section 2 where we are interested in the risk-neutral conditional return and volatility distribution over a period of one month, i.e., we set $\Delta = 1/12$. Spot volatility levels are filtered using realized volatilities based on a 10 minute frequency. Concerning option prices, we use these at a daily frequency only. These choices are the same as in Section 4 below.

Figure 1 shows our estimated risk-neutral return distributions, conditional on various initial volatility levels. The low level of $\sigma_t = 8.8\%$ corresponds to the first quartile of the objective volatility distribution (as measured by intraday realized volatilities), while the high level of $\sigma_t = 14.5\%$ corresponds to the third quartile. Recall that $\sigma_t = 12.4\%$ corresponds to the (objective) average volatility. Moreover, we provide the Aït-Sahalia and Lo (1998) estimate, which is unconditional with respect to the initial spot volatility level. It is provided as a benchmark and, indeed, lies in between our estimates for low and high initial volatility and is close to the estimate conditional on an average initial volatility level. Note, moreover, that in all cases the return distribution is somewhat left-skewed as induced by the negative return/volatility correlation parameter ρ.

A possible issue in our methodology is the use of intraday realized volatilities instead of (unobserved) actual spot volatilities. In this simulation exercise it is possible to study the effect of using intraday realized volatilities. Figure 2 shows (1) the true marginal risk-neutral return distributions, (2) the estimated marginal risk-neutral volatility distributions in case we use actual (simulated) Heston spot-volatilities and (3) our estimate using intraday realized volatilities instead of spot-volatility levels. The graphs clearly show that the effect of using intraday realized volatilities instead of the true underlying spot volatilities is negligible. Moreover, we see that our estimate closely follows the true risk-neutral volatility distribution. The graphs are conditional on an initial average volatility level $\sigma_t = 12.4\%$. For other initial volatility levels the results are comparable. Similarly, Figure 3 provides the risk-neutral volatility density for the same cases. Once more, we observe that our method succeeds in recovering the risk-neutral volatility distribution with some precision. This precision is due to the large number of observations that are available in this kind of analysis: we use, over a period of time, all traded options with more than one month to maturity. Unlike the risk-neutral return distributions, the effect of filtering spot volatilities is visible in these graphs. This shows that there are some statistically significant biases in the filtered volatilities. From a financial point of view, however, the differences are small, especially in the tails of the distributions. As an aside, note that the Pan (2002) parameters imply some right-skewness in the risk-neutral volatility distribution, although much less than we will find empirically in Section 4.

4 Empirical risk-neutral return and volatility distributions

This section provides risk-neutral return and volatility distribution estimates based on S&P-500 data. Section 4.1 discusses the data we use in some detail. In Section 4.2 we confirm negative risk-neutral return skewness and a negative volatility risk premium. Additionally, we find positive risk-neutral volatility skewness especially when volatility is high. Moreover, we see that (also risk-neutrally) returns and volatilities are negatively correlated. Finally, that section shows that the Heston (1993) model calibrated to the risk-neutral return distribution significantly overestimates the risk-neutral volatility of volatility. As our method also leads to risk-neutral return distributions conditional on *future* volatility levels, Section 4.3 presents these results and shows that in situations of decreasing volatility the negative return skewness disappears.

4.1 Data description

The empirical results in the present paper are based on European options traded on the Chicago Board Options Exchange over the period July, 1999, – December, 2003. The option data are extracted from the ABN-Amro Asset Management database and contain daily closing quotes of SPX options for all trading days in the sample period. In addition, the closing S&P-500 index levels are provided. Following Jackwerth and Rubinstein (1996), dividend rates are calculated from the actual dividends paid out by the SPX stocks. The methodology presented in Section 2 does not treat dividends explicitly but in the empirical analysis index prices are replaced by index prices discounted by the dividend rate. Finally, interpolated LIBOR rates are employed as a proxy for the risk free rate.

Following Bakshi, Cao, and Chen (1997), we only use options that satisfy a number of criteria. More precisely, we restrict attention to calls and puts that

1. have time-to-expiration greater than six calendar days,
2. have a bid price greater than or equal to 0.05$,
3. have a bid-ask spread less than or equal to 1$,
4. have a Black-Scholes implied volatility greater than zero and less than or equal to 80% (annualized),
5. satisfy the arbitrage restriction,

$$C_t(T) \geq \max\left(0, S_t e^{-q(T-t)} - K e^{-r(T-t)}\right),$$

   for call options and a similar restriction for put options. In this formula $C_t(T)$ denotes the call option price at time t for maturity T, S_t is the prevailing spot price, q the dividend rate, K the option exercise price, and r the continuously compounded intraday risk-free rate.

Table 1 provides descriptive statistics on the resulting set of options. From this table well-known patterns in implied volatilities across strikes and maturities are recognized. The volatility skew or smile is clearly present for all option categories. Unreported statistics on

CALLS		days to expiration			subtotal
	moneyness	<60	60–180	>180	
ITM	< 0.97	31.5%	27.4%	25.0%	
		12734	11877	4251	28862
ATM	0.97–1.03	22.1%	22.1%	22.7%	
		7949	6624	2350	16923
OTM	> 1.03	13.6%	11.4%	24.8%	
		12902	12879	1837	27618
	subtotal	33585	31380	8438	73403

PUTS		days to expiration			subtotal
	moneyness	<60	60–180	>180	
OTM	< 0.97	32.3%	28.0%	24.6%	
		14304	13086	4334	31724
ATM	0.97–1.03	21.9%	21.9%	22.2%	
		7947	6603	2350	16900
ITM	> 1.03	17.6%	14.7%	23.9%	
		9877	9824	1800	21501
	subtotal	32128	29513	8484	70125

Table 1: Summary statistics on SPX call and put option implied volatilities. We report average annualized implied volatilities of options on the S&P-500 index corresponding to the last tick before 3:00 PM and the total number of observations for each maturity category. The sample period is July 9, 1999, to November 27, 2003.

return data show that in the sample period the annualized standard deviation of returns equals 20.6%.

Following the literature, realized volatility is used as a proxy for (integrated) spot volatility. Continuously compounded returns over 10 minute intervals for the S&P-500 index are constructed. From the series of 10-minute returns on day t the estimated spot variance is calculated as

$$\hat{\sigma}_t^2 = \frac{1}{n} \sum_{i=1}^{n} \left\{ \log \left(S_{t+i\tau} / S_{t+(i-1)\tau} \right) \right\}^2 , \tag{4.1}$$

with τ representing a 10-minute trading period and $n = 6.5 \times 6 = 39$ the number of 10-minutes trading periods on one trading day. Figure 4 shows the estimated intraday volatilities over the complete sample period. The (annualized) volatility during the sample period varies between 3.7% and 69.0% with an (objective) average of 18.0%. Figure 4 shows that, after the turbulence in 2001 (September 11) and 2003 (start of Gulf War II), volatility has decreased to relatively low levels in 2003. Figure 5 shows the (objective) histogram of estimated volatilities. This figure confirms positive *objective* volatility skewness. In Section 4.2 we will, just as for the Heston model, nonparametrically estimate risk-neutral densities for three initial volatility levels: the average volatility level of 18.0%, the first quartile of 11.8%, and the third quartile of 19.2%. Note that the third quartile and the average volatility level are very close due to the strong volatility skewness.

4.2 Risk-neutral return and volatility densities

This section presents risk-neutral distributions for both returns and volatilities. Note that these distributions are conditional on an initial spot volatility level, as implied by Assumption 1. The dependence between *future* returns and volatilities will be discussed in Section 4.3. Figure 6 presents the risk-neutral excess return distribution as implied by our data. The initial volatility level of $\sigma_t = 18.0\%$ corresponds to the average volatility level as follows from the intraday realized volatilities. The low and high volatility levels of $\sigma_t = 11.8\%$ and $\sigma_t = 19.2\%$ correspond to the 25% and 75% quantiles, respectively. Observe that our volatility levels are higher than those in Pan (2002) as we consider the (much) more volatile 1999–2003 period, while Pan (2002) covers the January, 1989, – December, 1996, period. The figure clearly confirms negative skewness in the risk-neutral return distribution for all initial volatility levels. Alternatively, while (risk-neutral) expected excess returns are zero by definition, we see that the modes are slightly positive in all cases.

Our method also provides nonparametric risk-neutral volatility densities. For the same initial volatility levels, these are presented in Figure 7. The expected risk-neutral future spot volatility is, also using an initial average volatility level of $\sigma_t = 18.0\%$, (much) larger than this objective average. Therefore, the nonparametric results in Figure 7 confirm the existence of a negative volatility risk premium. This negative volatility risk premium results in option values higher than they would be in case of idiosyncratic volatility risk. The higher value is a compensation for unhedged volatility risk that option traders typically face because they only delta hedge their short options positions, see also Bakshi and Kapadia (2003). We also confirm that higher initial volatility leads to a right-shift in the future volatility distribution of about the same size. In addition to these long-established facts, Figure 7 depicts clear evidence of positive skewness in the risk-neutral volatility distribution.

This positive skewness is most pronounced for average or above average initial volatility levels. Another observation is that for higher initial volatility levels the risk-neutral probability of high future volatility remains significant. This indicates that, also risk-neutrally, volatility is highly persistent. The large probabilities for high volatility states are necessary to fit the prices of out-the-money put options in our sample. Otherwise stated, in order to price out-of-the-money puts correctly we need significant probability mass in low return and high volatility combinations. All these conclusions are consistent with a market aversion towards high volatility levels and an even larger aversion towards unexpected positive volatility shocks when volatility is already high to start with.

We can use our nonparametric technique to infer the accuracy of the popular Heston (1993) model. To that extent, we choose parameter values such as to provide the best (in least squares sense) fit of the estimated risk-neutral *return* distribution in Figure 6. This leads, in the notation of Section 3.3, to $\kappa + \eta = 1.47$, $\kappa\sigma^2 = 0.56$, $\sigma_\sigma = 1.12$, and $\rho = -38\%$. Note that κ and σ^2 cannot be identified separately from the risk-neutral return distribution alone. The resulting return distribution, conditional on an average initial volatility, is depicted in Figure 8. It is clear that the Heston (1993) model is capable of providing a very accurate description of risk-neutral return distribution for our sample. However, it fails in describing simultaneously the risk-neutral volatility distribution.The induced risk-neutral volatility distribution, using the same parameter values as above, is depicted in Figure 9. From this figure we see that an accurate fit of the return distribution leads to a severe overestimation of the risk-neutral volatility of volatility. Moreover, the Heston (1993) volatility distribution does not provide the correct positive risk-neutral volatility skewness that is apparent from the nonparametric estimates, despite the chosen negative instantaneous correlation parameter. Correspondingly, an optimal fit of the risk-neutral volatility distribution was determined. Not surprisingly, in that case the risk-neutral return distribution did not fit the observed one. The estimated volatility of volatility (σ_σ) was indeed much lower (0.48), but still larger than values reported in the time series literature, see for instance Eraker, Johannes, and Polson (2003). This suggests that more components in the volatility process are necessary to describe asset return and volatility distributions simultaneously. In parametric models this could be accomplished by, for instance, an additional Brownian component (see Chernov, Gallant, Ghysels, and Tauchen (2003)) or a jump component (see Broadie, Chernov, and Johannes (2007)). Once more, these are possible parametric adjustments, while our focus of interest is nonparametric estimates.

Our method also leads to joint risk-neutral return and volatility distributions which can be used to study the risk-neutral return and volatility dependence. Figure 10 graphs this joint estimate for an average initial volatility level of $\sigma_t = 18.0\%$. The graph shows that standard Gaussian and other elliptical distributions are unlikely to provide a good fit to the bivariate distribution. Moreover, we note that the bivariate risk-neutral return/volatility distribution exhibits negative dependence. While this leverage effect is well-known for objective distributions, our method establishes this nonparametrically for risk-neutral distributions. In general, however, it is difficult to asses the financial significance of these deviations from bivariate density plots. Therefore, we discuss in the next section risk-neutral return distributions conditional on *future* volatility levels.

4.3 Conditional risk-neutral return distributions

In order to assess the risk-neutral dependence structure of returns and volatilities, Figure 11 presents the return distribution conditional on three *future* level of the spot volatility. In line with previous parametric results, see, for instance, Bakshi, Cao, and Chen (1997), we find once more evidence of negative risk-neutral correlation between (excess) returns and volatility. This follows from the fact that higher future volatility levels in Figure 11 lead to negative shifts in the return distribution. At the same time, our method provides evidence that return distribution skewness depends on volatility changes. For a future volatility level of 24.0%, which is close to the risk-neutral average spot volatility, the return distribution is clearly skewed to the left. Left skewness of the risk-neutral return distribution is usually associated with a significant crash risk premium. However, for low future volatility levels the return distribution not only shifts to the right, but it loses its negative skewness and even shows some slight positive skewness. Apparently the crash risk premium (the product of both the price of crash risk and its likelihood) is higher when there is more uncertainty in the market. To the best of our knowledge, this latter result has not been established empirically before.

An obvious question is which parametric return/volatility models can produce joint return/volatility densities as depicted in Figure 11. This figure seems to indicate different behavior of the risk-neutral return distribution for moderate and high levels of volatility. This could suggest using a regime switching model. Specification of such a parametric model is outside the scope of this paper.

5 Conclusions

We present a nonparametric technique to infer risk-neutral return and volatility distributions from plain vanilla option prices. Using this technique and based on S&P-500 data, we confirm negative skewness in the risk-neutral return distribution and negative volatility risk premiums. It is important to note that we obtain these results without using a parametrically specified model, which implies that we are less prone to possible misspecification bias. At the same time, as we can assess the joint risk-neutral return and volatility distribution, we find evidence of positive skewness in the risk-neutral volatility distribution, which seems to increase with volatility levels. Moreover, conditionally on low future volatility levels, the return distribution is no longer negatively skewed but shows some slight positive skewness. These effects are consistent with volatility dependent risk premiums. Finally, our nonparametric estimates confirm negative risk-neutral return and volatility correlation.

References

Aït-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *The Journal of Finance* 53, 499–547.

———, 2000, Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9–51.

Amin, Kaushik I., 1993, Jump diffusion option valuation in discrete time, *The Journal of Finance* 48, 1833–1863.

Anagnou, Iliana, Mascia Bedendo, Stewart Hodges, and Robert Tompkins, 2002, The relation between implied and realised probability density functions, *Working Paper University of Technology Vienna.*

Andersen, Leif, and Jesper Andreasen, 2000, Jump-diffusion processes: Volatility smile fitting and numerical methods for option pricing, *Review of Derivatives Research* 4, 231–262.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical performance of alternative option pricing models, *The Journal of Finance* 52, 2003–2049.

Bakshi, Gurdip, and Nikunj Kapadia, 2003, Delta-hedged gains and the negative market volatility risk premium, *The Review of Financial Studies* 16, 527–566.

Bates, David S., 2000, Post-'87 crash fears in S&P-500 futures options, *Journal of Econometrics* 94, 181–238.

Bliss, Robert R., and Nikolaos Panigirtzoglou, 2004, Option-implied risk aversion estimates, *The Journal of Finance* 59, 407–446.

Breeden, Douglas, and Robert Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of Business* 51, 621–651.

Britten-Jones, Mark, and Anthony Neuberger, 2000, Option prices, implied processes, and stochastic volatility, *The Journal of Finance* 55, 839–866.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premiums: The evidence from the futures options, *The Journal of Finance, forthcoming.*

Brown, David P., and Jens C. Jackwerth, 2001, The pricing kernel puzzle: Reconciling index option data and economic theory, *Working Paper University of Wisconsin.*

Canina, Linda, and Stephen Figlewski, 1993, The informational content of implied volatility, *Review of Financial Studies* pp. 659–681.

Chernov, Michael, A. Ron Gallant, Eric Ghysels, and George Tauchen, 2003, Alternative models of stock price dynamics, *Journal of Econometrics* 116, 225–258.

Coutant, Sylvie, Eric Jondeau, and Michael G. Rockinger, 1998, Reading interest rate and bond future options' smiles around the 1997 french snap election, *CEPR* p. no. 2010.

Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–408.

Duffie, Darrell, and Rui Kan, 1996, A yield-factor model of interest rates, *Mathematical Finance* 6, 379–406.

Eraker, Bjørn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in returns and volatility, *Journal of Finance* 53, 1269–1300.

Heston, Steven L., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *The Review of Financial Studies* 6, 327–343.

Hull, John, and Alan White, 1987, The pricing of options on assets with stochastic volatilities, *The Journal of Finance* 42, 281–300.

Jackwerth, Jens C., 1999, Option implied risk-neutral distributions and implied binomial trees: A literature review, *Journal of Derivatives* 7, 66–82.

——— , 2000, Recovering risk aversion from option prices and realized returns, *The Review of Financial Studies* 13, 433–451.

——— , and Mark Rubinstein, 1996, Recovering probability distributions from option prices, *The Journal of Finance* 51, 1611–1631.

Jiang, George J., and Yisong S. Tian, 2005, The model-free implied volatility and its information content, *The Review of Financial Studies* 18, 1305–1342.

Kou, Steven G., 2002, A jump-diffusion model for option pricing, *Management Science* 48, 1086–1101.

Lamoureux, Christopher G., and William D. Lastrapes, 1993, Forecasting stock-return variance: Toward an understanding of stochastic implied volatilities, *The Review of Financial Studies* pp. 293–326.

Pan, Jun, 2002, The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* 63, 3–50.

Panigirtzoglou, Nikolaos, and George Skiadopoulos, 2004, A new approach to modeling the dynamics of implied distributions: Theory and evidence from the s&p 500 options, *Journal of Banking and Finance* 28, 1499–1520.

Renault, Eric, 1997, Econometric models of option pricing errors, *Advances in Economics and Econometrics: Theory and Applications, Vol. 3, D. Kreps and K. Wallis (Eds.), Cambridge University Press.*

Rubinstein, Mark, 1994, Implied binomial trees, *The Journal of Finance* 49, 771–818.

Scott, Louis O., 1997, Pricing stock options in a jump-diffusion model with stochastic volatility and interest rates: Applications of fourier inversion methods, *Mathematical Finance* 7, 413–426.

Stein, Elias, and Jeremy Stein, 1991, Stock price distributions with stochastic volatility: An analytical approach, *The Review of Financial Studies* 4, 727–752.

Weinberg, Steven A., 2001, Interpreting the volatility smile: An examination of the informational content of options prices, *Federal Reserve Board Washington*.

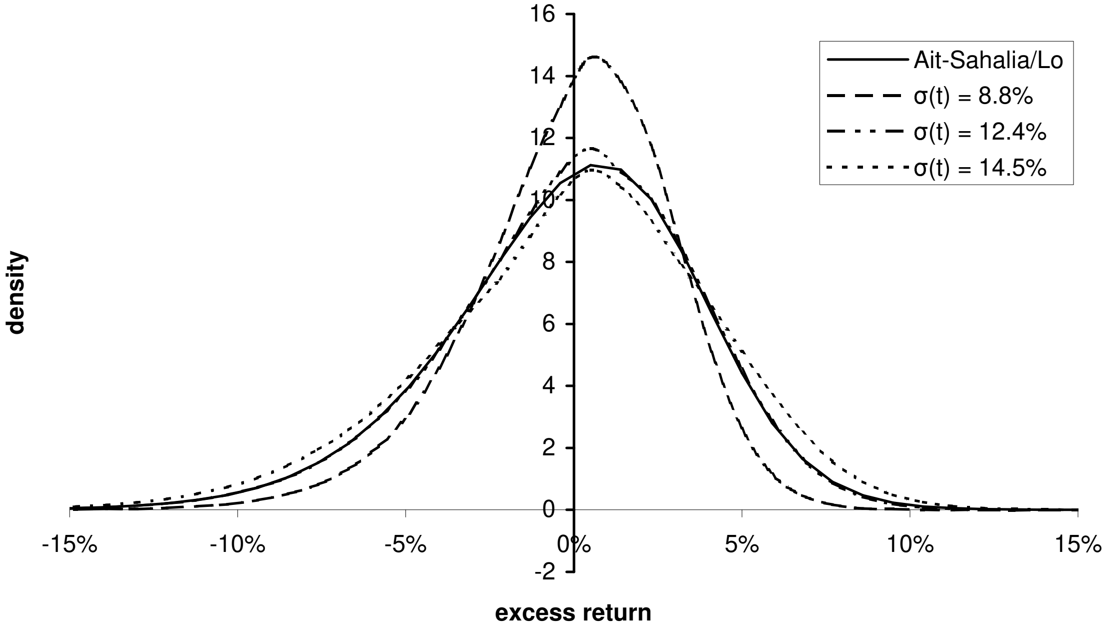

Figure 1: Estimated risk-neutral excess return distribution over a horizon of one month based on five years of simulated data in the Heston (1993) model with the Pan (2002) parameters. The distributions are conditional on the initial volatility level σ_t. The dotted-dashed, dashed, and dotted lines correspond to average, low (first quartile of unconditional objective volatility distribution), and high (third quartile) initial volatility, respectively. The solid line shows the (unconditional) Aït-Sahalia and Lo (1998) estimate.

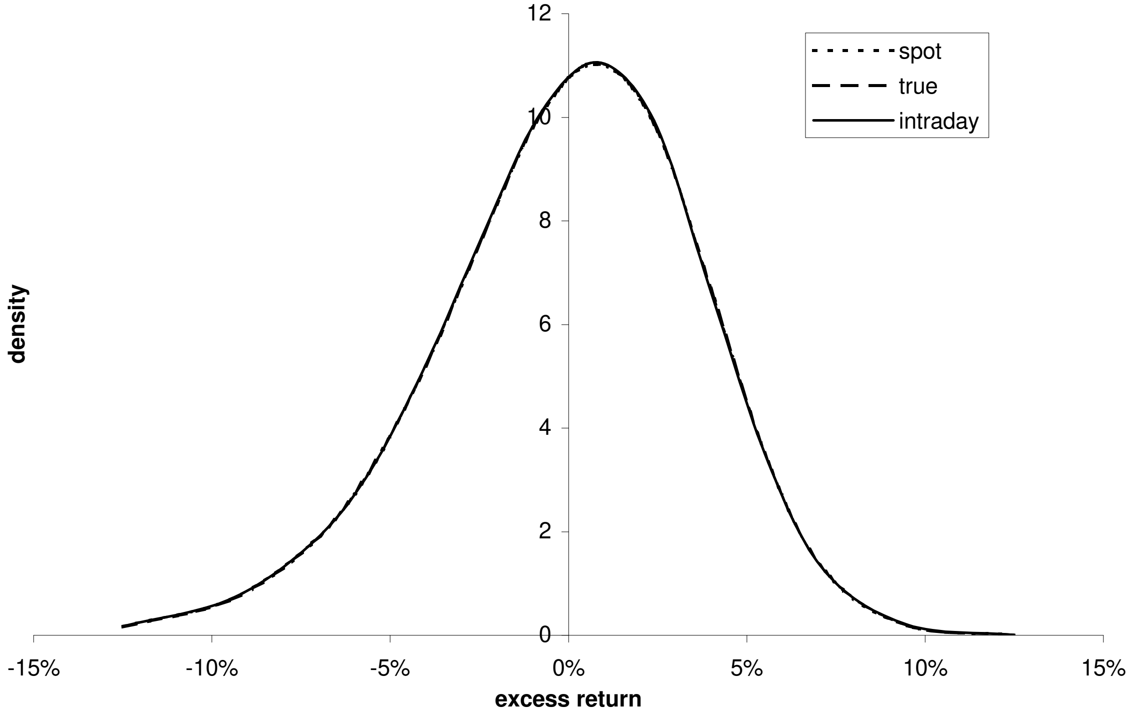

Figure 2: Estimated risk-neutral excess return distribution over a horizon of one month based on five years of simulated data in the Heston (1993) model with the Pan (2002) parameters. The solid line denotes the estimate based on realized intraday volatilities. The dotted line shows the estimate based on the actual spot volatilities. The dashed line shows the theoretical risk-neutral excess return distribution. All distributions are conditional on an average initial level of spot volatility $\sigma_t = 12.4\%$.

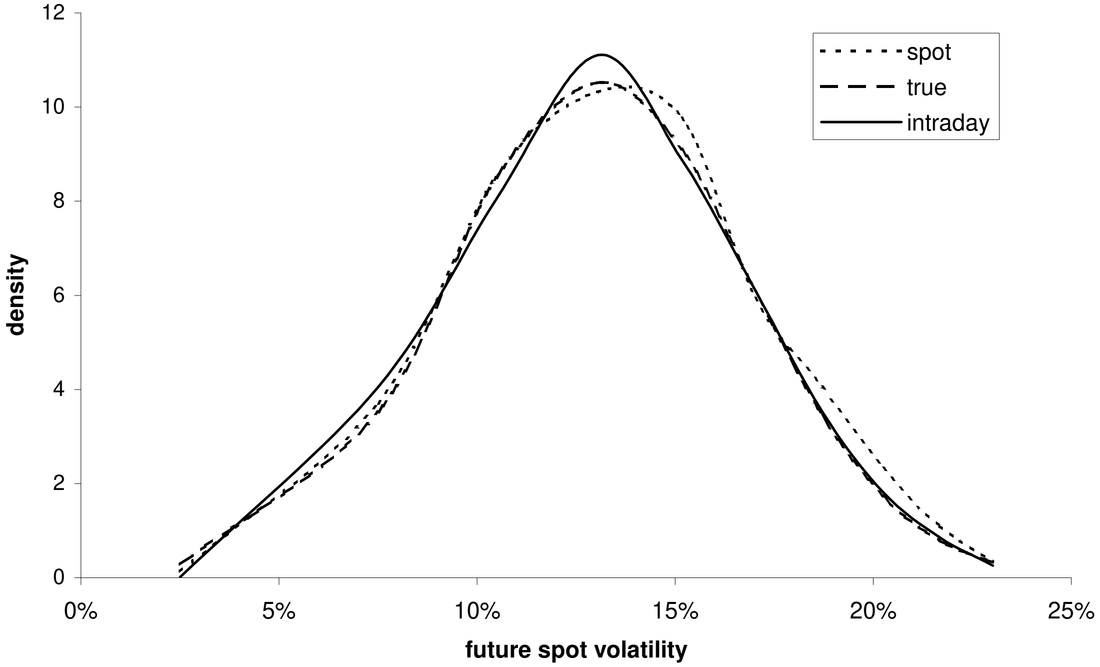

Figure 3: Estimated risk-neutral volatility distribution over a horizon of one month based on five years of simulated data in the Heston (1993) model with the Pan (2002) parameters. The solid line denotes the estimate based on realized intraday volatilities. The dotted line shows the estimate based on the actual spot volatilities. The dashed line shows the theoretical risk-neutral volatility distribution. All distributions are conditional on an average initial level of spot volatility $\sigma_t = 12.4\%$.

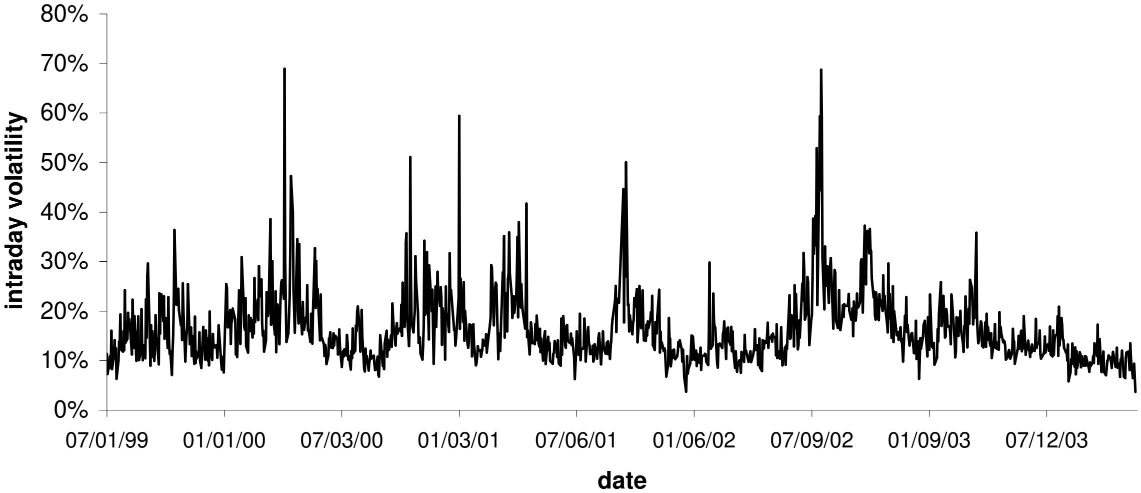

Figure 4: Estimated (realized) volatilities over the sample period July, 1999, until December, 2003, for the S&P500 index.

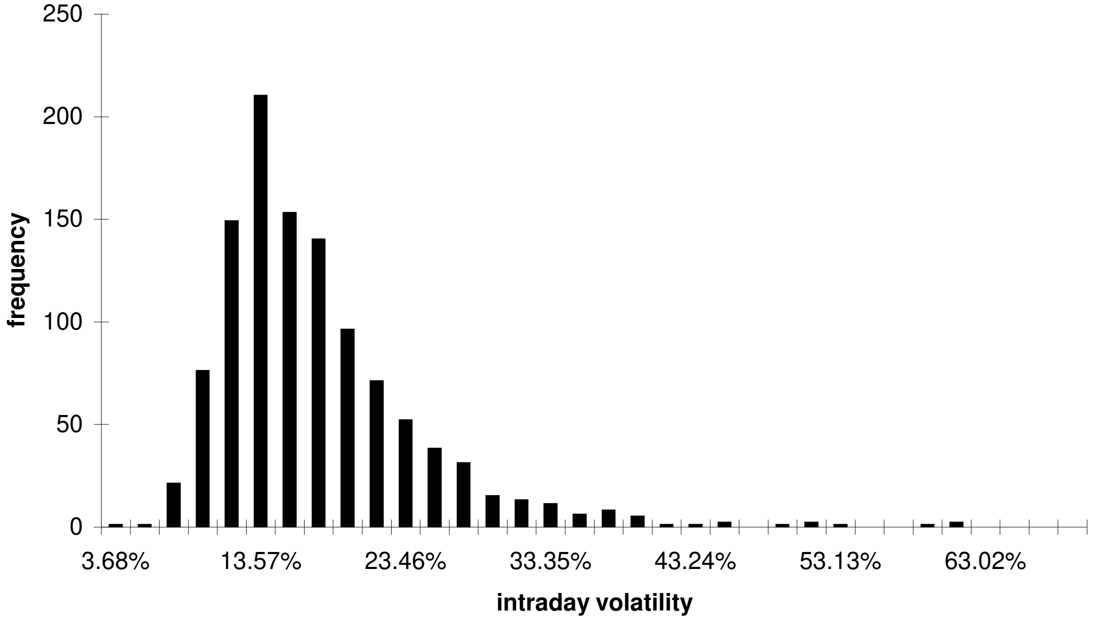

Figure 5: Histogram of estimated (realized) volatilities over the sample period July, 1999, until December, 2003for the S&P500 index. Modus is 13.6%, median is 14.7%, and average is 18.0%. The first and third quartile are 11.8% and 19.2%, respectively.

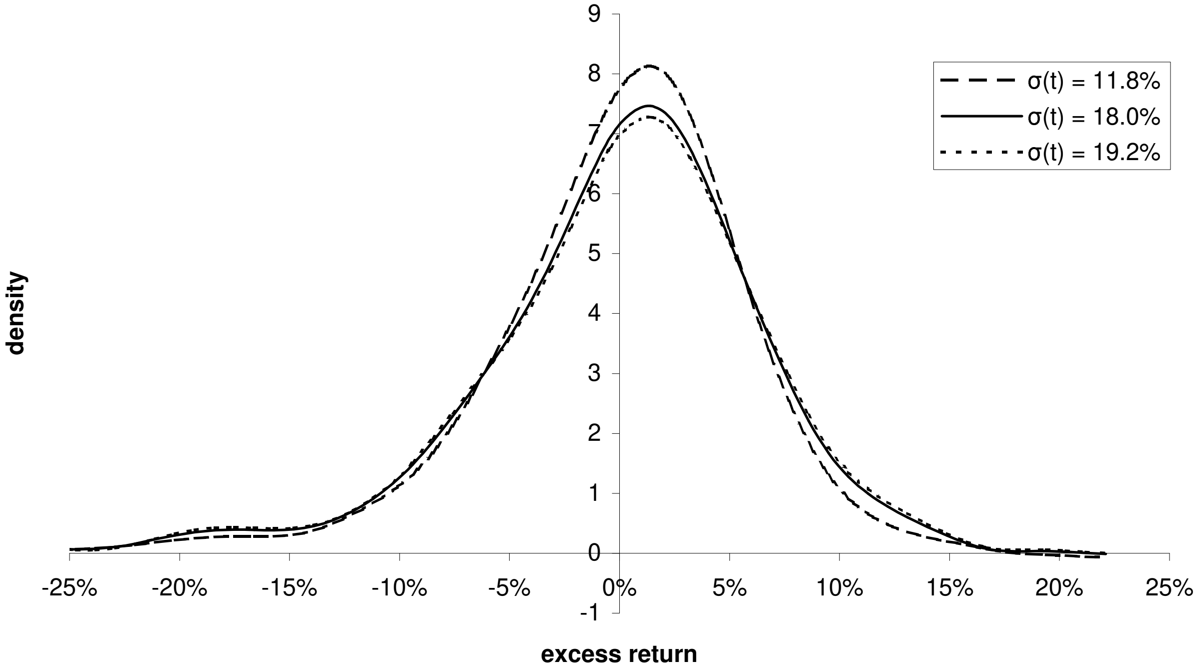

Figure 6: Estimated risk-neutral marginal excess return density over one month horizon based on S&P-500 data as described in the main text for three initial spot volatility levels.

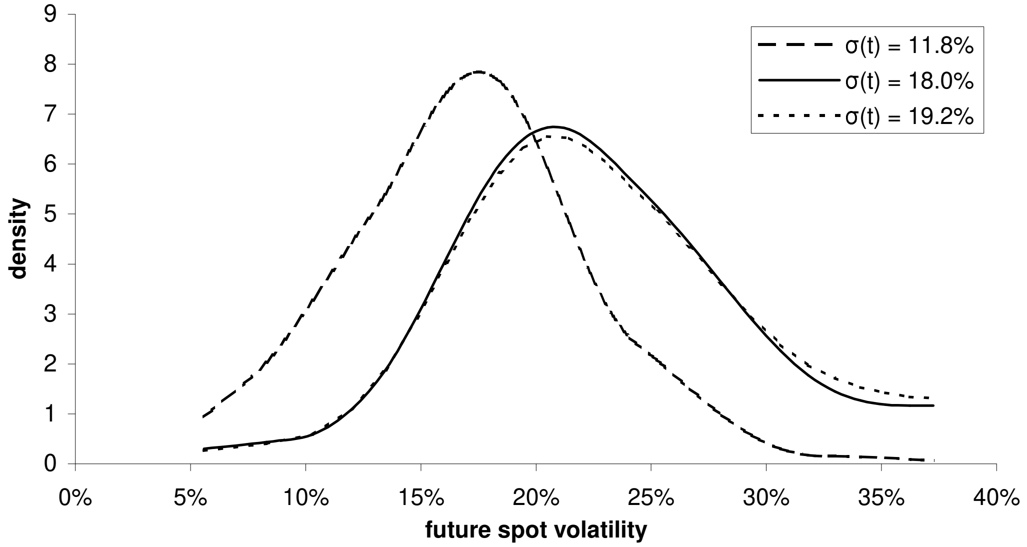

Figure 7: Estimated risk-neutral marginal volatility density over one month horizon based on S&P-500 data as described in the main text for three initial spot volatility levels.

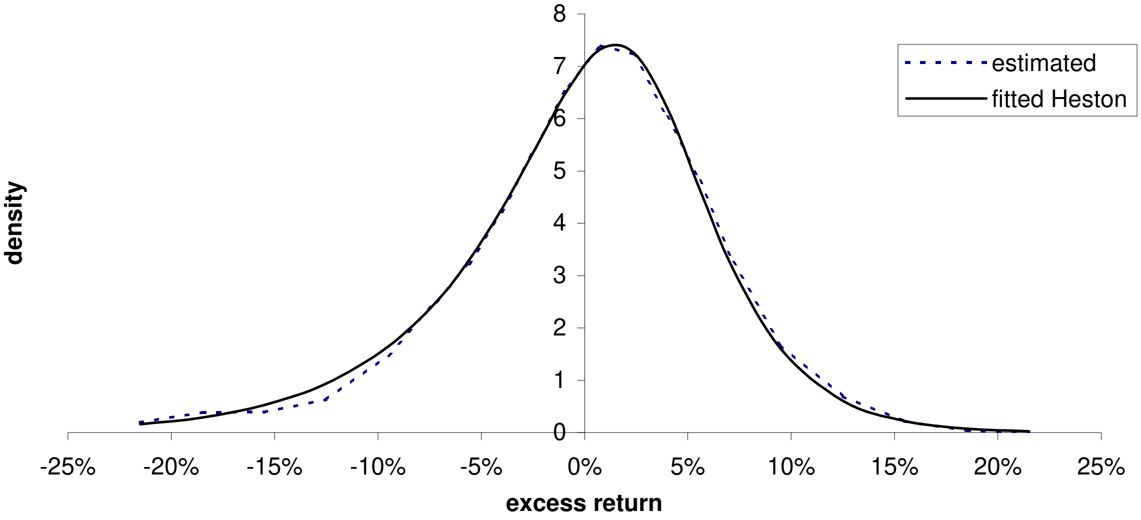

Figure 8: Risk-neutral return density from Heston (1993) models using parameters that fit the nonparametric estimate best (in quadratic mean sense), for on an (objective) average initial volatility level. See main text for details.

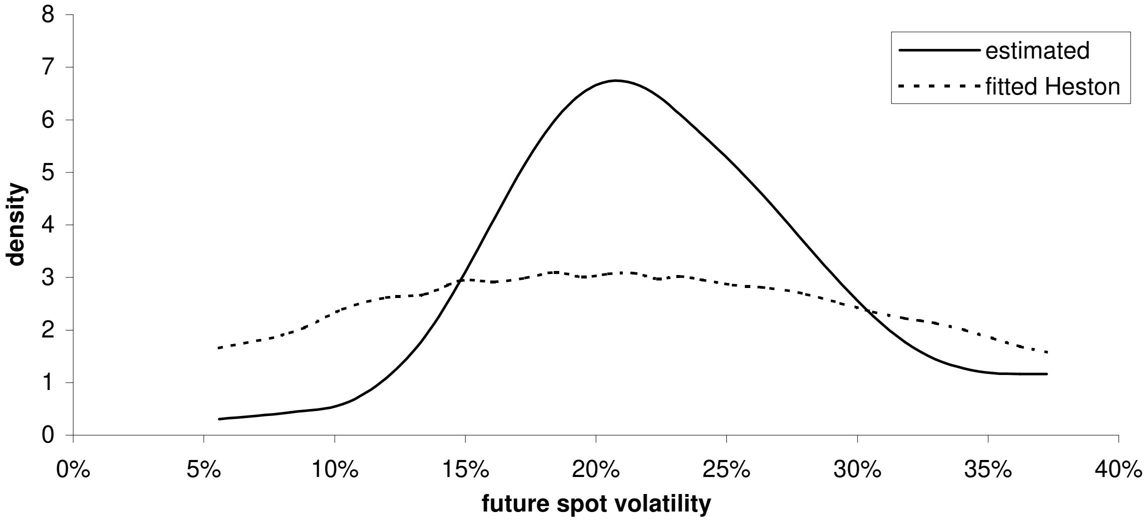

Figure 9: Risk-neutral volatility density from Heston (1993) models using the same parameters as in Figure 8.

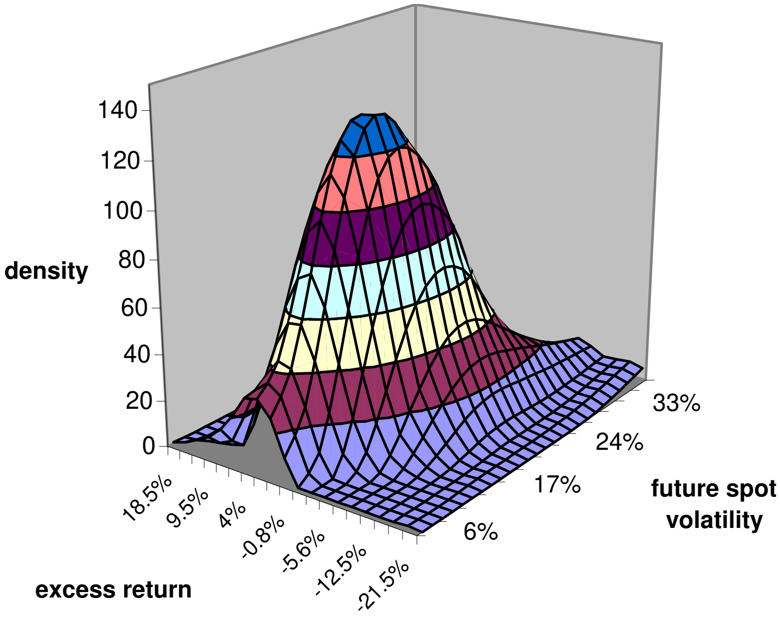

Figure 10: Estimated bivariate risk-neutral excess return and volatility density over one month horizon based on S&P500 data as described in the main text for an average initial volatility level of 18.0%

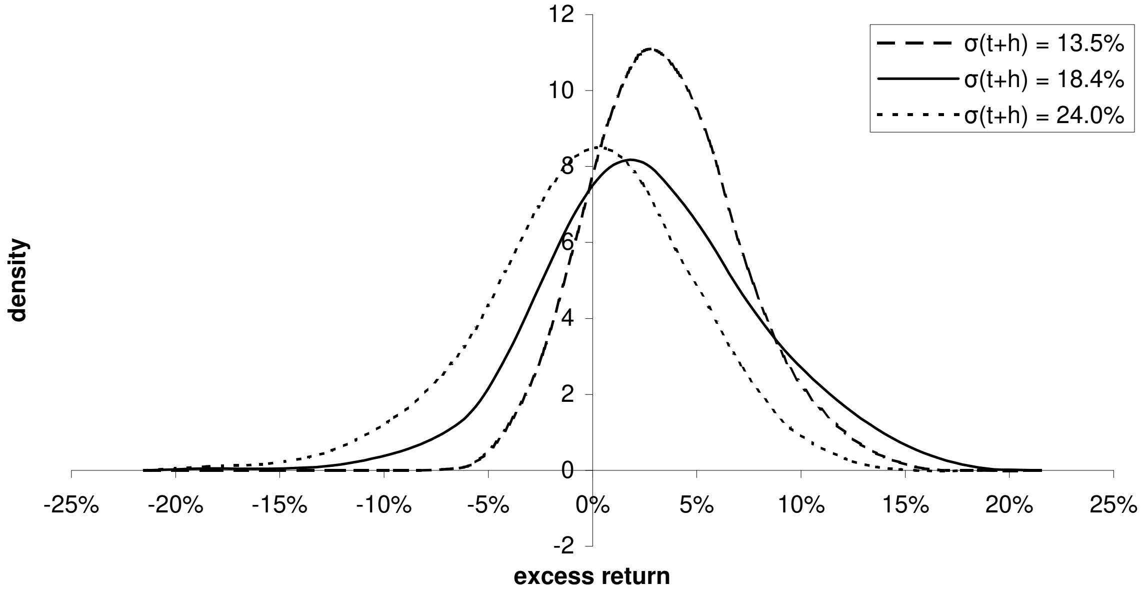

Figure 11: Risk-neutral conditional excess return distributions for an initial spot volatility level of 18.0% and several future spot volatility levels. Returns and volatilities are over a period of $h = 1$ month.