



ELSEVIER

Journal of Econometrics 94 (2000) 9–51

JOURNAL OF
Econometrics

www.elsevier.nl/locate/econbase

Nonparametric risk management and implied risk aversion

Yacine Aït-Sahalia^{a,*}, Andrew W. Lo^b

^a*Department of Economics, Princeton University, Princeton, NJ 08544-1021, USA and NBER*

^b*Sloan School of Management, MIT, Cambridge, MA 02142-1347, USA and NBER*

Abstract

Typical value-at-risk (VaR) calculations involve the probabilities of extreme dollar losses, based on the statistical distributions of market prices. Such quantities do not account for the fact that the same dollar loss can have two very different economic valuations, depending on business conditions. We propose a nonparametric VaR measure that incorporates economic valuation according to the state-price density associated with the underlying price processes. The state-price density yields VaR values that are adjusted for risk aversion, time preferences, and other variations in economic valuation. In the context of a representative agent equilibrium model, we construct an estimator of the risk-aversion coefficient that is implied by the joint observations on the cross-section of option prices and time-series of underlying asset values. © 2000 Elsevier Science S.A. All rights reserved.

JEL classification: G12; C13; C22

Keywords: Value-at risk; Risk aversion; Representative agent preferences; Nonparametric regression

1. Introduction

One of the most pressing economic issues facing corporations today is the proper management of financial risks. In response to a series of recent financial

*Corresponding author.

E-mail address: yacine@princeton.edu (Y. Aït-Sahalia)

catastrophes,¹ regulators, investment bankers, and chief executive officers have now embraced the notion of *risk management* as one of the primary fiduciary responsibilities of the corporate manager. Because financial risks often manifest themselves in subtle and nonlinear ways in corporate balance sheets and income statements, recent attention has focused on quantifying the fluctuations of market valuations in a statistical sense. These value-at-risk (VaR) measures lie at the heart of most current risk management systems and protocols. For example, JP Morgan's (1995) *RiskMetrics* system documentation describes VaR in the following way:

Value at Risk is an estimate, with a predefined confidence interval, of how much one can lose from holding a position over a set horizon. Potential horizons may be one day for typical trading activities or a month or longer for portfolio management. The methods described in our documentation use historical returns to forecast volatilities and correlations that are then used to estimate the market risk. These statistics can be applied across a set of asset classes covering products used by financial institutions, corporations, and institutional investors.

By modeling the price fluctuations of securities held in one's portfolio, an estimate and confidence interval of how much one can lose is readily derived from the basic principles of statistical inference. However, in this paper we argue that statistical notions of value-at-risk are, at best, incomplete measures of the true risks facing investors. In particular, while statistical measures do provide some information about the range of uncertainty that a portfolio exhibits, they have little to do with the economic valuation of such uncertainty. For example, a typical VaR statistic might indicate a 5% probability of a \$15M loss for a \$100M portfolio over the next month, which seems to be a substantial risk exposure at first glance. But if this 15% loss occurs only when other investments of similar characteristics suffer losses of 25% or more, such a risk may seem rather mild after all. This simplistic example suggests that a one-dollar loss is not always worth the same, and that circumstances surrounding the loss can affect its *economic* valuation, something that is completely ignored by purely statistical measures of risk.

In this paper, we propose an alternative to statistical VaR (henceforth S-VaR) that is based on *economic* valuations of value-at-risk, and which incorporates many other aspects of market risk that are central to the practice of risk management. Our alternative is based on the seminal ideas of Arrow (1964) and Debreu (1959), who first formalized the economics of uncertainty by introducing elementary securities each paying \$1 in one specific state of nature and nothing

¹ For example, the multimillion-dollar losses suffered by Gibson Greetings, Metallgesellschaft, Orange County, Proctor and Gamble, Barings Securities, etc.

in any other state. Now known as *Arrow–Debreu* securities, they are widely recognized as the fundamental building blocks of all modern financial asset-pricing theories, including the CAPM, the APT, and the Black and Scholes (1973) and Merton (1973) option-pricing models.

By construction, Arrow–Debreu prices have a probability-like interpretation – they are nonnegative and sum to unity – but since they are *market* prices determined in equilibrium by supply and demand, they contain much more information than statistical models of prices. Arrow–Debreu prices are determined by the combination of investors’ preferences, budget dynamics, information structure, and the imposition of market-clearing conditions, i.e., general equilibrium. Moreover, we shall show below that under certain special conditions, Arrow–Debreu prices reduce to the simple probabilities on which statistical VaR measures are based, hence the standard measures of value-at-risk are special cases of the Arrow–Debreu framework.

The fact that the market prices of these Arrow–Debreu securities need not be equal across states implies that a one-dollar gain need not be worth the same in every state of nature – indeed, the worth of a one-dollar gain in a given state is precisely the Arrow–Debreu price of that security. Therefore, we propose to use the prices of Arrow–Debreu securities to measure economic VaR (henceforth E-VaR).

Despite the fact that pure Arrow–Debreu securities are not yet traded on any organized exchange,² Arrow–Debreu prices can be estimated from the prices of traded financial securities using recently developed nonparametric techniques such as kernel regression, artificial neural networks, and implied binomial trees. Nonparametric techniques are particularly useful for value-at-risk calculations because departures from standard parametric assumptions, e.g., normality, can have dramatic consequences for tail probabilities.³ Using such techniques, we compare the performance of S-VaR and E-VaR measures and develop robust statistical methods to gauge the magnitudes of their differences.

Moreover, to provide an economic interpretation for the differences between S-VaR and E-VaR, we show how to combine S-VaR and E-VaR to yield a measure of the aggregate risk aversion of the economy, i.e., the risk aversion of the representative investor in a standard dynamic asset-pricing model. We propose to extract (unobservable) aggregate risk-preferences, what we call *implied risk aversion*, from (observable) market prices of traded financial securities. In particular, we are inferring the aggregate preferences that are compatible with the pair of option and index values.

² This is changing as derivatives markets become more sophisticated. For example, it is now possible to construct a limited set of Arrow–Debreu securities by forming portfolios of “digital” or “binary” options. See, also, the “supershares” security proposed by Garman (1978) and Hakansson (1976) which has been test-marketed recently by Leland, O’Brien, and Rubinstein Associates, Inc.

³ For other examples of the use of nonparametric methods in finance, see Ait-Sahalia (1996a, 1996b) and Stutzer (1996).

When applied to daily S&P 500 option prices and index levels from 1993, our nonparametric analysis uncovers substantial differences between S-VaR and E-VaR (see Fig. 2). A comparison of S-VaR and E-VaR densities shows that aggregate risk aversion is not constant across states or maturity dates, but changes in important nonlinear ways (see Fig. 4).

Of course, risk management is a complex process that is unlikely to be driven by any single risk measure, E-VaR or S-VaR. In particular, despite the fact that E-VaR provides information not contained in S-VaR, both measures belong in the arsenal of tools that risk managers can bring to bear on assessing and controlling risk. For certain purposes – regulatory reporting requirements, quick summaries of corporate exposure, or comparisons across companies and industries – S-VaR may be a simpler measure to compute and interpret. Nevertheless, E-VaR brings a new dimension into the risk management process and should be integrated into any complete risk management system.

In Section 2, we present a brief review of the theoretical underpinnings of Arrow–Debreu prices and their relation to dynamic equilibrium models of financial markets. In Section 3, we formally introduce the notion of economic value-at-risk, describe its implementation, and propose statistical inference procedures that can quantify its accuracy and relevance over statistical VaR. An explicit comparison of E-VaR with S-VaR, along with the appropriate statistical inference, is described and developed in Section 4. We construct an estimator of implied risk aversion in Section 5 and propose tests for risk neutrality and for specific preferences based on this estimator. To illustrate the empirical relevance of E-VaR, we apply our estimators to daily S&P 500 options data in Section 6. We conclude in Section 7 and collect our technical assumptions and results in the Appendix.

2. DGP, SPD, MRS, and VaR

Denote by S_t the price at time t of a security or portfolio of securities whose risk we wish to manage and let $u_{t,\tau} \equiv \log(S_{t+\tau}/S_t)$ denote its return between t and $t + \tau$. The usual statistical VaR measures are based on the probability distribution of $u_{t,\tau}$.⁴ For example, one common VaR measure is the standard deviation of returns $u_{t,\tau}$. Another is the 95% confidence interval of $u_{t,\tau}$ centered at its historical mean. More sophisticated VaR measures incorporate conditioning information and dynamics in specifying and estimating the probability distribution of $u_{t,\tau}$, i.e., they are based on conditional probabilities obtained from the *data-generating process* (DGP) of $\{S_t\}$.

⁴ See, for example, Smithson et al. (1995).

Although such VaR measures do capture important features of the uncertainty surrounding $u_{t,\tau}$, they fall short in one crucial respect: they are *statistical* evaluations of uncertainty, not *economic* valuations. In particular, one investor may be quite willing to bear a one-standard deviation drop in $u_{t,\tau}$, while another investor may be devastated by such an event. Therefore, although the dollar loss is the same for both investors, their personal valuations of such a risk can differ dramatically. More importantly, the *market* valuation of this risk – the value assigned by the interactions of many heterogeneous investors in a market setting – can differ substantially from purely statistical measures.

2.1. Dynamic equilibrium models

This distinction between the DGP and market valuations lies at the heart of dynamic equilibrium asset-pricing models in economics – beginning with Arrow (1964) and Debreu (1959) – in which the valuation of securities with uncertain payoffs is determined by the interaction and equilibration of market forces and market conditions.⁵ In such models, the specific DGP for prices is not assumed, but rather is derived from first principles as the (stochastic) sequence of prices that equates supply and demand at each point in time.

More importantly, unlike a purely statistical model of prices, e.g., geometric Brownian motion, a DGP that is derived from equilibrium prices contains an enormous amount of information about market conditions and investors' preferences that is critical for risk management. To see why, consider a standard dynamic exchange economy (Lucas, 1978; Rubinstein, 1976; Bick, 1990; He and Leland, 1993) in which securities markets are dynamically complete, there is a single consumption good, no exogenous income, and all investors seek to maximize at date t a state-independent utility function, subject to the usual budget constraints. There is one risky stock (the market portfolio, in total supply normalized to one share, and with price S_t at date t) and one riskless bond (in zero net supply) with continuously compounded rate of return r , assumed constant for simplicity. Both are available for trading at any date between t and T . Assume that the stock price process follows the stochastic differential equation

$$\frac{dS_t}{S_t} = \mu(S_t, t)dt + \sigma(S_t, t)dZ_t \quad (2.1)$$

where Z denotes a standard Brownian motion, and let $f_t(S_T)$ denote the conditional density of S_T given S_t implied by Eq. (2.1).

⁵ See Merton (1992) for a review of these and related models.

Under suitable assumptions for preferences and endowment shocks, it is well-known that market completeness allows us to introduce a representative investor with utility function U [see, for example, Constantinides (1982)]. Assuming that he can consume at date t and at the fixed future date T (allowing intermediate consumption or labor income does not affect the basic results), and that he receives one share of the stock as endowment at date t , the representative investor adjusts the dollar amount q_t invested in the stock at each intermediary date τ to solve the Merton (1971) optimization problem

$$\begin{aligned} & \max_{\{q_s \mid t \leq s \leq T\}} E[U(W_T)] \\ & \text{subject to} \quad dW_s = \{rW_s + q_s(\mu(S_s, s) - r)\} ds + q_s\sigma(S_s, s) dZ_s, \\ & \qquad \qquad \qquad W_s \geq 0, \quad t \leq s \leq T \end{aligned}$$

where W_s denotes his wealth at date s .

In equilibrium, the investor optimally invests all his wealth in the risky stock at every instant prior to T ($W_s = S_s$ for all $t \leq s \leq T$) and then consumes the terminal value of the stock at T ($C_T = W_T = S_T$). Let $J(W, S, t)$ denote the investor’s indirect utility function. The first-order condition for the investor’s problem takes the form

$$\frac{\partial J(W_s, S_s, s)}{\partial W} = e^{-r(s-t)} \frac{\partial J(W_t, S_t, t)}{\partial W} \zeta_s \tag{2.2}$$

with terminal condition at date $s = T$

$$U'(W_T) = e^{-r(T-t)} U'(W_t) \zeta_T \tag{2.3}$$

where

$$\zeta_s \equiv \exp \left\{ \int_t^s \left(\frac{\mu(S_u, u) - r}{\sigma(S_u, u)} \right) dZ_u - \frac{1}{2} \int_t^s \left(\frac{\mu(S_u, u) - r}{\sigma(S_u, u)} \right)^2 du \right\}. \tag{2.4}$$

Clearly, Eq. (2.2) implies that the right-hand-side of Eq. (2.4) must be path independent, i.e., it depends only on the values of the stock price at t and s and not on the path taken by the stock price in between these dates. The price at date t of a security with a single date- T liquidating payoff of $\psi(W_T)$ is then given by

$$E_t[\psi(W_T)M_{t,T}] \tag{2.5}$$

where $M_{t,T} \equiv U'(W_T)/U'(W_t)$ is the *stochastic discount factor* or *marginal rate of substitution* (MRS) between dates t and T . If we define the *state-price density* to be

$$f_t^*(S_T) \equiv f_t(S_T) \times E[\zeta_T | S_t, S_T] = f_t(S_T) \times \zeta_T$$

(since ζ_T is a function of (S_t, S_T)) and recall that $W_T = S_T$, then we can rewrite Eq. (2.5) as

$$\begin{aligned} E_t[\psi(W_T)M_{t,T}] &= \int_0^\infty \psi(W_T) \frac{U'(W_T)}{U'(W_t)} f_t(S_T) dW_T \\ &= e^{-r(T-t)} \int_0^\infty \psi(W_T) f_t^*(S_T) dW_T \\ &= e^{-r(T-t)} E^*[\psi(W_T)]. \end{aligned} \quad (2.6)$$

This version of the Euler equation shows that the price of any asset can be expressed as a discounted expected payoff, discounted at the riskless rate of interest. However, the expectation must be taken with respect to f^* , an MRS-weighted probability density function, not the original probability density function f of future consumption.

This density f^* is called the *state-price density* (SPD) and it is the continuous-state counterpart to the prices of Arrow–Debreu state-contingent claims that pay \$1 in a given state and nothing in all other states. Under market completeness, f^* is unique. In particular, Arrow (1964) and Debreu (1959) showed that if there are as many state-contingent claims as there are states, then the price of any security can be expressed as a weighted average of the prices of these state-contingent claims, now known as *Arrow–Debreu state prices*. In a continuous-state setting, f^* satisfies the same property – any arbitrary security can be priced as a simple expectation with respect to f^* .

This underscores the importance of f^* for risk management: the SPD aggregates all economically pertinent information regarding investors' preferences, endowments, asset price dynamics, and market clearing, whereas purely statistical descriptions of the DGP of prices do not.

Within certain restrictions, it is possible to characterize the class of DGP of prices that are compatible with an equilibrium model (see, for example Bick, 1990; Wang, 1993; He and Leland, 1993). Fixing the utility function, however, is not sufficient to identify uniquely the DGP of the price process. If parametric restrictions are imposed on the DGP of asset prices, e.g., geometric Brownian motion, the SPD may be used to infer the preferences of the representative agent

in an equilibrium model of asset prices (see, for example, Bick, 1987). Alternatively, if specific preferences are imposed, e.g., logarithmic utility, the SPD may be used to infer the DGP of asset prices. Indeed, in equilibrium, any two of the following imply the third: (1) the representative agent’s preferences; (2) asset price dynamics; and (3) the SPD.

2.2. No-arbitrage models

The practical relevance of SPDs for derivative pricing and hedging applications has also become apparent in *no-arbitrage* or *dynamically complete-markets* models in which sophisticated dynamic trading strategies involving a set of “fundamental” securities can perfectly replicate the payoffs of more complex “derivative” securities. For example, suppose that we observe a set of n_1 asset prices following Itô diffusions driven by n_2 independent Brownian motions:

$$dS_t = \mu_t dt + \sigma_t dZ_t \tag{2.7}$$

where S_t and μ_t are $(n_1 \times 1)$ -vectors, Z_t is an $(n_2 \times 1)$ -vector of independent Brownian motions, σ_t is an $(n_1 \times n_2)$ -matrix, and $n_1 \geq n_2$. Suppose that there exists a riskless asset with instantaneous rate of return r . Then path-independent derivative securities on an asset with payoff function $\psi(S_T)$ are spanned by certain dynamic trading strategies, i.e., derivatives are redundant assets hence they may be priced by arbitrage.⁶ In such applications the asset price dynamics are specified explicitly and conditions are imposed to ensure the existence of an SPD and dynamic completeness of markets (Harrison and Kreps, 1979; Duffie and Huang, 1985; Duffie, 1996).

For example, the system of asset prices S_t in Eq. (2.7) supports an SPD if and only if the system of linear equations $\sigma_t \cdot \lambda_t = \mu_t$ admits at every date a solution λ_t such that

$$\exp \left[\int_t^T \lambda_\tau \cdot \lambda_\tau d\tau / 2 \right]$$

has finite expectation, and

$$\exp \left[- \int_t^T \lambda_\tau dZ_\tau - \int_t^T \lambda_\tau \cdot \lambda_\tau d\tau / 2 \right]$$

⁶ Additional assumptions are, of course, required such as frictionless markets, unlimited riskless borrowing and lending opportunities at the same instantaneous rate $r_{t,\tau}$, a known diffusion coefficient, etc. See Merton (1973, 1992) for further discussion.

has finite variance. In the presence of an SPD, markets are complete if and only if $\text{rank}(\sigma_t) = n_2$ almost everywhere. Then the SPD can be characterized explicitly without reference to preferences. In the particular case of geometric Brownian motion, with constant volatility σ , interest rate $r_{t,\tau}$ and dividend yield $\delta_{t,\tau}$ over the interval $(t, t + \tau)$, the SPD or risk-neutral pricing density is given by the conditional distribution of the risk-neutral stochastic process with dynamics

$$dS_t^* = (r_{t,\tau} - \delta_{t,\tau})S_t^* dt + \sigma S_t^* dZ_t$$

which is a lognormal distribution with mean $((r_{t,\tau} - \delta_{t,\tau}) - \sigma^2/2)\tau$ and variance $\sigma^2\tau$.

More generally, denote by S_t the price of an underlying asset and by $f_t^*(S_T)$ the SPD of the asset price S_T at a future date T , conditioned on the current price S_t [for simplicity, we leave implicit the dependence of f_t^* on $(S_t, \tau, r_{t,\tau}, \delta_{t,\tau})$]. Consider now a European-style derivative security with a single liquidating payoff $\psi(S_T)$. To rule out arbitrage opportunities among the asset, the derivative and a risk-free cash account, the price of the derivative at t must be equal to

$$e^{-r_{t,\tau}\tau} \int_0^{+\infty} \psi(S_T) f_t^*(S_T) dS_T \tag{2.8}$$

For example, a European call option with maturity date T and strike price X has a payoff function $\psi(S_T) = \max[S_T - X, 0]$ hence its date- t price H_t is simply

$$H(S_t, X, \tau, r_{t,\tau}) = e^{-r_{t,\tau}\tau} \int_0^{+\infty} \max[S_T - X, 0] f_t^*(S_T) dS_T \tag{2.9}$$

Even the most complex path-independent derivative security can be priced and hedged according to Eq. (2.8).⁷

3. Economic VaR

The relevance of the SPD for risk management is clear: the MRS-weighted probability density function f^* provides a more economically complete measure of value-at-risk – *economic* value – than the probability density function f of the DGP. Therefore, f^* can be used in the same ways that statistical VaR measures such as standard deviation, 95% confidence intervals, tail probabilities, etc. are used. To distinguish the more traditional method of risk management from this

⁷ See also Cox and Ros (1976), Goldberger (1991), Madan and Milne (1994) and Rady (1994).

approach, we shall refer to the statistical measure of value-at-risk as ‘S-VaR’ since it is based on a purely statistical model of the DGP, and call the SPD-based measure “E-VaR” since it is based on economic considerations as well.

Now if the MRS in Eq. (2.5) were observable, implementing E-VaR measures and comparing them to S-VaR measures would be a simple matter. However, in practice obtaining f^* can be quite a challenge, especially for markets that are more complex than the pure-exchange economy described in Section 2.1. Fortunately, several accurate and computationally efficient estimators of f^* have been developed recently and we provide a brief review of these estimators in Section 3.1 and derive their asymptotic distributions in Section A.1 of Appendix A. With these estimators in hand, we show in Section 4 how to gauge the relative importance of E-VaR empirically by examining the ratio f^*/f .

3.1. SPDs and option prices

Banz and Miller (1978), Breeden and Litzenberger (1978), and Ross (1976) were among the first to suggest that Arrow–Debreu prices may be estimated or approximated from the prices of traded financial securities. In particular, building on Ross’s (1976) insight that options can be used to create pure Arrow–Debreu state-contingent securities, Banz and Miller (1978) and Breeden and Litzenberger (1978) provide an elegant method for obtaining an explicit expression for the SPD from option prices: the SPD is the second derivative (normalized to integrate to unity) of a call option pricing formula with respect to the strike price.

To see why, consider the portfolio obtained by selling two call options struck at X and buying one struck at $X - \varepsilon$ and one at $X + \varepsilon$. Consider $1/\varepsilon^2$ shares of this portfolio, often called a ‘butterfly’ spread because of the shape of its payoff function $\psi(S_T)$ which pays nothing outside the interval $[X - \varepsilon, X + \varepsilon]$. Letting ε tend to zero, the payoff function of the butterfly tends to a Dirac delta function with mass at X , i.e., in the limit the butterfly becomes an elementary Arrow–Debreu security paying \$1 if $S_T = X$ and nothing otherwise. The limit of its price as ε tends to zero should therefore be equal to $e^{r \cdot t} f^*(X)$. Now denote by $H(S_t, X, \tau)$ the market price of a call option at time t with strike price X , time-to-maturity τ , and underlying asset price S_t . Then, by construction, the price of the butterfly spread must be

$$\frac{1}{\varepsilon^2} [- 2H(S_t, X, \tau) + H(S_t, X - \varepsilon, \tau) + H(S_t, X + \varepsilon, \tau)] \tag{3.1}$$

which has, as its limit as $\varepsilon \rightarrow 0$, $\partial^2 H(S_t, X, \tau) / \partial X^2$.

3.2. The Black–Scholes SPD

For example, recall that under the hypotheses of Black and Scholes (1973) and Merton (1973),⁸ the date- t price H of a call option maturing at date $T \equiv t + \tau$, with strike price X , written on a stock with date- t price P_t and dividend yield $\delta_{t,\tau}$, is given by:⁹

$$\begin{aligned}
 H_{BS}(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}; \sigma) &= e^{-r_{t,\tau}\tau} \int_0^\infty \max[S_T - X, 0] f_{BS,t}^*(S_T) dS_T \\
 &= S_t \Phi(d_1) - X e^{-r_{t,\tau}\tau} \Phi(d_2)
 \end{aligned}
 \tag{3.2}$$

where

$$d_1 \equiv \frac{\log(S_t/X) + (r_{t,\tau} - \delta_{t,\tau} + \sigma^2/2)\tau}{\sigma\sqrt{\tau}}, \quad d_2 \equiv d_1 - \sigma\sqrt{\tau}.
 \tag{3.3}$$

In this case the corresponding SPD is a log-normal density with mean $(r_{t,\tau} - \delta_{t,\tau} - \sigma^2/2)\tau$ and variance $\sigma^2\tau$:

$$\begin{aligned}
 f_{BS,t}^*(S_T) &= e^{r_{t,\tau}\tau} \left. \frac{\partial^2 H_{BS}}{\partial X^2} \right|_{X=S_T} \\
 &= \frac{1}{S_T \sqrt{2\pi\sigma^2\tau}} \exp \left[-\frac{[\log(S_T/S_t) - (r_{t,\tau} - \delta_{t,\tau} - \sigma^2/2)\tau]^2}{2\sigma^2\tau} \right].
 \end{aligned}
 \tag{3.4}$$

⁸ Frictionless markets, continuous trading, geometric Brownian motion for S_t , specific bond price dynamics including as a special case constant interest rates, and investors' agreement on the value of σ and the distributional characteristics, not necessarily including the asset's expected rate of return [see Merton (1973, Section 6)]. The assumption of geometric Brownian motion can be relaxed to allow for mean reversion without changing the functional form of the Black–Scholes formula (as long as the diffusion coefficient σ is constant), however, equilibrium prices (hence the SPD) will not be unaffected [see He and Leland (1993, Section 3) and Lo and Wang (1995) for further discussion].

⁹ As discussed in Ait-Sahalia and Lo (1998, Section 3), the empirical analysis is greatly simplified by inferring dividend rates through observed prices on forward contracts written on the underlying asset. Let $F_{t,\tau} = S_t e^{(r_{t,\tau} - \delta_{t,\tau})\tau}$ denote the value at t of a forward contract written on the asset, with the same maturity τ as the option. At the maturity of the futures, the futures price equals the asset's spot price. Thus a European call option on the asset has the same value as a European call option on the futures contract with the same maturity. As a result, we will often rewrite the Black–Scholes formula as $H_{BS}(F_{t,\tau}, X, \tau, r_{t,\tau}; \sigma) = e^{-r_{t,\tau}\tau} (F_{t,\tau} \Phi(d_1) - X \Phi(d_2))$, with $d_1 \equiv (\log(F_{t,\tau}/X) + (\sigma^2/2)\tau)/(\sigma\sqrt{\tau})$ and $d_2 \equiv d_1 - \sigma\sqrt{\tau}$.

This expression shows that the SPD can depend on many quantities in general, and is distinct from but related to the PDF of the terminal stock price S_T . More generally, while sufficiently strong assumptions on the underlying asset price dynamics can often characterize the SPD uniquely, in most cases the SPD cannot be computed in closed form and numerically intensive methods must be used to calculate it. It is clear from Eq. (3.4) that the SPD is inextricably linked to the parametric assumptions underlying the Black–Scholes option pricing model. If those parametric assumptions do not hold, e.g., if the dynamics of $\{S_t\}$ contain Poisson jumps, then Eq. (3.4) will yield incorrect prices, prices that are inconsistent with the dynamic equilibrium model or the hypothesized stochastic process driving $\{S_t\}$. Given the general lack of success in fitting highly parametric models to financial data (see, for example, Campbell et al., 1997, Chapters 2 and 12), combined with the availability of the data and the large effects of differences in specification, it is quite natural to focus on nonparametric methods for estimating SPDs.

3.3. Nonparametric SPD estimators

Ait-Sahalia and Lo (1998) propose to estimate the SPD nonparametrically by exploiting Breeden and Litzenberger's (1978) insight that $f_t^*(S_T) = \exp(r_{t,\tau})\partial^2 H(\cdot)/\partial X^2$. They suggest using market prices to estimate an option-pricing formula $\hat{H}(\cdot)$ nonparametrically, which can then be differentiated twice with respect to X to obtain $\partial^2 \hat{H}(\cdot)/\partial X^2$. They use kernel regression to construct $\hat{H}(\cdot)$.¹⁰ Assuming that the option-pricing formula H to be estimated is an arbitrary nonlinear function of a vector of option characteristics or “explanatory” variables, $Y \equiv [F_{t,\tau} \ X \ \tau \ r_{t,\tau}]'$ where $F_{t,\tau}$ is the forward price of the asset.

In practice, they propose to reduce the dimension of the kernel regression by using a *semiparametric* approach. Suppose that the call pricing function is given by the parametric Black–Scholes formula (3.2) except that the implied volatility parameter for that option is a nonparametric function $\sigma(X/F_{t,\tau}, \tau)$:

$$H(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}) = H_{BS}(F_{t,\tau}, X, \tau, r_{t,\tau}; \sigma(X/F_{t,\tau}, \tau)). \quad (3.5)$$

We assume that the function H defined by Eq. (3.5) satisfies all the required conditions to be a ‘rational’ option-pricing formula in the sense of Merton

¹⁰ See Härdle (1990) and Wand and Jones (1995) for a more detailed discussion of nonparametric regression. There are other alternatives to Eq. (3.5) that can be used to obtain option-pricing formulas: see Derman and Kani (1994), Dupire (1994), Hutchinson et al. (1994), and Rubinstein (1994). For an extension to American options and the nonparametric estimation of the early exercise boundary, see Broadie et al. (1996). For tests of the volatility models, see Dumas et al. (1995).

(1973, 1992).¹¹ In this semiparametric model, we only need to estimate non-parametrically the regression of σ on a subset \tilde{Y} of the vector of explanatory variables Y . The rest of the call pricing function $H(\cdot)$ is parametric, thereby considerably reducing the sample size n required to achieve the same degree of accuracy as the full nonparametric estimator. We partition the vector of explanatory variables $Y \equiv [\tilde{Y}' F_{t,\tau} r_{t,\tau}]'$ where \tilde{Y} contains \tilde{d} nonparametric regressors. As a result, the effective number of nonparametric regressors d is given by \tilde{d} .

In our empirical application, we will consider $\tilde{Y} \equiv [X/F_{t,\tau} \tau]'$ ($\tilde{d} = 2$) and form the Nadaraya–Watson kernel estimator of $E[\sigma | X/F_{t,\tau}, \tau]$ as:

$$\hat{\sigma}(X/F_{t,\tau}, \tau) = \frac{\sum_{i=1}^n k_{X/F} \left(\frac{X/F_{t,\tau} - X_i/F_{t_i,\tau_i}}{h_{X/F}} \right) k_{\tau} \left(\frac{\tau - \tau_i}{h_{\tau}} \right) \sigma_i}{\sum_{i=1}^n k_{X/F} \left(\frac{X/F_{t,\tau} - X_i/F_{t_i,\tau_i}}{h_{X/F}} \right) k_{\tau} \left(\frac{\tau - \tau_i}{h_{\tau}} \right)} \tag{3.6}$$

where σ_i is the volatility implied by the option price H_i , and the univariate kernel functions $k_{X/F}$ and k_{τ} and the bandwidth parameters $h_{X/F}$ and h_{τ} are chosen to optimize the asymptotic properties of the second derivative of $\hat{H}(\cdot)$, i.e., of the SPD estimator. We then estimate the call pricing function as:

$$\hat{H}(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}) = H_{BS}(F_{t,\tau}, X, \tau, r_{t,\tau}, \delta_{t,\tau}; \hat{\sigma}(X/F_{t,\tau}, \tau)). \tag{3.7}$$

The SPD estimator follows by taking the second partial derivatives of $\hat{H}(\cdot)$ with respect to X :

$$\hat{f}_t^*(S_T) = e^{r_{t,\tau} T} \left[\frac{\partial^2 \hat{H}(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau})}{\partial X^2} \right]_{X=S_T}. \tag{3.8}$$

We give the asymptotic distribution of this estimator in Section A.1 of Appendix A.

3.4. Other estimators of the SPD

Several other estimators of the SPD have been proposed in the recent literature [see Ait-Sahalia and Lo (1998) for a more detailed discussion and an empirical comparison]. Hutchinson et al. (1994) employ several nonparametric

¹¹ In particular, see Merton 1992, Section 8.2). These conditions imply that $\sigma(X/F, \tau)$ cannot be an arbitrary function but must yield an $H_{BS}(F_{t,\tau}, X, \tau, r_{t,\tau}; \sigma(X/F_{t,\tau}, \tau))$ that satisfies all the conditions of a rational option-pricing formula.

techniques to estimate option-pricing models that they describe collectively as *learning networks* – artificial neural networks, radial basis functions, and projection pursuit – and find that all these techniques can recover option-pricing models such as the Black–Scholes model. Taking the second derivative of their option-pricing estimators with respect to the strike price yields an estimator of the SPD.

Another estimator is Rubinstein's (1994) *implied binomial tree*, in which the risk-neutral probabilities $\{\pi_i^*\}$ associated with the binomial terminal stock price S_T are estimated by minimizing the sum of squared deviations between $\{\pi_i^*\}$ and a set of *prior* risk-neutral probabilities $\{\tilde{\pi}_i^*\}$, subject to the restrictions that $\{\pi_i^*\}$ correctly price an existing set of options and the underlying stock, in the sense that the optimal risk-neutral probabilities yield prices that lie within the bid-ask spreads of the options and the stock [see also Jackwerth and Rubinstein (1996) for smoothness criteria].

This approach is similar in spirit to Jarrow and Rudd's (1982) and Longstaff's (1995) method of fitting risk-neutral density functions using a four-parameter Edgeworth expansion. However, Rubinstein (1994) points out several important limitations of Longstaff's method when extended to a binomial model, including the possibility of negative probabilities. Derman and Kani (1994) and Shimko (1993) have proposed related estimators of the SPD.

There are several important differences between kernel estimators and implied binomial trees. Implied binomial trees require a prior $\{\tilde{\pi}_i^*\}$ for the risk-neutral probabilities; kernel estimators do not. Implied binomial trees are typically estimated for *each* cross-section of options; kernel estimators aggregate options prices over time to get a single SPD. This implies that implied binomial tree is completely consistent with all option prices at each date, but is not necessarily consistent across time. In contrast, the kernel SPD estimator is consistent across time, but there may be some dates for which the SPD estimator fits the cross-section of option prices poorly and other dates for which the SPD estimator performs very well.

Whether or not consistency over time is a useful property depends on how well the economic variables used in constructing the kernel SPD can account for time variation in risk-neutral probabilities. In addition, the kernel SPDs take advantage of the data temporally surrounding a given date. Tomorrow's and yesterday's SPDs contain information about today's SPD – this information is ignored by the implied binomial trees but not by kernel-estimated SPDs.

Finally, and perhaps most importantly, statistical inference is virtually impossible with learning-network estimators and implied binomial trees, because of the recursive nature of the former approach (White, 1992), and the non-stationarities inherent in the latter approach (recall that implied binomial trees are estimated for each cross-section of option prices). In contrast, the statistical inference of kernel estimators is well developed and computationally quite tractable.

4. Statistical VaR

Having obtained an estimator \hat{f}^* of the SPD, we can now gauge its importance for risk management by studying the behavior of the ratio of \hat{f}^* to \hat{f} , where \hat{f} is an estimator of the conditional density of the DGP, i.e., S-VaR. If the ratio $\hat{\zeta} \equiv \hat{f}^*/\hat{f}$ exhibits substantial variation over its domain, this indicates that E-VaR measures contain important economic information that are not captured by their S-VaR counterparts. Of course, because of estimation error, $\hat{\zeta}$ will never be constant in any given dataset even if ζ is. Therefore some measure of the statistical fluctuations inherent in $\hat{\zeta}$ is required, and we now propose estimators for f and f^* and describe how to combine them to estimate and conduct statistical inference on $\hat{\zeta}$.

4.1. Parametric S-VaR

There are several methods of estimating the S-VaR or statistical distribution of the future price S_T conditional on the current price S_t . Currently, the most common approach is to assume that this distribution belongs to some parametric family, e.g., lognormal, which is characterized by a small number of parameters. Estimating the distribution is then reduced to estimating these parameters. This underlies virtually all of the S-VaR measures currently used in practice. However, there is considerable empirical evidence to suggest that the most popular parametric models are generally inconsistent with the behavior of recent financial data. Therefore, a nonparametric approach to S-VaR may be preferred.

4.2. Nonparametric S-VaR

In keeping with the spirit of our nonparametric approach, we propose to estimate the S-VaR of S_T without making any parametric assumptions. In particular, we collect the time series of the index values, calculate the τ -period continuously compounded returns, $u_\tau \equiv \log(S_T/S_t)$, and construct a kernel estimator of the density function $g(\cdot)$ of these returns:

$$\hat{g}(u_\tau) \equiv \frac{1}{NH_{ui=1}} \sum^N k_u\left(\frac{u_\tau - u_{t_i,\tau}}{H_u}\right). \tag{4.1}$$

From the density of the continuously compounded returns we can then calculate

$$\Pr(S_T \leq S) = \Pr(S_t e^{u_\tau} \leq S) = \Pr(u_\tau \leq \log(S/S_t)) = \int_{-\infty}^{\log(S/S_t)} g(u_\tau) du_\tau \tag{4.2}$$

and recover the price density $f_t(\cdot)$ corresponding to return density as

$$f_t(S) = \frac{\partial}{\partial S} \Pr(S_T \leq S) = \frac{g(\log(S/S_t))}{S}. \tag{4.3}$$

Our estimator of S-VaR is

$$\hat{f}_t(S_T) = \frac{\hat{g}(\log(S_T/S_t))}{S_T} \tag{4.4}$$

which can be computed directly from the estimator (4.1) of the density function $g(\cdot)$.

At first glance, deriving the asymptotic distribution of the ratio $\hat{\zeta} \equiv \hat{f}^*/\hat{f}$ seems to involve a number of complex steps due to the different sample sizes used to construct the numerator (based on the entire panel data of option prices) and the denominator (based only on the time series of index returns), as well as the cross-correlation between \hat{f}^* and \hat{f} . Fortunately, this problem has a crucial characteristic that simplifies things considerably: the rates of convergence of the two estimators are different. Specifically, the S-VaR estimator \hat{f} converges faster than the E-VaR estimator \hat{f}^* .

There are indeed two opposing effects that only partially offset each other: on the one hand, \hat{f}^* is obtained by twice differentiating a nonparametric regression function, while \hat{f} is a direct nonparametric density estimator, thus \hat{f}^* converges slower, *ceteris paribus*; on the other hand, \hat{f}^* is estimated using a larger sample size (based on the entire panel data of option prices, $n = NJ$) while \hat{f} is estimated using only N observations (based only on the time series of index returns), thus \hat{f}^* converges faster, *ceteris paribus*. Details are in Section A.2 (of Appendix A) for the S-VaR estimator and in Section A.3 (of Appendix A) for the ratio of f^* to f .

5. Implied risk aversion

It is apparent from (2.3) that the ratio of f^* to f is proportional to the MRS of the representative agent:

$$\zeta_T \equiv \frac{f_t^*(S_T)}{f_t(S_T)} = e^{r(T-t)} \frac{U'_T(C_T)}{U'_t(C_t)} = e^{r(T-t)} M_{t,T}. \tag{5.1}$$

Therefore, their ratio $\hat{\zeta} = \hat{f}^*/\hat{f}$ is an estimator – up to a scale factor – of the MRS itself. If $\hat{\zeta}(\cdot)$ is a nearly constant function over its domain (recall that both $\hat{f}^*(\cdot)$ and $\hat{f}(\cdot)$ are estimated *functions* with several arguments), this suggests that the representative agent of Section 2.1 is approximately risk neutral and that S-VaR

and E-VaR measures will be close.¹² This observation will be the basis for a test of risk neutrality described below.

5.1. Interpretations and qualifications

In interpreting the quantity \hat{f}^*/\hat{f} as a scaled estimator of the MRS, we are implicitly assuming that an equilibrium asset-pricing model (such as the representative-agent model of Section 2.1) holds. Of course, the particular model of economic equilibrium used most often in these studies – the representative-agent model of Lucas (1978) – has been criticized on a number of theoretical and empirical grounds, e.g., Kirman (1992), Kocherlakota (1996) and Rogoff (1996).¹³ Moreover, many of these models focus on aggregate consumption, whereas risk-management issues involve other quantities as well, e.g., the S&P 500 index.¹⁴ Nevertheless, we can still conduct the simple thought-experiment of a simple economy in which market-wide financial aggregates like the S&P 500 proxy for aggregate consumption, as for instance in Brown and Gibbons (1985). Recall that in the economy of Section 2.1, in equilibrium the investor optimally invests all his wealth in the single risky stock at every instant prior to the terminal date and then consumes the terminal value of the stock at T . While this model of equilibrium is restrictive, it has the advantage that \hat{f}^*/\hat{f} measures the

¹² Of course, if ζ is constant, then its only possible value is one, since both f^* and f integrate to one.

¹³ See Constantinides (1982) for a discussion of the generality that the representative-agent model affords.

¹⁴ This distinction is an important one in our application for another reason: our use of the S&P 500 as a proxy for aggregate consumption does not account for inflation. In particular, the standard economic paradigm in which a utility function $U(C)$ is specified over consumption bundles C , consumption is generally taken to mean *real* consumption, i.e., quantities of goods and services. In contrast, the simple framework of Section 2.1 assumes that the representative agent's preferences is defined for end-of-lifetime wealth W_T , and by choosing the S&P 500 index as a proxy for W_T , we have implicitly defined $U(C)$ over *nominal* (i.e., dollar-denominated) consumption. In principle, it is a simple matter to convert nominal quantities to real quantities – simply deflate nominal values by a price deflator – but since we are extracting the MRS from options on the S&P 500, we cannot perform this conversion easily (we would require market prices for options on the real value of the S&P 500 or other real consumption indexes, and these do not yet exist).

However, our nominal approach is consistent with the notion of pure Arrow–Debreu securities, securities that pay \$1 in a single state of nature and nothing in all other states. It is the relationship between the prices of Arrow–Debreu securities and the marginal utilities of real consumption that is confounded by changes in purchasing power.

Since our empirical application involves rather short time horizons – less than one year – during a period of relatively low inflation (January to December, 1993), the distinction between real and nominal quantities may not be as serious an issue as for longer-horizon applications. Also, under certain conditions – for example, homothetic preferences and inflation that is statistically independent of the real economy – inflation will not affect the interpretation of our E-VaR estimators. And our test for risk neutrality in Section 5.3 is unaffected by our use of nominal quantities (although our test for specific preferences in Section 5.4 is affected).

risk preferences implicit in traded financial securities, and reflects the market's aggregation of such preferences. However, we recognize that this need not be the case in more general models. For instance, if the MRS contained two factors, the S&P 500 and one factor that is orthogonal to it, then the SPD \hat{f}^* that we estimated for the S&P 500 could be misleading as an instrument for pricing, say, a fixed-income portfolio.

Moreover, our approach assumes that markets are complete: starting from a finite set of call options with different strike prices, we are able to extract enough information to price other derivative securities, and this can only be done generically if markets are complete. The incomplete markets case may be a more plausible alternative, but in that case, the precise nature of incompleteness – borrowing constraints, uninsurable shocks to labor income, other market frictions – must be modeled explicitly before any meaningful implications for risk management can be developed.¹⁵ Since we do not have strong priors about the potential sources of market incompleteness, we take the simpler but more restrictive approach outlined in Section 2.1.

Although the differences between f and f^* can be quite large for certain types of preferences, implying important differences between S-VaR and E-VaR measures, there is one set of preferences for which there is no difference: risk neutrality. If the representative agent is risk neutral, his utility function must be linear in W_T hence his MRS is unity and $f^*/f = 1$. Therefore, S-VaR measures are all special cases of the more general E-VaR measures in which the MRS is allowed to vary, and the relevance of the latter over the former hinges on how risk averse the representative agent is in practice. In other words, the ratio f^*/f carries information that is relevant for risk management, but also more generally for the literature on tests of dynamic asset-pricing models, MRS bounds, the equity-premium puzzle, variance-bounds tests, etc.¹⁶

5.2. A characterization of implied risk aversion

Recall that in equilibrium any two of the following imply the third: (1) the representative agent's preferences U ; (2) the asset price dynamics, or equivalently

¹⁵ A certain degree of generality may be obtained by considering a consumption-based asset-pricing model such as those in Breeden (1979), Hansen and Singleton (1983), and Grossman and Shiller (1982). However, this approach requires that consumption be observable, a nontrivial requirement in light of the time-averaged and error-prone nature of aggregate consumption data. Alternatively, Cochrane and Saa-Requejo (1996) develop bounds for the possible values of the missing securities in an incomplete market by imposing a priori limits on the possible Sharpe ratios attainable by forming portfolios with these securities.

¹⁶ See, for example, Hansen and Singleton (1983), Hansen and Jagannathan (1991), Mehra and Prescott (1985), and Shiller (1981). It should be emphasized that our approach does not make use of aggregate consumption data, but uses nominal data instead (S&P 500 option prices). See footnote 12 for further discussion.

the conditional density function f ; and (3) the SPD f^* . In this section, we implement this relationship by inferring from the market prices of options and index values the information on f^* and f , respectively, that is needed to characterize the representative agent’s preferences. Recall that the first-order condition (2.3) for the investor’s optimization problem is

$$\zeta_t(S_T) \equiv \frac{f_t^*(S_T)}{f_t(S_T)} = \lambda e^{-r_{i,t}\tau} \frac{U'(S_T)}{U'(S_t)}$$

where λ is a constant independent of the index level. Rather than extract from the ratio f^*/f information about the MRS (which would require knowledge of both the initial level of marginal utility $U'(S_t)$ and the constant λ), we can directly infer the Arrow–Pratt measure of relative risk aversion $\rho_t(S_T)$ by observing that

$$\zeta'_t(S_T) = \lambda e^{-r_{i,t}\tau} \frac{U''(S_T)}{U'(S_t)} \Rightarrow \rho_t(S_T) \equiv -\frac{S_T U''(S_T)}{U'(S_T)} = -\frac{S_T \zeta'_t(S_T)}{\zeta_t(S_T)}. \tag{5.2}$$

This suggests a very natural estimation strategy for the measure of risk aversion $\rho(\cdot)$: we estimate the first derivatives of $f^*(\cdot)$ and $f(\cdot)$, and then calculate

$$\hat{\rho}_t(S_T) \equiv -\frac{S_T \hat{\zeta}'_t(S_T)}{\hat{\zeta}_t(S_T)} = \frac{S_T \hat{f}'_t(S_T)}{\hat{f}_t(S_T)} - \frac{S_T \hat{f}^*{}'(S_T)}{\hat{f}^*(S_T)}. \tag{5.3}$$

Before we go any further, it is useful to examine what answer we would obtain if we were to apply our estimation strategy based on Eq. (5.2) to the case where the data were generated by the Black–Scholes model. It is well-known that Constant Relative Risk Aversion (CRRA) preferences sustain the Black–Scholes model in equilibrium (see, for instance, Rubinstein, 1976; Breeden and Litzenberger, 1978; Brennan, 1979; Bick, 1987; He and Leland, 1993). We now show that our characterization (5.2) of the implied risk aversion through the ratio of E-VaR and S-VaR reproduces this result.¹⁷ In the Black–Scholes case, $f_t^*(S_T)$ is given by Eq. (3.4), while $f_t(S_T)$ is given by the same Eq. (3.4) with $(r - \delta)$ replaced by the actual drift μ of the DGP:

$$dS_t = \mu S_t dt + \sigma S_t dZ_t. \tag{5.4}$$

¹⁷ Recall that in the economy we have hypothesized, the equilibrium is such that the investor optimally invests all his wealth in the single risky stock at every instant prior to the terminal date and then consumes the terminal value of the stock. Should that assumption be relaxed, the Black–Scholes model could be supported by preferences other than CRRA.

Indeed, in the Black–Scholes model, S_T follows a geometric Brownian motion under both the actual and risk-neutral probability measures; the only difference lies in the expected return of the asset. It follows that

$$\begin{aligned} \zeta_{BS,t}(S_T) &= \left(\frac{S_T}{S_t}\right)^{-(\mu-r+\delta)/\sigma^2} \exp\left[\frac{(\mu-r+\delta)(\mu+r-\delta-\sigma^2)\tau}{2\sigma^2}\right] \\ &\equiv c\left(\frac{S_T}{S_t}\right)^{-a} \end{aligned} \tag{5.5}$$

where c denotes a constant and $a \equiv (\mu - r + \delta)/\sigma^2$. Depending upon the value of $a \geq 0$, we obtain (up to an irrelevant constant)

$$\begin{aligned} a = 1: \quad &U_{BS}(S_T) = \log(S_T) \\ a \neq 1: \quad &U_{BS}(S_T) = S_T^{1-a}/(1-a) \end{aligned} \tag{5.6}$$

where $a = 0$ corresponds to a risk-neutral representative agent (in that case, the expected rate of return on the asset is $\mu = r - \delta$), and $a = 1$ corresponds to logarithmic utility. More generally, the class of representative-agent utility functions which are implied by the Black–Scholes model all belong to the class,

$$\rho_{BS,t}(S_T) = a \tag{5.7}$$

where a is the constant *coefficient* of relative risk aversion. This can be contrasted to our estimation strategy, based on Eq. (5.3), where the relative risk aversion ρ is obtained as a *function* of the level of aggregate wealth S_T , with the function identified by forcing the joint observations on option prices (E-VaR) and index returns (S-VaR) to be compatible with one another. Eq. (5.7) corresponds to the special case where the function reduces to a constant.

Eq. (5.7) can be linked quite naturally to the known necessary and sufficient condition for $\{S_t\}$ to be an equilibrium price process in a Black–Scholes economy (Bick, 1990; Wang, 1993; and He and Leland, 1993, Proposition 1): when σ is constant, the asset’s expected return $\mu(S, t)$ in $dS_t = \mu(S_t, t)S_t dt + \sigma S_t dZ_t$ must satisfy the partial differential equation

$$\frac{\partial \mu}{\partial t} + S_t \mu \frac{\partial \mu}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 \mu}{\partial S^2} = 0$$

with the boundary condition

$$\mu(S, T) = r - \delta - \sigma^2 \frac{S U''(S)}{U'(S)}.$$

In the Black–Scholes model, where μ is independent of the level of the stock price, this equation reduces to our characterization (5.7) of the representative agent’s preferences jointly implied by f_{BS}^* and f_{BS} . Note that for each specification of the utility function, a drift function can be found to construct an equilibrium price process with constant volatility (subject to certain regularity conditions). In other words, the equilibrium price dynamics of the underlying asset are not fully identified from the utility function alone. Knowledge of the SPD is required.

Of course, we wish to avoid relying on assumptions such as constancy of σ and/or μ and instead infer the representative agent’s preferences directly from the market prices of options and index values in a nonparametric fashion – this is precisely what Eq. (5.3) allows us to do.

Using implied binomial trees, Jackwerth (1997) also proposes to estimate aggregate risk aversion levels that are consistent with option prices. Indeed, any SPD estimator can be used to extract information about risk aversion. The relative strengths and weaknesses of these approaches are determined by the relative strengths and weaknesses of the corresponding SPD estimators on which they are based. See Section 3.4 and Ait-Sahalia and Lo (1998) for further details. In Section A.4 (of Appendix A), we calculate the distribution of our estimator of the relative risk aversion function $\hat{\rho}$. Its distribution will be used to construct pointwise confidence intervals for the estimator.

5.3. A test of risk neutrality

In our context, risk neutrality of the representative agent can be characterized equivalently as $f^* = f$ or as $\rho = 0$. A test based on comparing either the $f^*(\cdot)$ to $f(\cdot)$ only involves the computation of the E-VaR and S-VaR, whereas comparing $\rho(\cdot)$ to 0 also involves calculating their first derivatives, and it can be shown that this results in a loss of power. We therefore propose to test risk neutrality in the form of the hypothesis

$$H_0: \Pr(f_i^*(S_T) = f_i(S_T)) = 1 \quad \text{vs.} \quad H_A: \Pr(f_i^*(S_T) = f_i(S_T)) < 1. \tag{5.8}$$

A natural test statistic is

$$R(f^*, f) \equiv E[(f_i^*(S_T) - f_i(S_T))^2 \omega(\tilde{Y})] \tag{5.9}$$

where $\omega(\tilde{Y})$ is a weighting function. An estimator for $R(f^*, f)$ is the sample analog of the righthand side of Eq. (5.9)

$$\hat{R}(\hat{f}^*, \hat{f}) \equiv \frac{1}{n} \sum_{i=1}^n (\hat{f}_i^* - \hat{f}_i)^2 \omega_i \tag{5.10}$$

where $\hat{f}_i^* \equiv \hat{f}_i^*(S_T)$, $\hat{f}_i \equiv \hat{f}_i(S_T)$ and ω_i is a trimming index. Any other evaluation of the integral on the righthand side of Eq. (5.9) can be used.

If the representative agent's preferences implied by the joint data on option prices (f^*) and index dynamics (f) are indeed risk neutral, then the two density functions should be close to one another, and R close to zero. If we can take into account the sampling variation due to data noise, then we will be able to use Eq. (5.10) as the basis for a test statistic of Eq. (5.8). For that purpose, we derive in Section A.5 (of Appendix Appendix A) the distribution of our test statistic.

5.4. A test of specific preferences

More generally, we can test whether the index and option data, as summarized by our estimated E-VaR and S-VaR, are compatible with a specific set of preferences for the representative agent. For example, the null hypothesis might be that the representative agent has CRRA preferences, $\rho_t(S_T) = a/S_T$. Let $\rho_t(S_T; \theta)$ be the Arrow–Pratt measure of risk aversion specified by the null hypothesis, where θ is a parameter vector in a compact subset of R^b . Note that we cannot test the specific preferences hypothesis at the level of f^*/f since that would involve an unknown factor (the initial value of the agent's marginal utility). However, the hypothesis is unambiguous if taken in the form of the risk aversion measure. We can test

$$H_0: \Pr(\rho_t(S_T) = \rho_t(S_T; \theta)) = 1 \text{ vs. } H_A: \Pr(\rho_t(S_T) = \rho_t(S_T; \theta)) < 1$$

using the following test statistic:

$$P(\rho_t, \rho_t(\cdot; \theta)) \equiv E[(\rho_t(S_T) - \rho_t(S_T; \theta))^2 \omega(\mathbf{Y})]$$

where $\omega(\mathbf{Y})$ is a weighting function. The distribution of this statistic has the same form as in Section 5.3, except that its convergence rate is slower due to the fact that $\hat{\rho}_t(\cdot)$ converges even slower than $\hat{f}_i^*(\cdot)$. Just as the estimation of $\hat{f}_i^*(\cdot)$ did not affect the distribution of the statistic in the previous case, the estimation of the unknown parameter vector θ will not affect the distribution of this statistic. Any consistent estimator $\hat{\theta}$ of θ may then be used for the purpose of calculating

$$\hat{P}(\hat{\rho}_t, \hat{\rho}_t(\cdot; \hat{\theta})) \equiv \frac{1}{n} \sum_{i=1}^n (\hat{\rho}_t, \rho_{i, \hat{\theta}})^2 \omega_i$$

6. An empirical example

To gauge the empirical relevance of E-VaR, we compare it to S-VaR in the case of S&P 500 index options using data obtained from the CBOE for the

sample period from 4 January 1993 to 31 December 1993. In particular, we estimate $\zeta \equiv f^*/f$ by taking the ratio of nonparametric estimators of f^* and f , where f^* is estimated according to Ait-Sahalia and Lo (1998) (see Section 3) and f is estimated by standard kernel density estimation techniques. We also estimate the coefficient of risk aversion ρ as described in Section 5.2, and test the null hypothesis of risk neutrality as proposed in Section 5.3.

6.1. S&P 500 index options

We use the same dataset of option prices and characteristics as Ait-Sahalia and Lo (1998) hence our discussion of its properties shall be brief. The dataset contains 16,923 daily pairs of call- and put-option prices for S&P 500 Index Options (symbol: SPX), traded on the Chicago Board Options Exchange between 4 January 1993 and 31 December 1993. We take averages of bid- and ask-prices as our raw data. Observations with time-to-maturity less than one day, implied volatility greater than 70%, and price less than 1/8 are dropped, which yields a final sample of 14,431 observations and this is the starting point for our empirical analysis.

To address problems of infrequent trading, nonsynchronous prices, and dividends, we process the raw data using the following procedure. Since all option prices are recorded at the same time on each day, we require only one temporally matched index price per day. To circumvent the unobservability of the dividend rate $\delta_{t,\tau}$, we infer the futures price $F_{t,\tau}$ for each maturity τ . By the spot-futures parity, $F_{t,\tau}$ and S_t are linked through

$$F_{t,\tau} = S_t e^{(r_{t,\tau} - \delta_{t,\tau})\tau} \tag{6.1}$$

To derive the implied futures, we use the put–call parity relation which must hold if arbitrage opportunities are to be avoided, independently of any parametric option pricing model:¹⁸

$$H(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}) + X e^{-r_{t,\tau}\tau} = G(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}) + F_{t,\tau} e^{-r_{t,\tau}\tau} \tag{6.2}$$

where G denotes the put price. To infer the futures price $F_{t,\tau}$ from this expression, we require reliable call and put prices – prices of actively traded options – at the same strike price X and time-to-expiration τ . To obtain such reliable pairs, we must use calls and puts that are closest to at-the-money (recall that in-the-money options are illiquid relative to out-of-the-money ones, hence any matched

¹⁸ Put–call parity is a pure arbitrage relationship, and as a result is almost never violated in practice. See also Black and Scholes (1973), Harvey and Whaley (1992), Kamara and Miller (1995) and Rubinstein (1985).

pair that is not at-the-money would have one potentially unreliable price). On every day t , we do this for all available maturities τ to obtain for each maturity the implied futures price from put–call parity.

Given the derived futures price $F_{t,\tau}$, we then replace the prices of all illiquid options, i.e., in-the-money options, with the price implied by put–call parity at the relevant strike prices. Specifically, we replace the price of each in-the-money call option with $G(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau}) + F_{t,\tau}e^{-r_{t,\tau}\tau} - Xe^{-r_{t,\tau}\tau}$ where, by construction, the put with price $G(S_t, X, \tau, r_{t,\tau}, \delta_{t,\tau})$ is out-of-the-money and therefore liquid. After this procedure, all the information contained in liquid put prices has been extracted and resides in corresponding call prices via put–call parity, therefore put prices may now be discarded without any loss of reliable information.

To estimate S-VaR, we collect the time series of S&P 500 index returns from the Center for Research in Security Prices up to 31 December 1993. We calculate the continuously compounded returns corresponding to holding the S&P 500 index for τ days, and use them as the inputs $\{u_\tau\}$ to construct the estimator (4.1) for each horizon τ . Each series contains $N = 1008$ observations, i.e., four years of daily observations.

6.2. The E-VaR and S-VaR of the S&P 500

We focus our empirical analysis on the six-month ($\tau = 126$ days) horizon. We report in Fig. 1 the estimated implied volatility curve $X/F \mapsto \hat{\sigma}(X/F, \tau)$ and its first three derivatives with respect to moneyness, estimated from the semiparametric model. The bandwidth values are given in Table 1. The curves in Fig. 1 are the basic inputs to our subsequent estimators of \hat{f}^* , $\hat{\zeta}$ and $\hat{\rho}$. The confidence intervals are calculated based on Proposition 1.

Fig. 2 plots the estimated E-VaR and S-VaR, in addition to a 95% confidence interval for E-VaR, for the four traded option maturities. The main difference between E-VaR and S-VaR is the difference in skewness between the two densities: E-VaR is strongly negatively skewed, while S-VaR is slightly positively skewed.

These facts combine to make the estimated ratio $\hat{\zeta}$ in Fig. 3 generally decreasing as a function of the S&P 500 value, and statistically different from one (which corresponds to risk neutrality of the representative agent).¹⁹ That $\hat{\zeta}$ is decreasing in S_T can be understood intuitively by noting that the expected value

¹⁹Because we use bandwidths that are optimized for each derivative, the derivative of the optimally estimated density is not equal to the optimal estimate of the derivative of the density. To estimate the derivative (Fig. 4) the bandwidth is larger, hence the estimate is smoother, and for that bandwidth the estimate of the MRS would have been nonincreasing throughout. However, this is not the case for the bandwidth that is optimal for the MRS (depicted in Fig. 3).

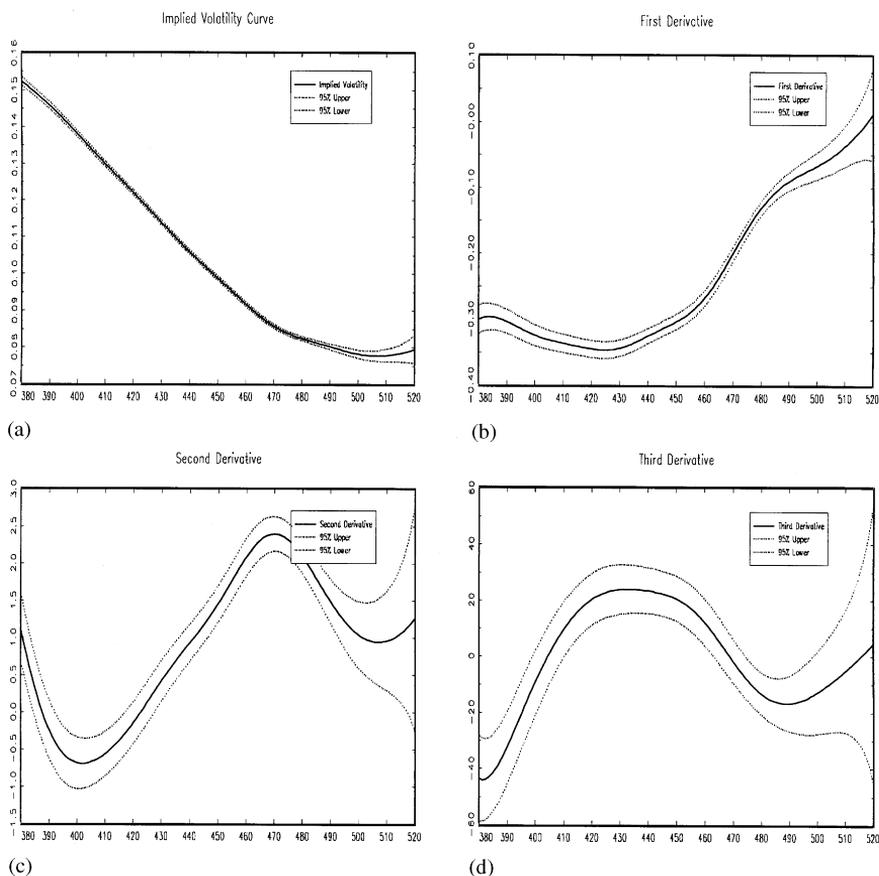


Fig. 1.

of S_T under the actual measure is empirically higher than its expected return under the risk-neutral measure, i.e.,

$$E[S_T|S_t] = \int_0^{+\infty} S_T f_t(S_T) dS_T > \int_0^{+\infty} S_T f_t^*(S_T) dS_T = E^*[S_T|S_t] \quad (6.3)$$

which, in terms of instantaneous rates of return, corresponds to $\mu > r - \delta$. Therefore, for Eq. (6.3) to hold, we expect $f_t(S_T) > f_t^*(S_T)$ for the high values of S_T and conversely for the low values of S_T ; that is, $\hat{\zeta}(S_T) > 1$ for low values of S_T and $\hat{\zeta}(S_T) < 1$ for large values of S_T . This downward-sloping pattern for $\hat{\zeta}$ is confirmed by Fig. 3.

Table 1
Bandwidth values for E-VaR and S-VaR

Variable	Estimator	Kernel	n	N	q	p	m	d	s	h	h_{cv}	H
X/F	\hat{f}^*	$k_{(2)}$	14,431	—	2	5	2	2	0.074	0.040	0.021	—
τ	\hat{f}^*	$k_{(2)}$	14,431	—	2	5	0	2	72.32	20.52	20.52	—
u	\hat{f}	$k_{(2)}$	—	1008	2	5	0	1	0.071	—	—	0.0196

Note: Bandwidth selection for the second derivative of the estimator $\hat{\sigma}(X/F_{t,\tau}, \tau)$ required in addition to the previous one to compute \hat{f}^* , and bandwidth selection for \hat{f} . The rules for $h_{X/F}$ and h_τ are $h_{X/F} = c_{X/F}s(X/F)n^{-1/(\hat{d} + 2(q_{X/F} + m))}$ and $h_\tau = c_\tau s(\tau)n^{-1/(\hat{d} + 2q)}$ where $k_{X/F}$ is of order $q_{X/F}$ and k_τ is of order q_τ , $s(X/F)$ and $s(\tau)$ are the unconditional standard deviations of the nonparametric regressors, $c_{X/F} \equiv \gamma_{X/F}/\log(n)$ with $\gamma_{X/F}$ constant, and $c_\tau \equiv \gamma_\tau \log(n)$, with γ_τ constant. h_{cv} denotes the bandwidth used to estimate the conditional variance of the nonparametric regression. The rule for H_u is $H_u = c_u s(u)N^{-1/(1 + 2(q_u + m_u))}$ where q_u is the order of the kernel k_u to estimate the m_u th derivative of $g_u(s(u))$ is the unconditional standard deviation of the returns, and $c_u \equiv \gamma_u/\log(N)$ with γ_u constant.

To quantify the relative preferences of the representative agent between a \$1 payoff in different states, consider the future value of the inverse of the ratio $\hat{\zeta}$ plotted in Fig. 3:

$$\frac{1}{e^{-r_{t,\tau}\zeta_t(S_T)}} = \frac{f_t(S_T)}{e^{-r_{t,\tau}f_t^*(S_T)}} \tag{6.4}$$

The denominator of the right-hand side is the price at t of an Arrow–Debreu security paying \$1 at $T = t + \tau$ if the S&P 500 state is between $\$S_T$ and $\$(S_T + 1)$. The numerator is the probability of that event actually being realized, that is, the expected payoff from buying the Arrow–Debreu security. Hence their ratio is one plus the expected rate of return from buying the security at t and selling it at T . We report these returns in Table 3 for different states S_T , and contrast them with those of the Black–Scholes model. Not surprisingly, the Black–Scholes model makes the negative states substantially less valuable than the nonparametric estimates do, reflecting the lack of skewness in its E-VaR for continuously compounded returns.

We then implement empirically the test for risk neutrality (Table 2) of the representative agent that was developed in Proposition 5. The null hypothesis of risk neutrality ($f^* = f$) is rejected with a p -value of 0.00, which confirms Fig. 3 (recall that risk neutrality corresponds to $\zeta = 1$, i.e., a horizontal line in Fig. 3).

6.3. The implied risk aversion of S&P 500 options

To estimate the relative risk aversion function ρ , we need to select the bandwidth values to estimate the third derivative of the call pricing function,

Comparison of E-VAR and S-VAR

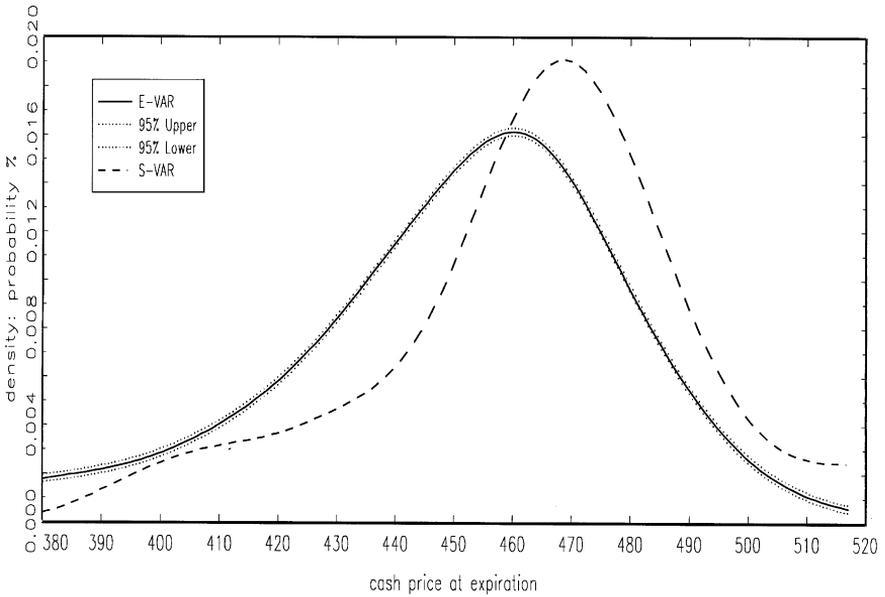


Fig. 2.

and the first derivative of S-VaR. Bandwidth values to estimate \hat{f}^{*} and \hat{f}' are reported in Table 4, and we plot the implied relative risk aversion function in Fig. 4. The confidence interval is constructed from the asymptotic distribution theory derived in Proposition 4.

The notable feature of Fig. 4 is that $\hat{\rho}$ exhibits a U-shaped pattern as a function of S_T . In other words, the market prices of S&P 500 options and the market returns on the S&P 500 index are such that the representative agent becomes more averse as the index goes down in value, as well as for very high values of the index. This phenomenon cannot be captured by CRRA preferences [see (5.7)], providing yet another characterization of the differences between market prices and the Black–Scholes model. Note also that the estimated $\hat{\rho}$ in Fig. 4 is everywhere positive, thereby implying a concave utility function.

Comparing the line (5.7) to the nonparametric estimate of $\hat{\rho}$ in Fig. 4 provides a comparison of the constant CRRA preference, where $\rho = a$ is constant, to the nonparametrically implied function of relative risk aversion. We find that the CRRA ranges between 1 and 60, depending upon the level of the S&P 500, i.e., aggregate wealth. The weighted average, over the range of S&P 500 values, of the function $\hat{\rho}$ is 12.7 and is represented by the horizontal line in Fig. 4. An interesting comparison is whether this range of values for a is “reasonably”

Scaled Marginal Rate of Substitution

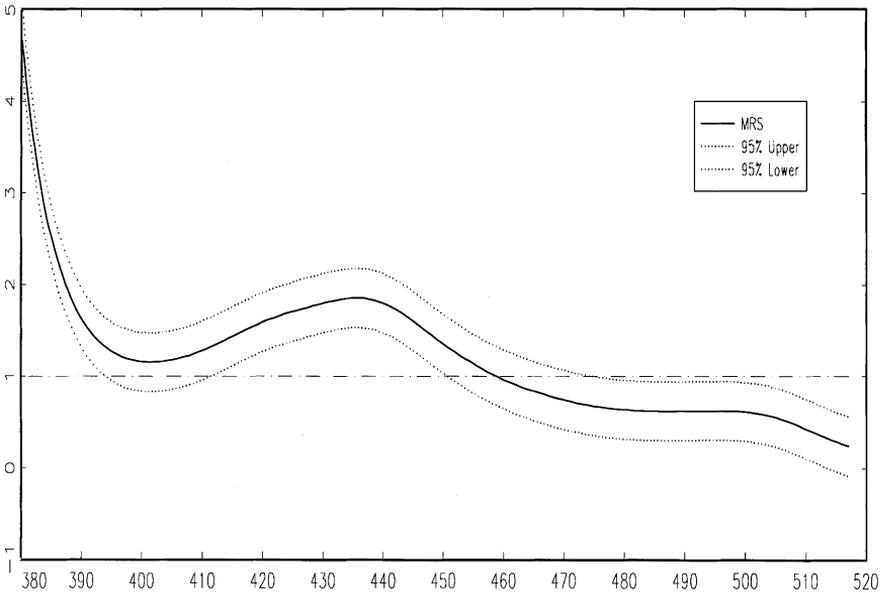


Fig. 3.

compatible with the values found using consumption, but no option, data. Is there an equity-premium-like puzzle at the levels of option prices, or do they imply coefficients of risk aversion that are less extreme than those typically found in the equity-premium literature?²⁰ This estimate for a is higher than the typical values used in theoretical models (where the range is typically 1 to 5), higher than $(\mu - r + \delta)/\sigma^2$ (the value of a in the Black–Scholes model evaluated at values of μ , δ and σ given by the S&P 500 returns and the riskfree rate r), yet generally lower than the estimates found by studies of the Euler equation in consumption-based asset pricing models.²¹ Table 5 reports the range of values of the constant CRRA that have been estimated in the literature. While the implied CRRA coefficient is informative, it is important to keep in mind that

²⁰ See Renault and Garcia (1996) for a different attempt to confront option data with the Euler equation.

²¹ However, it should be emphasized that our estimate is based on nominal data, whereas estimates obtained from consumption-based asset-pricing models use real data. For our relatively short time-horizon (less than one year), this distinction may not matter; see footnote 14 for further discussion.

Table 2
Test of risk neutrality

	Kernel	δ	η	h	h_{cv}	Trim	$H_0: f^* = f$
X/F	$k_{(2)}$	8.75	0.48	0.011856	0.01496	50/5	
τ	$k_{(2)}$	4.75	0.48	4.62255	14.6577	50/5	
Integral							54.5
Test statistic							19.7
p -value							0.00

Note: Nonparametric tests of the null hypothesis of risk neutrality of the representative agent, based on comparing globally E-VaR to S-VaR. The bandwidths are $h_{X/F} = \eta_{X/F} n^{-1/\delta_{X/F}}$ and $h_{\tau} = \eta_{\tau} n^{-1/\delta_{\tau}}$, with $\eta_{X/F}$ and η_{τ} constant, $\delta_{X/F} = \delta_{\tau} + 4$ and δ_{τ} satisfying the inequalities (A.20) and (A.21). Further, $k_{X/F} = k_{\tau} = k$ is a kernel of order $q = 2$. The bandwidth to estimate the S-VaR \hat{f} in Eq. (5.9) is H_u given by the rule (A.8) with $m_u = 0$. The average value of H_u across all maturities is 0.015833. The bandwidths to estimate the conditional variance of the nonparametric regression are $h_{X/F,cv}$ and $h_{\tau,cv}$ respectively, and are optimally smoothed to produce consistent estimates (rate $\delta_{X/F,cv} = (2q_{X/F} + d) = 6 = \delta_{cv,\tau}$). The weighting function ω is a trimming index, i.e., only observations with estimated density above a certain level, and away from the boundaries of the integration space, are retained. The two numbers in the column “Trim” refer respectively to the trimming level (as a percentage of the mean estimated density value), and the percentage trimmed at the boundary of the integration space when calculating the test statistics. For instance, if the latter is 5%, the trimming index retains the values between 1.05 times the minimum evaluation value and 0.95 times the maximum value. “Integral” refers to the percentage of the estimated density mass on the integration space that is kept by the trimming index, i.e., $\int \pi(\tilde{Y}) \omega(\tilde{Y}) d\tilde{Y}$, where π is the marginal density of the nonparametric regressors $\tilde{Y} = [X/F, \tau]$. “Test statistic” refers to the standardized distance measure between the E-VaR and S-VaR estimates (remove the bias term, divide by the standard deviation). The integral defining R is calculated over the integration space given by the rectangle $[0.85, 1.10] \times [10, 136]$ in the moneyness \times days-to-expiration space. The kernel weights are constructed using the binning method with 30 bins in the moneyness dimension and 20 in the days-to-expiration dimension.

Table 3
Expected rates of returns of S&P 500 states

State: S_T	375	400	425	450	475	500
Nonparametric rate of return	- 4.65	- 0.17	- 0.93	- 0.51	0.75	0.85
Black–Scholes rate of return	- 3.09	- 2.09	- 1.34	- 0.27	0.56	1.35

Note: Expected rate of return (annualized and continuously compounded over a horizon of $\tau = 126$ days) from an investment in an Arrow–Debreu claim that pays if the state at T falls between $\$ S_T$ and $\$ (S_T + 1)$. The nonparametric estimates are obtained with the bandwidths given in Table 1. The Black–Scholes estimate is based on $\hat{\mu} = 7.64\%$, $\hat{\sigma} = 9.74\%$, $r_{t,\tau} = 3.10\%$ and $\delta_{t,\tau} = 2.78\%$. The resulting Black–Scholes value of the CRRA coefficient is 7.72. Negative values in the table correspond to states which are more likely to occur under the risk neutral than the actual probabilities; they are correspondingly “expensive” based on the expected rate of return measure.

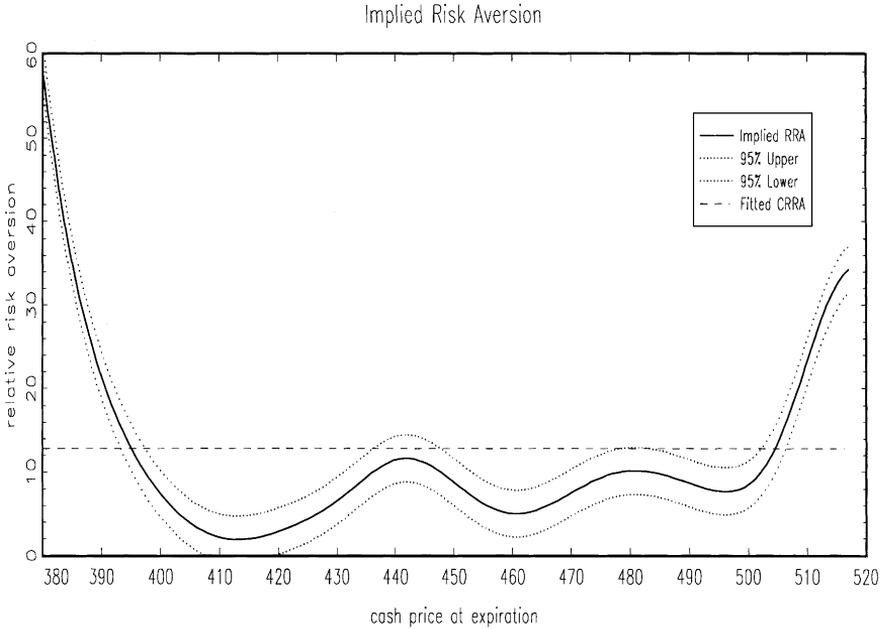


Fig. 4.

Table 4
Bandwidth values for the risk aversion coefficient

Variable	Kernel	Estimate	n	N	q	p	m	d	s	h	H
X/F	$k_{(2)}$	$\hat{f}^{* \prime}$	14,431	–	2	5	3	2	0.074	0.047	–
τ	$k_{(2)}$	$\hat{f}^{* \prime}$	14,431	–	2	5	0	2	72.324	20.52	–
u	$k_{(2)}$	\hat{f}^{\prime}	–	1008	2	5	1	1	0.071	–	0.029

Note: Bandwidth selection for the third derivative of the estimator $\hat{\sigma}(X/F_{t,\tau})$ required in addition to the previous two to compute $\hat{f}^{* \prime}$, and bandwidth selection for \hat{f}^{\prime} . The rules for $h_{X/F}$ and h_{τ} are $h_{X/F} = c_{X/F} s(X/F) n^{-1/(d+2(q_{X/F}+m))}$ and $h_{\tau} = c_{\tau} s(\tau) n^{-1/(d+2q)}$ where $k_{X/F}$ is of order $q_{X/F}$ and k_{τ} is of order q_{τ} , $s(X/F)$ and $s(\tau)$ are the unconditional standard deviations of the nonparametric regressors, $c_{X/F} \equiv \gamma_{X/F}/\log(n)$, with $\gamma_{X/F}$ constant, and $c_{\tau} \equiv \gamma_{\tau}/\log(n)$, with γ_{τ} constant. The rule for H_u to estimate the m_u -th derivative of g is $H_u = c_u s(u) N^{-1/(1+2(q_u+m_u))}$ where q_u is the order of the kernel $k_{m_u}(u)$ is the unconditional standard deviation of the returns, and $c_u \equiv \gamma_u/\log(N)$, with γ_u constant.

Fig. 4 shows that CRRA preferences are quite misspecified: the CRRA best-fitted line is quite far from the nonparametrically estimated risk-aversion curve.

Table 5
Estimated values of the constant coefficient of relative risk aversion

Study	CRRA range
Arrow (1971)	1
Friend and Blume (1975)	2
Hansen and Singleton (1982, 1984)	0–1
Mehra and Prescott (1985)	55
Ferson and Constantinides (1991)	0–12
Cochrane and Hansen (1992)	40–50

Note: Representative values of the CRRA coefficient reported in the literature, to be compared to then the fitted CRRA found here. The value in Arrow (1971) is based on a summary of a number of studies as well as theoretical considerations. Friend and Blume (1975) study individual portfolio holdings. Other values are based on Table I in Hansen and Singleton (1984); Footnote 5 in Mehra and Prescott (1985); Table 4 in Ferson and Constantinides (1991); and in Cochrane and Hansen (1992), the range 40–50 is required to fit the Hansen and Jagannathan (1991) bound (their Fig. 1), i.e., values of a of at least 40 are required to generate a variance of the stochastic discount factor implied by the equity premium region on the graph. Even at that value, the mean-standard deviation pairs do not lie inside the Hansen–Jagannathan “cup” however.

7. Conclusion

Risk management has become a first-order concern for financial managers and in this paper, we propose a new measure of value-at-risk: E-VaR. The difference between economic and statistical value-at-risk lies in the fact that E-VaR incorporates and reflects the combined effects of aggregate risk preferences, supply and demand, and probabilities; S-VaR involves only one of these effects. If aggregate preferences are risk neutral, E-VaR reduces to S-VaR. But if preferences are not risk neutral, E-VaR provides incremental information that cannot be gleaned from purely statistical measures.

Of course, this does not suggest that E-VaR is a substitute for S-VaR. S-VaR is simpler to compute and interpret, and it has found its way into certain regulatory reporting conventions. In contrast, E-VaR is considerably more difficult to interpret since it relies on a specific model of economic equilibrium and is therefore subject to the criticisms that can always be leveled against any economic model. Moreover, E-VaR is computationally more demanding, particularly for the nonparametric estimators that we have proposed. Although a parametric version of E-VaR is readily available – as in the Black–Scholes case – there is so much mounting empirical evidence against the standard parametric models for so many types of assets that we are reluctant to propose any approach other than a nonparametric one. Therefore, a distinct disadvantage of E-VaR is its computational complexity. Nevertheless, given the very nature of the risk management function, the potential benefits of E-VaR would seem to outweigh the computational costs.

In our empirical example, we have demonstrated that S-VaR and E-VaR yield considerably different risk assessments for the S&P 500 index. This suggests that E-VaR captures aspects of market risk that S-VaR does not (recall that S-VaR is a special case of E-VaR). However, conclusive evidence of the superiority of E-VaR must lie in its applications to specific risk-management processes, and we hope to develop such applications in future research.

Acknowledgements

We thank Petr Adamek, John Cochrane, George Constantinides, John Heaton, Jens Jackwerth, and Jiang Wang for helpful discussions, and are grateful to Eric Ghysels and two referees for comments on an earlier draft. Research support from MIT's Laboratory for Financial Engineering, the National Science Foundation (Grants No. SBR-9709976 and SBR-9996023), and the University of Chicago's Center for Research in Security Prices is gratefully acknowledged. Part of this research was conducted during the first author's tenure as an Alfred P. Sloan Research Fellow.

Appendix A

In this appendix, we derive the asymptotic distributions of the estimators and test statistics proposed in the text. In particular, Section A.1 contains results for E-VaR, Section A.2 contains results for S-VaR, Section A.3 contains results for the ratio $\hat{\zeta}$, Section A.4 contains results for implied risk aversion $\hat{\rho}$, and Section A.5 contains results for the test of risk neutrality.

A.1. Statistical inference for E-VaR

We collect option prices in the form of panel data, consisting of N observation periods and J options per period. The sample size is $n = NJ$. We make the following assumptions on the data used to construct the nonparametric regression (3.6), i.e., (σ, \tilde{Y}) where $\tilde{Y} \equiv [X/F_{t,\tau} \ \tau]'$ [the nonparametric regression function is $\hat{\sigma}(\tilde{Y})$, and we wish to estimate its m -th partial derivative with respect to the first component $X/F_{t,\tau}$ of the vector \tilde{Y}]:

Assumption 1.

1. The process $\{\mathbf{v}_i \equiv (\sigma_i, \tilde{Y}_i): i = 1, \dots, n\}$ is strictly stationary with $E[\sigma_i^4] < \infty$ and $E[\|\tilde{Y}_i\|^2] < \infty$, and is β -mixing with mixing coefficients β_j that decay at a rate at least as fast as j^{-b} , $b > 4$, as $j \rightarrow \infty$. The joint density of $(\mathbf{v}_1, \mathbf{v}_{1+j})$ exists for all j and is continuous.

2. The density $\pi(\sigma, \tilde{\mathbf{Y}})$ is p -times continuously differentiable with respect to $\tilde{\mathbf{Y}}$, with $p > m$, and π and its derivatives are bounded and in $L_2(\mathfrak{R}^{1+\tilde{d}})$. The marginal density of the nonparametric regressors, $\pi(\tilde{\mathbf{Y}})$, is bounded away from zero on every compact set in $\mathfrak{R}^{\tilde{d}}$.
3. $\sigma(\tilde{\mathbf{Y}})\pi(\tilde{\mathbf{Y}})$ and its derivatives are bounded. The conditional variance

$$s^2(\tilde{\mathbf{Y}}) \equiv E[(\sigma - \sigma(\tilde{\mathbf{Y}}))^2 | \tilde{\mathbf{Y}}] \tag{A.1}$$

is bounded and satisfies $s^4(\tilde{\mathbf{Y}}) \in L_2(\mathfrak{R}^{\tilde{d}})$. The conditional fourth moment $E[(\sigma - \sigma(\tilde{\mathbf{Y}}))^4 | \tilde{\mathbf{Y}}]$ is bounded.

Definition 1. A kernel function k is of order q if:

$$\int_{-\infty}^{+\infty} z^l k(z) dz = \begin{cases} 1 & \text{if } l = 0, \\ 0 & \text{if } 0 < l < q, \\ (-1)^q q! \chi_q & \text{if } l = q, \end{cases}$$

where l is an integer and $\int_{-\infty}^{+\infty} |z|^l |k(z)| dz < \infty$ for all $0 \leq l \leq q$.

Assumption 2. The kernel functions $k_{X/F}$ and k_τ are bounded, three-times continuously differentiable, and have derivatives which are bounded and in $L_2(\mathfrak{R})$. $k_{X/F}$ is of order $q_{X/F}$ and k_τ is of order q_τ . The bandwidths are given by

$$h_{X/F} = c_{X/F} s(\mathbf{X}/F) n^{-1/(\tilde{d} + 2(q_{X/F} + m))}, \quad h_\tau = c_\tau s(\tau) n^{-1/(\tilde{d} + 2q_\tau)} \tag{A.2}$$

where $s(\mathbf{X}/F)$ and $s(\tau)$ are the unconditional standard deviations of the nonparametric regressors, \mathbf{X}/F denotes the vector of observations of $X_i/F_{t_i, \tau}$, $c_{X/F} \equiv \gamma_{X/F} / \log(n)$, with $\gamma_{X/F}$ constant, and $c_\tau \equiv \gamma_\tau / \log(n)$, with γ_τ constant.

In practice, we use the kernel functions

$$k_{(2)}(z) \equiv e^{-z^2/2} / \sqrt{2}, \quad k_{(4)}(z) \equiv (3 - z^2)e^{-z^2/2} / \sqrt{8\pi} \tag{A.3}$$

which are of order $q = 2$ and $q = 4$ respectively. We then obtain

Proposition 1. Under Assumptions 1 and 2:

$$n^{1/2} h_{X/F}^{(2m+1)/2} h_\tau^{1/2} \left[\frac{\partial^m \hat{\sigma}}{\partial X^m}(\tilde{\mathbf{Y}}) - \frac{\partial^m \sigma}{\partial X^m}(\tilde{\mathbf{Y}}) \right] \xrightarrow{d} \mathcal{N}(0, \sigma_{\tilde{d}^m \sigma}^2) \tag{A.4}$$

where

$$\sigma_{d^m \sigma}^2 \equiv \frac{s^2(\tilde{Y})(\int_{-\infty}^{+\infty} (k_{X/F}^{(m)})^2(\omega) d\omega)(\int_{-\infty}^{+\infty} k_{\tau}^2(\omega) d\omega)}{\pi(\tilde{Y})F_{t,\tau}^{2m}}. \tag{A.5}$$

Therefore the E-VaR estimator is distributed asymptotically as

$$n^{1/2} h_{X/F}^{5/2} h_{\tau}^{1/2} [\hat{f}^*(S_T) - f^*(S_T)] \xrightarrow{d} \mathcal{N}(0, \sigma_{f^*}^2) \tag{A.6}$$

where $\partial H_{BS}/\partial \sigma$ is the option's gamma evaluated at $\hat{\sigma}(\tilde{Y})$ and

$$\sigma_{f^*}^2 \equiv \left[e^{r_{t,\tau}} \frac{\partial H_{BS}}{\partial \sigma}(\hat{\sigma}(\tilde{Y}), Y) \right]^2 \sigma_{d^2 \sigma}^2. \tag{A.7}$$

Here $k_{X/F}^{(m)}$ denotes the m -th derivative of the univariate kernel function k . The term $F_{t,\tau}^{2m}$ in the denominator of Eq. (A.5) is due to

$$\frac{\partial^m \hat{\sigma}(\tilde{Y})}{\partial X^m} = \frac{\partial^m \hat{\sigma}(\tilde{Y})}{\partial (X/F_{t,\tau})^m} \frac{1}{F_{t,\tau}^m}.$$

We give in Table 6 the values of the integrals of the kernel functions that appear in the expressions above for the functions (A.3) to be used in our empirical estimation of the S&P 500 E-VaR.

This proposition follows from the functional delta method in Ait-Sahalia (1997): the expression $H_{BS}(\hat{\sigma}(\tilde{Y}), Y) - H_{BS}(\sigma(\tilde{Y}), Y)$ behaves asymptotically like $\partial H_{BS}/\partial \sigma(\hat{\sigma}(\tilde{Y}) - \sigma(\tilde{Y}))$, and because derivatives of $\hat{\sigma}(\tilde{Y})$ converge at a progressively slower rate (by comparing Eq. (A.4) for $m = 0, 1$ and 2) the asymptotic distribution of $\partial^2 H_{BS}/\partial X^2(\hat{\sigma}(\tilde{Y}), Y) - \partial^2 H_{BS}/\partial X^2(\sigma(\tilde{Y}), Y)$ is that of $\partial H_{BS}/\partial \sigma(\partial^2 \hat{\sigma}/\partial X^2 - \partial^2 \sigma/\partial X^2)$.

A.2. Statistical inference for S-VaR

To estimate S-VaR, we make the following assumptions:

Assumption 3. The returns $\{u_{i,\tau}; i = 1, \dots, N\}$ are strictly stationary, with $E[u_{i,\tau}^4] < \infty$. Their marginal density g admits $p_u = 4$ continuous derivatives, and is bounded away from zero on every compact set in \mathfrak{R} . In addition, they are β -mixing at a rate β_j decaying as $j \rightarrow \infty$ at a rate at least as fast as j^{-b} , $b > 4$. The joint density of $(u_{t_1,\tau}, u_{t_1+j,\tau})$ exists for all j and is continuous.

Assumption 4. The kernel function k_u is in $L_2(\mathfrak{R})$ and is of order q_u . The bandwidth H_u to estimate the m_u th derivative of g , $m_u \leq p_u$, is given by

$$H_u = c_u s(u) N^{-1/(1+2(q_u+m_u))} \tag{A.8}$$

where $s(u)$ is the unconditional standard deviation of the returns, $c_u \equiv \gamma_u/\log(N)$, with γ_u constant.

We then obtain

Proposition 2. Under Assumptions 3 and 4,

$$N^{1/2} H_u^{1/2} [\hat{g}(u) - g(u)] \xrightarrow{d} \mathcal{N}(0, \sigma_g^2) \tag{A.9}$$

where $\sigma_g^2 \equiv g(u)(\int_{-\infty}^{+\infty} k_u^2(\omega) d\omega)$. Therefore

$$N^{1/2} H_u^{1/2} [\hat{f}_i(S_T) - f_i(S_T)] \xrightarrow{d} \mathcal{N}(0, \sigma_f^2) \tag{A.10}$$

where

$$\sigma_f^2 \equiv \left(\int_{-\infty}^{+\infty} k_u^2(\omega) d\omega \right) f_i(S_T) / S_T.$$

The asymptotic distribution of the S-VaR estimator follows from the same technique as in Section A.1. To obtain σ_f^2 , simply multiply through Eq. (A.9) by $1/S_T$, and replace according to Eq. (4.3).

A.3. Statistical inference for the ratio $\hat{\zeta}$

As discussed in Section 4, deriving the asymptotic distribution of the ratio $\hat{\zeta} \equiv \hat{f}^*/\hat{f}$ is made simpler by the fact that the S-VaR estimator \hat{f} converges faster than the E-VaR estimator \hat{f} . Indeed, compare Eq. (A.6) with Eq. (A.10). While \hat{f}^* is obtained by twice differentiating a nonparametric regression function, \hat{f} is a direct nonparametric density estimator. Thus \hat{f}^* converges slower, *ceteris paribus*; on the other hand, \hat{f}^* is estimated using a larger sample size (based on the entire panel data of option prices, $n = NJ$) while \hat{f} is estimated using only N observations (based only on the time series of index returns), thus \hat{f}^* converges faster, *ceteris paribus*. By substituting the specific bandwidth rules given in Eqs. (A.2) and (A.8), and $\bar{d} = 2$, $p_{X/F} = p_\tau = 5$ and $p_u = 5$, we find that the ratio

of the convergence rates in Eqs. (A.6) and (A.10) is given by

$$\begin{aligned} \frac{n^{1/2} h_{X/F}^{5/2} h_{\tau}^{1/2}}{N^{1/2} H_u^{1/2}} &= O\left(\frac{\log(N)n^{(2p_{X/F}-5)/2(\bar{d}+2p_{X/F})}}{\log(n)N^{p_u/(1+2p_u)}}\right) \\ &= O(N^{(2p_{X/F}-5)/2(\bar{d}+2p_{X/F})-p_u/(1+2p_u)}) = o(1) \end{aligned}$$

As a result, the asymptotic distribution of \hat{f}^*/\hat{f} is identical to that of \hat{f}^*/f — that is, we can treat \hat{f} as fixed at its true value f for the purpose of conducting inference on $\hat{\zeta}$. More specifically, the correlation between the two estimators induced by the use of the same data is of second order relative to the sampling variation of the slowly converging E-VaR estimator, and is therefore discarded asymptotically. It follows immediately that

Proposition 3. Under Assumptions 1–4:

$$n^{1/2} h_{X/F}^{5/2} h_{\tau}^{1/2} [\hat{\zeta}(S_T) - \zeta(S_T)] \xrightarrow{d} \mathcal{N}(0, \sigma_{\zeta}^2) \tag{A.11}$$

where

$$\sigma_{\zeta}^2 \equiv \sigma_{f^*/f}^2(S_T). \tag{A.12}$$

A.4. Statistical inference for implied risk aversion $\hat{\rho}$

We now derive the asymptotic distribution of our estimator of the Arrow–Pratt risk aversion measure $\hat{\rho}(\cdot)$. We are fortunate to be again in a context where one estimator, namely \hat{f}^* , converges slower than every other

Table 6
Kernel constants in asymptotic distributions

m	0	1	2	3
$\int_{-\infty}^{+\infty} (k_{(2)}^{(m)})^2(w)dw$	$\frac{1}{2\sqrt{\pi}}$	$\frac{1}{4\sqrt{\pi}}$	$\frac{3}{8\sqrt{\pi}}$	$\frac{31}{16\sqrt{\pi}}$
$\int_{-\infty}^{+\infty} (k_{(4)}^{(m)})^2(w)dw$	$\frac{27}{32\sqrt{\pi}}$	$\frac{175}{64\sqrt{\pi}}$	$\frac{273}{128\sqrt{\pi}}$	$\frac{2025}{256\sqrt{\pi}}$

Note: Kernel constants that characterize the asymptotic variances of the nonparametric estimators in Propositions 1–4. The kernel functions $k_{(2)}$ and $k_{(4)}$ are defined in (A.3). $k^{(m)}$ denotes the m -th derivative of k .

Table 7
Kernel constants in test statistic

m	0	1	2
$\int_{-\infty}^{+\infty} (\int_{-\infty}^{+\infty} k_{(2)}^{(m)}(w) k_{(2)}^{(m)}(w+v) dw)^2 dv$	$\frac{1}{2\sqrt{2\pi}}$	$\frac{3}{32\sqrt{2\pi}}$	$\frac{105}{512\sqrt{2\pi}}$

Note: Kernel convolution constants that characterize the asymptotic variance of the test statistic in Proposition 5. The kernel function $k_{(2)}$ is defined in (A.3), and $k_{(2)}^{(m)}$ denotes its m -th derivative.

one (\hat{f}^* , \hat{f}' and \hat{f}) appearing in the ratio (5.3), all of which can then be taken as fixed for the purpose of computing the asymptotic distribution of $\hat{\rho}$. That is, $\hat{\rho}$ behaves asymptotically like $-\hat{f}^*/\hat{f}'$, which in turn behaves asymptotically like $-\hat{f}^*/\hat{f}'$.

To estimate the first derivative of f^* , we select the bandwidth as in Eq. (A.2), except that we now assume $p_{X/F} = 5$ and have $m = 3$. The kernel choices are identical to those used to compute \hat{f}^* itself, that is $k_\tau = k_{(4)}$ and $k_{X/F} = k_{(2)}$. We can apply Eqs. (A.4)–(A.5) to $m = 3$ to obtain the distribution of $\partial^3 \hat{\sigma}(\tilde{Y})/\partial X^3 - \partial^3 \sigma(\tilde{Y})/\partial X^3$ with asymptotic variance $\sigma_{\hat{\sigma}}^2$. Under the same assumptions as Proposition 1, the asymptotic distribution of $\hat{f}^{* \prime}$ is given by that of $\partial H_{BS}/\partial \sigma(\partial^3 \hat{\sigma}/\partial X^3 - \partial^3 \sigma/\partial X^3)$, i.e.:

$$n^{1/2} h_X^{7/2} h_\tau^{1/2} [\hat{f}^{* \prime}(S_T) - f^{* \prime}(S_T)] \xrightarrow{d} \mathcal{N}(0, \sigma_{\hat{f}^*}^2) \tag{A.13}$$

where

$$\sigma_{\hat{f}^*}^2 \equiv \left[\frac{\partial H_{BS}}{\partial \sigma}(\hat{\sigma}(\tilde{Y}), Y) \right]^2 \sigma_{\hat{\sigma}}^2 \tag{A.14}$$

To estimate f'_i , we take the total derivative of (4.4) with respect to S_T , i.e., $f'_i(S_T) = (-1/S_T^2)g(\log(S_T/S_i)) + (1/S_T^2)g'(\log(S_T/S_i))$ and note that the leading term in the asymptotic distribution will be the slow converging $(1/S_T^2)\hat{g}'(\log(S_T/S_i))$. We derive the distribution of \hat{f}'_i for completeness only since, as discussed above, it does not influence the distribution of $\hat{\rho}_i$. Set the bandwidth for \hat{f}'_i as in Eq. (A.8), except that we now assume $p_u = 3$ and have $m_u = 1$. The kernel choice is $k_u = k_{(2)}$, and from

$$N^{1/2} H_u^{3/2} [\hat{g}'(u) - g'(u)] \xrightarrow{d} \mathcal{N}(0, \sigma_g^2) \tag{A.15}$$

where $\sigma_{g'}^2 \equiv g(u)(\int_{-\infty}^{+\infty} (k'_u)^2(\omega) d\omega)$ we obtain, under the same assumptions as Proposition 2 – simply multiply through Eq. (A.15) by $1/S_T^2$ and recall Eq. (4.3):

$$N^{1/2} H_u^{3/2} [\hat{f}_t'(S_T) - f_t'(S_T)] \xrightarrow{d} \mathcal{N}(0, \sigma_{g'}^2) \tag{A.16}$$

where

$$\sigma_{f'}^2 \equiv \frac{f_t(S_T)}{S_T^3} \left(\int_{-\infty}^{+\infty} (k'_u)^2(\omega) d\omega \right). \tag{A.17}$$

We therefore obtain:

Proposition 4. Under Assumptions 1–4

$$n^{1/2} h_{X/F}^{7/2} h_\tau^{1/2} [\hat{\rho}_t(S_T) - \rho_t(S_T)] \xrightarrow{d} \mathcal{N}(0, \sigma_\rho^2) \tag{A.18}$$

where

$$\sigma_\rho^2 \equiv \sigma_{f'}^2 \cdot S_T^2 / f^{*2}(S_T). \tag{A.19}$$

A.5. A test of risk neutrality

In practice, we evaluate numerically the test statistic (an integral) on a rectangle of values of \tilde{Y} representing the support of the density π , and use the binning method to evaluate the kernels [see e.g., Wand and Jones (1995) for a description of the binning method]. To construct the estimators involved in Eq. (5.9), we use the following bandwidth and kernel functions:

Assumption 5. The bandwidths $h_{X/F}$ and h_τ to estimate the E-VaR \hat{f}^* in Eq. (5.9) are given by

$$h_{X/F} = \eta_{X/F} n^{-1/\delta_{X/F}}, \quad h_\tau = \eta_\tau n^{-1/\delta_\tau}$$

where $\eta_{X/F}$ and η_τ are constant, and $\delta_{X/F} = \delta_\tau + 4$ and δ_τ satisfy

$$\left(\frac{3}{\delta_\tau + 6} \right) + \left(\frac{1}{\delta_\tau} \right) < \frac{1}{2} < \left(\frac{q + 9/2}{\delta_\tau + 4} \right) + \left(\frac{1/2}{\delta_\tau} \right) \tag{A.20}$$

and

$$\left(\frac{3/2}{\delta_\tau + 4}\right) > \left(\frac{1/2}{\delta_\tau}\right). \tag{A.21}$$

Furthermore, $k_{X/F} = k_\tau = k$ is a kernel of order q . The bandwidth to estimate the S-VaR \hat{f} in Eq. (5.9) is identical to H_u given by (A.8) with $m_u = 0$.

The distribution of the test statistic is given by

Proposition 5. Under Assumptions 1, 3 and 5, and under H_0

$$nh_{X/F}^{9/2}h_\tau^{1/2}\hat{R}(\hat{f}^*, \hat{f}) - h_{X/F}^{-1/2}h_\tau^{-1/2}B_R \xrightarrow{d} \mathcal{N}(0, \Sigma_R^2) \tag{A.22}$$

where

$$B_R = \left(\int_{-\infty}^{+\infty} k''^2(w) dw\right)\left(\int_{-\infty}^{+\infty} k^2(w) dw\right)\int_{\tilde{Y}} \sigma^2(\tilde{Y})\tilde{\omega}(\tilde{Y}) d\tilde{Y}, \tag{A.23}$$

$$\Sigma_R^2 = 2\left[\int_{-\infty}^{+\infty}\left(\int_{-\infty}^{+\infty} k''(w)k''(w+v) dw\right)^2 dv\right] \left[\int_{-\infty}^{+\infty}\left(\int_{-\infty}^{+\infty} k(w)k(w+v) dw\right)^2 dv\right]\int_{\tilde{Y}} \sigma^4(\tilde{Y})\tilde{\omega}^2(\tilde{Y}) d\tilde{Y} \tag{A.24}$$

and

$$\tilde{\omega}(\tilde{Y}) \equiv \left(e^{r_{t,t}}\frac{\partial H_{BS}(\sigma(\tilde{Y}), \mathbf{Y})}{\partial \sigma}\right)^2 \omega(\tilde{Y}). \tag{A.25}$$

We report in Table 7 the values of the new kernel integrals that appear in the expressions above for the actual choice $k = k_{(2)}$. Distribution (A.22) follows again from the functional delta method, applied now to derive the second-order term in the expansion of the test statistic (under the null, the liner term is degenerate). It is remarkable that the estimation of S-VaR does not contribute any term to the asymptotic distribution of the test statistic. In addition, only the higher-order derivative of the nonparametric regression function matter. This follows from the fact that $(e^{r_{t,t}}\{\partial H_{BS}(\sigma_{BS}, \mathbf{Y})/\partial \sigma\}\{\partial^2 \hat{\sigma}(\tilde{Y})/\partial X^2 - \partial^2 \sigma(\tilde{Y})/\partial X^2\})^2$ is the leading term in the expansion of $(\hat{f}_t^*(S_T) - f_t^*(S_T))^2$, which is itself the leading

term in the expansion of

$$\begin{aligned}
 (\hat{f}_t^*(S_T) - \hat{f}_t(S_T))^2 &= (\hat{f}_t^*(S_T) - f_t^*(S_T) + f_t(S_T) - \hat{f}_t(S_T))^2 \\
 &= (\hat{f}_t^*(S_T) - f_t^*(S_T))^2 - 2(\hat{f}_t^*(S_T) \\
 &\quad - f_t^*(S_T))(\hat{f}_t(S_T) - f_t(S_T)) + (\hat{f}_t(S_T) - f_t(S_T))^2 \\
 &= (\hat{f}_t^*(S_T) - f_t^*(S_T))^2 + o(\|\hat{f}_t^* - f_t^*\|^2).
 \end{aligned}$$

To estimate consistently the conditional variance of the regression, $s^2(\tilde{Y})$, we calculate the difference between the kernel estimate of the regression of the squared dependent variable σ^2 on \tilde{Y} and the squared of the regression $\sigma(\tilde{Y})$ of the dependent variable σ . The regression $E[\sigma^2|\tilde{Y}]$ is estimated with bandwidth $h_{cv} = \eta_{cv}n^{-1/\delta_{cv}}$, $\delta_{cv} = \delta_\tau$ and η_{cv} constant. Prior to computing the estimates, we standardize the regressors \tilde{Y} by removing their respective sample means, and dividing by their sample unconditional standard deviations.

The test statistics are formed by standardizing the asymptotically normal distance measure \hat{R} : we estimate consistently the asymptotic mean B_R and variance Σ_R^2 , then subtract the mean and divide by the standard deviation. The test statistic then has the asymptotic $N(0, 1)$ distribution. Since the test is one-sided (we only reject when \hat{R} is too large, hence when the test statistic is large and positive), the 10% critical value is 1.28, while the 5% value is 1.64. We fix the variables in Y that are excluded from \tilde{Y} at their sample means.

References

- Ait-Sahalia, Y., 1996a. Nonparametric pricing of interest rate derivative securities. *Econometrica* 64, 527–560.
- Ait-Sahalia, Y., 1996b. Testing continuous-time models of the spot interest rate. *Review of Financial Studies* 9, 385–426.
- Ait-Sahalia, Y., 1997. The delta method for nonparametric kernel functionals. Working paper, University of Chicago.
- Ait-Sahalia, Y., Lo, A., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 53, 499–548.
- Arrow, K.J., 1964. The role of securities in the optimal allocation of risk bearing. *Review of Economic Studies* 31, 91–96.
- Arrow, K.J., 1971. *Essays in the Theory of Risk Bearing*. North-Holland, Amsterdam.
- Banz, R., Miller, M., 1978. Prices for state-contingent claims: some estimates and applications. *Journal of Business* 51, 653–672.
- Bick, A., 1987. On the consistency of the Black–Scholes model with a general equilibrium framework. *Journal of Financial and Quantitative Analysis* 22, 259–275.
- Bick, A., 1990. On viable diffusion price processes of the market portfolio. *Journal of Finance* 45, 673–689.

- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–659.
- Breeden, D., 1979. An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics* 7, 265–296.
- Breeden, D., Litzenberger, R., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621–651.
- Brennan, M.J., 1979. The pricing of contingent claims in discrete-time models. *Journal of Finance* 34, 53–68.
- Broadie, M., DeTemple, J., Ghysels, E., Torrès, O., 1996. Nonparametric estimation of American options exercise boundaries and call prices. Working paper, Columbia University.
- Brown, D.P., Gibbons, M.R., 1985. A simple econometric approach for utility-based asset pricing models. *Journal of Finance* 40, 359–382.
- Campbell, J., Lo, A., MacKinlay, C., 1997. *The Econometrics of Financial Markets*, Princeton University Press, Princeton, NJ.
- Cochrane, J.H., Hansen, L.P., 1992. Asset pricing explorations for macroeconomics. In: 1992 NBER Macroeconomics Annual, NBER, Cambridge, MA.
- Cochrane, J.H., Saa-Requejo, J., 1996. Good deals pricing of derivatives in incomplete markets. Working paper, University of Chicago.
- Constantinides, G.M., 1982. Intertemporal asset pricing with heterogeneous consumers and without demand aggregation. *Journal of Business* 55, 253–268.
- Cox, J., Ross, S., 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145–166.
- Debreu, G., 1959. *Theory of Value*. Wiley, New York.
- Derman, E., Kani, I., 1994. Riding on the smile. *RISK* 7, 32–39.
- Duffie, D., 1996. *Dynamic Asset Pricing Theory*, Second ed. Princeton University Press, Princeton, NJ.
- Duffie, D., Huang, C., 1985. Implementing Arrow-Debreu equilibria by continuous trading of few long-lived securities. *Econometrica* 53, 1337–1356.
- Dumas, B., Fleming, J., Whaley, R.E., 1995. Implied volatility functions: empirical tests. Working paper, Fuqua School of Business, Duke University.
- Dupire, B., 1994. Pricing with a smile. *RISK* 7, 18–20.
- Ferson, W.E., Constantinides, G.M., 1991. Habit persistence and durability in aggregate consumption: empirical tests. *Journal of Financial Economics* 29, 199–240.
- Friend, I., Blume, M.E., 1975. The demand for risky assets. *American Economic Review* 65, 900–922.
- Garman, M., 1978. The pricing of supershares. *Journal of Financial Economics* 6, 3–10.
- Goldberger, D.H., 1991. A unified method for pricing options on diffusion processes. *Journal of Financial Economics* 29, 3–34.
- Grossman, S., Shiller, R., 1982. Consumption correlatedness and risk measurement in economies with non-traded assets and heterogeneous information. *Journal of Financial Economics* 10, 195–210.
- Hansen, L.P., Jagannathan, R., 1991. Implications of security market data for models of dynamic economies. *Journal of Political Economy* 99, 225–262.
- Hansen, L.P., Singleton, K.J., 1982. Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica* 50, 1269–1286.
- Hansen, L.P., Singleton, K.J., 1983. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy* 91, 249–265.
- Hansen, L.P., Singleton, K.J., 1984. Errata: generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica* 52, 267–268.
- Hakansson, N., 1976. Purchasing power funds: A new kind of financial intermediary. *Financial Analysts Journal* 32, 49–59.

- Härdle, W., 1990. *Applied Nonparametric Regression*, Cambridge University Press, Cambridge, UK.
- Harrison, M., Kreps, D., 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20, 381–408.
- Harvey, C.R., Whaley, R.E., 1992. Market volatility prediction and the efficiency of the S&P 100 index option market. *Journal of Financial Economics* 31, 43–73.
- He, H., Leland, H., 1993. On equilibrium asset price processes. *Review of Financial Studies* 6, 593–617.
- Hutchinson, J., Lo, A., Poggio, T., 1994. A nonparametric approach to the pricing and hedging of derivative securities via learning networks. *Journal of Finance* 49, 851–889.
- Jackwerth, J.C., 1997. Recovering risk aversion from option prices and realized returns. Working paper, Haas School of Business, University of California, Berkeley.
- Jackwerth, J.C., Rubinstein, M., 1996. Recovering probability distributions from contemporary security prices. *Journal of Finance* 51, 1611–1631.
- Jarrow, R., Rudd, A., 1982. Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics* 10, 347–369.
- Kamara, A., Miller, T.W. Jr, 1995. Daily and intraday tests of European put-call parity. *Journal of Financial and Quantitative Analysis* 30, 519–539.
- Kirman, A., 1992. What or whom does the representative individual represent? *Journal of Economic Perspectives* 6, 117–136.
- Kocherlakota, N., 1996. The equity premium: it's still a puzzle. *Journal of Economic Literature* 34, 42–71.
- Lo, A.W., Wang, J., 1995. Option pricing models when asset returns are predictable. *Journal of Finance* 50, 87–129.
- Longstaff, F., 1995. Option pricing and the martingale restriction. *Review of Financial Studies* 8, 1091–1124.
- Lucas, R., 1978. Asset prices in an exchange economy. *Econometrica* 46, 1429–1446.
- Madan, D.B., Milne, F., 1994. Contingent claims valued and hedged by pricing and investing in a basis. *Mathematical Finance* 4, 223–245.
- Mehra, R., Prescott, E., 1985. The equity premium: a puzzle. *Journal of Monetary Economy* 15, 145–161.
- Merton, R., 1973. Rational theory of option pricing. *Bell Journal of Economics and Management Science* 4, 141–183.
- Merton, R., 1992. *Continuous-Time Finance*, Revised Ed. Basil Blackwell, Oxford, UK.
- Rady, S., 1994. State prices implicit in valuation formulae for derivative securities: A martingale approach. Discussion Paper No. 181, LSE Financial Markets Group, London, UK.
- Renault, E., Garcia, R., 1996. Risk aversion, intertemporal substitution and option pricing. Working paper, Université de Montreal.
- Rogoff, K., 1996. The purchasing power parity puzzle. *Journal of Economic Literature* 34, 647–668.
- Ross, S., 1976. Options and efficiency. *Quarterly Journal of Economics* 90, 75–89.
- Rubinstein, M., 1976. The valuation of uncertain income streams and the pricing of options. *Bell Journal of Economics* 7, 407–425.
- Rubinstein, M., 1985. Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from 23 August 1976 through 31 August 1978. *Journal of Finance* 40, 455–480.
- Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 49, 771–818.
- Shiller, R., 1981. Do stock prices move too much to be justified by subsequent changes in dividends? *American Economic Review* 71, 421–436.
- Shimko, D., 1993. Bounds of probability. *RISK* 6, 33–37.
- Smithson, C., Smith, C., Wilford, D., 1995. *Managing Financial Risk*, Richard D. Irwin, Chicago, IL.

- Stutzer, M., 1996. A simple nonparametric approach to derivative security valuation. *Journal of Finance* 51, 1633–1652.
- Waud, M.P., Jones, M.C., 1995. *Kernel Smoothing*. Chapman & Hall, London, UK.
- Wang, S., 1993. The integrability problem of asset prices. *Journal of Economic Theory* 59, 199–213.
- White, H., 1992. Nonparametric estimation of conditional quantiles using neural networks. In *Proceedings of the Symposium on the Interface*. Springer Verlag, New York, NY.