

Empirical option pricing: a retrospection

David S. Bates^{a,b,*}

^a*Department of Finance, Henry B. Tippie College of Business Administration, University of Iowa, 108 Pappajohns Bus. Ad. Bldg., Suite S252, Iowa City, IA 52242-1000, USA*

^b*National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA 02138, USA*

Abstract

This article provides an overview and discussion of empirical option pricing research: how we test models, what we have learned, and what are some key issues. Some suggestions for future research are provided.

© 2003 Elsevier B.V. All rights reserved.

JEL classification: G13

Keywords: Option pricing; Derivatives; Risk-neutral distributions; Time series analysis; Affine models

0. Introduction

In 1995, I was asked to write a survey on testing option pricing models for the *Handbook of Statistics*. I initially assumed the survey would focus mostly on tests of the latest theoretical models—those that extended the original Black–Scholes–Merton (BSM) model by adding jumps and/or stochastic volatility. However, I soon found that the overwhelming bulk of the empirical literature was set within the BSM paradigm, with relatively few papers that tested alternative models. With regard to those models, I was forced to speculate to some degree on what would be the result of tests that had not yet been run.

Times have changed. In the 6 years since that survey, there has been substantial empirical research into alternative option pricing models. Furthermore, there have been substantial parallel advances in time series econometrics; in particular, the development of numerically intensive techniques compatible with the continuous-time models commonly employed in derivatives research. The result has been a renewed focus on the degree to which the “risk-neutral” models consistent with observed bond and/or

* Tel.: +1-319-353-2288; fax: +1-319-335-3690.

E-mail address: david-bates@uiowa.edu (D.S. Bates).

option prices are also consistent with the underlying objective conditional distributions estimated from time series analysis.

It is my intent in this article to reflect on what we have achieved in empirical option pricing research, what are the key issues, and where we may be heading. While not intended as a detailed survey of recent empirical work,¹ an overview of issues and approaches may be of use to researchers interested in working in the area. Furthermore, since most of the econometric issues in option pricing are also relevant for bond pricing, some of the parallels will be emphasized. I begin in Section 1 with a basic framework for categorizing empirical derivatives research, and discuss some of what we have learned from proposing and testing alternative option pricing models. In Section 2, I discuss some major remaining puzzles in option pricing. I conclude in Section 3 with some suggestions for possible research topics.

1. Empirical derivatives research

1.1. A basic framework

Derivatives are products with payoffs and prices dependent upon the stochastic evolution of associated underlying financial variables. Option payoffs depend upon the evolution of the asset price upon which the options are written, while bond prices of various maturities are sensitive to the stochastic properties of short-term interest rates. Theoretical derivatives research is typically concerned with developing plausible models of the time series properties of the underlying financial variables, and identifying appropriate risk premia for various risks such as interest rate, volatility and jump risk that are not directly priced by other traded assets. These “objective” time series models and risk premia imply an associated “risk-neutral” probability measure that can be used to price any derivative at the expected discounted value of its future payoff. While the objective and “risk-neutral” probability measures are related, and may share some common parameters, they are identical only in the case of zero risk premia on all relevant risks.

Empirical tests of bond and option pricing models examine whether various facets of the joint transition densities of the underlying and of derivatives prices are in fact consistent with information contained in the cross-section of derivatives prices, after appropriate risk adjustments. The two major approaches have involved examining the derivatives pricing implications of models estimated from time series analysis, and examining particular model-specific predictions for joint asset and derivatives returns *conditional* on information inferred from the cross-section of derivatives prices. The latter approach has been more common, especially for models with latent variables such as stochastic volatility. The values of those latent variables are more easily inferred from derivatives prices than estimated from time series analysis, under the null hypothesis of correct model specification.

¹ See Garcia et al. (2001) for such a survey.

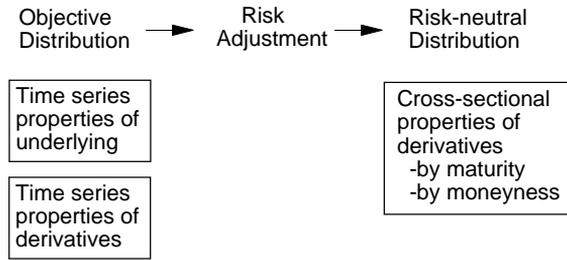


Fig. 1. Objective and risk-neutral distributions.

Within the empirical bond pricing literature, the key issue is whether the term structure of interest rates is consistent with the stochastic evolution of the short-term interest rate and of bond prices. For example, the expectation hypothesis tests discussed by Campbell (1995) examine whether yield spreads between n - and 1-period bonds correctly predict future interest rate and bond yield increases. Multifactor bond pricing models frequently use time series and cross-sectional implications when estimating stochastic processes. Examples include the 2-factor CIR model of Pearson and Sun (1994), the Heath et al. (1992) method of modeling stochastic forward rate curve evolution consistently with an initial term structure of forward rates, and the Ahn et al. (2003) and Jagannathan et al. (2003) papers in this issue.

Within the empirical option pricing literature, model-specific implications for objective and risk-neutral distributions and how they interact create a considerable variety of possible tests of any given model. For example, consider the myriad ways the BSM model has been tested over the last 30 years.

- (1) *Time series properties of the underlying asset price*: The assumption of constant instantaneous volatility of returns or of log-differenced asset prices has been rejected by the ARCH literature, while the assumption that asset prices follow a diffusion has been rejected by the fat-tailed properties of returns. The “leverage” effects noted by Black (1976) and subsequent researchers indicated more complicated interactions between stock returns and stock volatility than could be accommodated by the assumption of geometric Brownian motion.
- (2) *Cross-sectional properties of option prices*: BSM implies only one free parameter σ , so the implicit standard deviations (ISDs) inferred from option prices should be identical across different strike prices (moneyness) and maturities. This has been rejected by the non-flat “smile” and “smirk” patterns observed across different strike prices, and by non-flat term structure patterns.
- (3) *Time series properties of option prices*: BSM implies that ISDs should also be constant across time—a hypothesis clearly rejected by time series graphs of ISDs.
- (4) *Cross-sectional option prices and asset returns*: The fundamental compatibility between the ISD inferred from options and the volatility of asset returns has typically been tested by regressing realized volatility measures upon ISDs. Typical

results summarized by Bates (1996b, pp. 593–595) are that ISDs are frequently biased and possibly inefficient predictors of realized asset return volatility.

- (5) *Cross-sectional option prices versus option returns*: Stein (1989), Diz and Finucane (1993), and Campa and Chang (1995) examined whether the term structure of ISDs predicts subsequent evolution in short-term ISDs, with mixed results.
- (6) *Joint distribution of asset returns and option returns*: Bakshi et al. (2000) note that BSM and other univariate Markov models imply that call option returns should be directly related (and put options inversely related) to underlying asset returns, while call and put returns should be inversely related. They find these predictions are rejected 7–17% of the time at intradaily and daily data frequencies for S&P 500 index options in 1994.
- (7) *Option prices versus the joint distribution of asset returns and option returns*: The no-arbitrage foundations of BSM imply that option returns should be a levered function of asset returns:

$$\frac{\Delta \tilde{C}}{C} = \frac{SC_S^{\text{BS}}}{C^{\text{BS}}} \bigg|_{\hat{\sigma}} \frac{\Delta \tilde{S}}{S} + O(\Delta t) \quad (1)$$

with the leverage estimate based typically on an ISD inferred from recent option prices.² The null hypothesis is therefore that there is a *degeneracy* in the joint distribution of option and asset returns that is identifiable from the level of option prices. Profitability tests examine the average returns from self-financing portfolios of options and accompanying asset hedge positions, which since (theoretically) riskless should earn the riskless rate.³ Hedge ratio tests examine the magnitude of risk reduction from hedging. The Bakshi et al. (2000) results imply the BSM prediction of zero risk cannot be achieved.

Tests of the BSM model were often run in a somewhat ad hoc fashion: testing specific implications of the model while ignoring other obvious inconsistencies. For instance, the fact that options of different strike prices and maturities yielded different ISDs initially generated a debate on how to construct weighted ISD estimates for use in examining the conditional distribution of asset and/or option returns,⁴ rather than being taken as evidence against the model. Similarly, the use of contemporaneous or lagged ISDs when forecasting volatility and computing hedge ratios was inconsistent with the BSM prediction that ISDs should be constant over time. The focus of this literature was often on how well ad hoc versions of the model worked for practitioners; not on formally testing the null hypothesis of the BSM model.

1.2. Alternative models

The empirical deficiencies of the BSM model have prompted research along three dimensions. The *univariate diffusion* models maintained the diffusion and no-arbitrage

² The original Black and Scholes (1972) test of their model used volatility estimates from asset returns, which fall within the sixth test classification above.

³ See Coval and Shumway (2001) for a recent example.

⁴ See Bates (1996b, pp. 587–589).

foundations of BSM, but relaxed the assumption of geometric Brownian motion. Examples include the constant elasticity of variance model of Cox and Ross (1976) and Cox and Rubinstein (1985); the leverage models of Geske (1979) and Rubinstein (1983); and the more recent implied binomial and trinomial trees models. *Stochastic volatility* models allow the instantaneous volatility of asset returns to evolve stochastically over time—typically as a diffusion, but possibly as a regime-switching (Naik, 1993) or jump-diffusion process (Duffie et al., 2000). Jump models such as Merton (1976) relax the diffusion assumption for asset prices. Recent approaches have combined features of these three approaches. Stochastic generalizations of binomial tree models are surveyed by Skiadopoulos (2000), while the currently popular affine class of distributional models creates a structure that nests particular specifications of the stochastic volatility and jump approaches.

While alternatives to the BSM model began appearing within a few years of the original Black and Scholes (1973) paper, and continued in the 1980s with the stochastic volatility models, actually testing those alternatives using options data was handicapped by several factors. First, options data were not readily available until after the advent of option trading on centralized exchanges. Second, the new models were perforce more complicated. Pricing options was a substantial computational burden for the (main-frame) computers available at the time, while employing these option pricing models in empirical work was a major computational and programming challenge. Third, the extra risks introduced by jumps and stochastic volatility raised the issue of how to price those risks. The early literature tended to assume zero risk premia; e.g., Merton (1976) for jump risk, and the stochastic volatility models of Hull and White (1987), Johnson and Shanno (1987), Scott (1987), and Wiggins (1987).

The 1990s witnessed several important developments that have facilitated greater empirical scrutiny of alternative option pricing models. The most important has been the Fourier inversion approach to option pricing of Stein and Stein (1991) and Heston (1993). While other methods already existed for pricing options under general stochastic volatility processes—e.g., the Monte Carlo approach of Scott (1987), or the higher-dimensional finite-difference approach of Wiggins (1987)—Fourier inversion is substantially more efficient when applicable. Furthermore, as discussed by various authors,⁵ it can be applied to multiple types of risk (stochastic volatility, jumps, stochastic interest rates) and more general structures (multiple additive or concatenated latent variables; time-varying jump risk) without sacrificing option pricing tractability. Finally, the affine pricing kernels first developed by Cox et al. (1985) imply risk premia on untraded risks are linear functions of the underlying state variables, creating a method for going between objective and risk-neutral probability measures while retaining convenient affine properties for both measures.

Affine models have two major advantages for econometric work. First, European call (and put) option prices $c(S_t, Y_t, T; \theta, \Phi)$ with maturity T can be computed rapidly conditional upon the observed underlying asset price S_t , the values of relevant underlying state variables Y_t , and the various model parameters θ and risk premia Φ that

⁵ See, e.g., Bates (1996a, 2000), Bakshi et al. (1997), Scott (1997), Bakshi and Madan (2000), Duffie et al. (2000) and Dai and Singleton (2000).

determine the risk-neutral probability measure. Second, the joint characteristic function associated with the joint conditional transition density $p(S_{t+T}, Y_{t+T} | S_t, Y_t, \theta)$ has an analytic solution, implying objective transition densities can also be evaluated via Fourier inversion or other methods.⁶ Consequently, it becomes relatively straightforward to infer Y_t values from observed option prices, and to test whether the observed time series properties of asset and/or option prices *conditional* on those values are consistent with the predicted properties.

Reverse tests based upon the time series properties of asset and/or option prices and their option pricing implications have been handicapped by the difficulty of estimating continuous-time models with latent state variables such as stochastic volatility. However, the time series literature has advanced sufficiently over the last decade that this approach is now feasible. Examples include the Bayesian Monte Carlo Markov Chain approach of Jacquier et al. (1994), and Gallant and Tauchen's EMM/SNP approach.⁷ The latter approach derives moment conditions from a particular auxiliary time series model that captures discrete-time data properties, uses simulated method of moments to estimate the parameters of a continuous-time model consistent with those moment conditions, and uses "reprojection" to generate filtered and smoothed estimates \hat{Y}_t from times series data. As illustrated in the Ahn et al. (2003) paper in this issue, the methodology can be used to estimate model parameters and latent variable values from the joint time series properties of the underlying and of derivatives prices.

1.3. Testing alternative models

So what have we learned from proposing alternatives to the BSM model of geometric Brownian motion with constant volatility and interest rates? In terms of the criteria discussed in connection with Fig. 1, I would say we have learned a great deal about the cross-sectional option pricing implications of the alternative models, and some of the models' qualitative strengths and weaknesses. We are starting to learn the degree to which cross-sectional option pricing patterns are quantitatively consistent with the time series properties of the underlying asset and of option prices, but more work can be done. Finally, even though BSM is not the right model, results from the tests of that model are informative and should be kept in mind.

The early univariate diffusion alternatives to geometric Brownian motion were attempting to capture time-varying volatility and the "leverage" effects of Black (1976) in a simple fashion. The more recent implied binomial trees models, by contrast, were designed to exactly match observed cross-sectional option pricing patterns at any given instant, primarily in order to price over-the-counter exotic options. An exact fit is achieved by a large number of free instantaneous volatility parameters at the various nodes. One of the strengths of the approach—only asset price risk, implying no-arbitrage option pricing—is also one of its major weaknesses. First, as noted by

⁶ The Cox et al. (1985) square-root diffusion implies conditional transition densities Y_t have a non-central χ^2 distribution. Those transition densities are more easily evaluated by a series solution than by Fourier inversion.

⁷ The two approaches are surveyed by Johannes and Polson (2001) and Gallant and Tauchen (2001), respectively.

Bakshi et al. (2000), instantaneous option price evolution is not fully captured by underlying asset price movements, precluding the riskless hedging predicted by these models. Second, these models imply conditional distributions depend upon the level of the underlying asset price, which in turn counterfactually implies nonstationary objective and risk-neutral conditional distributions for stock, stock index, or exchange rate returns. Third, these models use time-dependent instantaneous volatilities $\sigma(S, t)$, which creates difficulties for implementing empirical tests premised on the stationarity assumption that calendar time does *not* matter. This last awkwardness is shared by bond pricing models that rely on time-dependent processes to exactly match an initial term structure of bond yields.

Standard jump models have primarily been used to match volatility smiles and smirks. This they can do fairly well for a single maturity, less well for multiple maturities. As discussed by Bates (2000, p. 205), the standard assumption of independent and identically distributed returns in jump models implies these models converge towards BSM option prices at longer maturities—in contrast to the still-pronounced volatility smiles and smirks at those maturities.⁸ The i.i.d. return structure is also inconsistent with time variation in implicit volatilities. Furthermore, only infrequent large jumps matter for option pricing deviations from Black–Scholes, and it is difficult to estimate such models on time series data.

The major strength of stochastic volatility models is their qualitative consistency with the stochastic but typically mean-reverting evolution of implicit standard deviations. The option pricing implications of standard stochastic volatility models are otherwise fairly close to an ad hoc Black–Scholes model with an updated volatility estimate, when the volatility process is calibrated using parameter values judged plausible given the time series properties of asset volatility or option returns. Consequently, standard stochastic volatility models with plausible parameters cannot easily match observed volatility smiles and smirks, while even unconstrained stochastic volatility models price options better after jumps are added. However, Bakshi et al. (1997) find that adding jumps or stochastic interest rates does not improve the unconstrained stochastic volatility model's assessments of how to hedge option price movements.

Since every simple option pricing model has its weaknesses, hybrid models are needed—and are feasible within the affine model structure. Having jump components addresses moneyness biases, while having stochastic latent variables allows distributions to evolve stochastically over time. Furthermore, Bates (2000) finds that modeling jump intensities as stochastic alleviates the maturity-related option pricing problems of standard jump models. It is interesting that the latest binomial tree models incorporate additional stochastic components as well, as discussed by Skiadopoulos (2000).

The use of hybrid models has been a mainstay of the time series literature since the t -GARCH model of Bollerslev (1987). The latest time series analyses suggest that even these models may be inadequate to describe discernable patterns in volatility evolution.

⁸ See Tompkins (2001) for an examination of observed ISD patterns from stock index, bond, currency, and interest rate options for various countries.

Volatility assessments from intradaily returns⁹ and from high–low ranges¹⁰ indicate longer-lasting volatility shifts than are typically estimated in the ARCH framework, and suggest either a long-memory or a multifactor volatility process. The paper by Chernov et al. (2003) in this issue examines various continuous-time specifications using a 37-year history of daily Dow Jones returns and the EMM/SNP methodology. They conclude that various diffusion-based specifications require at least a two-factor stochastic volatility model to simultaneously summarize volatility evolution and capture the fat-tailed properties of daily returns, while affine specifications require jumps in returns and/or in volatility, and probably both. Furthermore, they find evidence of substantial volatility feedback: volatility is itself more volatile at higher levels.¹¹ While the analysis by Chernov et al. is not confined to affine models, they argue that the comparable fit and analytic convenience may make affine specifications preferable.

1.4. Empirical methodologies

Given a particular model with objective and risk-neutral conditional distributions, researchers face many choices and tradeoffs regarding how to proceed. Which of the various sources of information should be employed in estimating the model: option prices (and if so, which ones), asset returns, and/or option returns? Should one recalibrate the model daily, or impose constant parameters over the full data interval? What tests and diagnostics should be used?

1.4.1. Broad versus narrow option data

Most studies exploit information contained in option prices when testing the consistency of objective and risk-neutral distributions. The issue is whether to use a broad cross-section of options for every day, with multiple strike prices and possibly multiple maturities as well, or to use only as many option prices as there are state variables. Each has strengths and weaknesses. The broad approach emphasizes cross-sectional properties of option prices (e.g., volatility smirks) when inferring parameters. Furthermore, more simultaneous observations are better when there is measurement error in the option data from, e.g., bid-ask spreads, synchronization error with underlying asset prices, or rounding error. Measurement error is the typical justification for using loss functions such as nonlinear ordinary least squares when inferring parameters from option prices.

While there is typically some measurement error in option data, there is not very much.¹² Consequently, implicit parameter “estimation” is actually closer to inversion, and key parameters of the risk-neutral distribution will be identified almost exactly from an appropriate cross-section of option data. Volatility smirks or smiles translate into unique values for the volatility of volatility and asset/volatility correlation in a stochastic

⁹ See Andersen et al. (2001a, b).

¹⁰ See Alizadeh et al. (2002).

¹¹ Jones (2003) in this issue presents parallel evidence that stock market variance estimates based substantially on at-the-money implicit variances also exhibit volatility feedback.

¹² See Bates (2000, Table 4) for a decomposition of option pricing residuals into measurement error and specification error components. The latter is one to four times the size of the former.

volatility model, and into unique jump distribution parameters in a jump-diffusion model. The term structure of implicit volatilities severely constrains the instantaneous volatility estimate and the risk-neutral parameters of volatility drift in a stochastic volatility model.

A recent trend has been to combine cross-sectional data with asset and/or option returns when estimating parameters; see, e.g., Bates (1996a, 2000) or the Chernov (2003) paper in this issue. While necessary for estimating risk premia on volatility or jump risk, and valid under the hypothesis of correct model specification plus measurement error, mixed data sources need to be treated with caution. First, the approach is essentially mixing different statistical techniques: parameter inference, and time series estimation. The option data effectively carry much more weight than asset or option returns for those parameters identifiable from cross-sectional option pricing patterns. Indeed, the estimation procedure is close to a two-step process of first inferring risk-neutral parameters and state variable realizations from option prices, and then using time series estimates for the remaining parameters. Standard error estimates for all parameters will be affected by the apparently high precision of a subset of the parameter vector.

Second, merging time series and cross-sectional implications can make it difficult to identify *why* a particular model is rejected—as it almost invariably is. Merged estimation techniques mix together multiple model implications, and it can be difficult to identify what is driving the parameter estimates and the model rejections. It is useful expositionally to structure tests in the form of cross-sectional derivatives price patterns *versus* time series properties, before building more complicated models premised upon those aspects being compatible.

Pearson and Sun (1994), Pan (2002), and Jones (2003) employ an alternate “narrow” approach of using only as many derivatives prices as state variables. The first two articles assume no measurement error for those specific derivatives prices (bonds and options, respectively); Jones allows for measurement error in option prices. The narrow approach in essence places heavier weight on time series information regarding implicit state variable evolution when estimating parameters. The testable implications are whether the model can correctly price those derivatives *not* used when inferring state variables—i.e., whether it can capture the cross-sectional properties of derivatives prices. This approach therefore also tests the consistency between cross-sectional derivatives prices and time series properties, but differs in how it allocates the data between initial estimation and final tests.

1.4.2. Daily recalibration?

Daily parameter recalibration was often used in early option studies essentially as a form of data description. Implicit volatilities from at-the-money or pooled cross-sectional option data roughly summarized the overall daily level of option prices, while implicit parameters from more general models provided a more accurate summary of cross-sectional option prices. Furthermore, recalibration is sometimes an approximation for a more consistent (and more complicated) model with stochastic latent variables, under the assumption that option prices are not especially sensitive to the neglected randomness. Examples include using BSM implicit volatility as an assessment of conditional (risk-neutral) volatility, and the common practice of updating short-term

interest rates daily without specifying an explicit stochastic interest rate option pricing model.

However, recalibration can hide fundamental flaws in the underlying model. For instance, nonhomogeneous option pricing models such as the constant elasticity of variance and implied binomial trees models imply nonstationary objective and risk-neutral conditional distributions, because of the dependency of instantaneous volatility $\sigma(S, t)$ upon the nonstationary asset price. Since conditional volatility estimates and ISDs do in fact appear to be roughly mean-reverting (apart from the October 1987 structural break in stock index ISD patterns) nonhomogeneous models would fit option prices and asset returns very poorly over time if not repeatedly recalibrated.

The affine class of option pricing models has advanced to the point that recalibration is possibly no longer necessary; or, more precisely, can be confined to resetting key state variables. Parameters are theoretically constant, current values of specific state variables theoretically determine cross-sectional option prices and the objective conditional joint distribution of asset and option returns, and the testable hypotheses are the consistency of the objective and risk-neutral measures. Furthermore, since the underlying state variables are assumed stationary, so are all the data: asset returns, option/asset price ratios, and option returns. The state space structure is consequently well-suited for estimation, especially under the common assumption that option prices are good proxies for the underlying state variables.

1.4.3. Which tests?

All tests are in principle informative; the choice depends to some degree on which issues the researcher wishes to examine. Whether cross-sectional option prices are consistent with the time series properties of the underlying asset returns is probably the most fundamental of tests. The consistency of option prices with option returns as well provides information regarding hypothesized stochastic evolution of latent variables; e.g., the Bates (1996a, 2000) and Bakshi et al. (1997) observation that the implicit volatility of volatility parameter from stochastic volatility models is too high. The joint evolution of options and asset returns is fundamental to hedging issues, and to assessments of volatility and jump risk premia.

Perhaps the one test that does not appear especially informative is short-horizon “out-of-sample” option pricing tests used inter alia by Bakshi et al. (1997) and Dumas et al. (1998), especially when combined with daily parameter recalibration. While it is claimed that such tests control for in-sample overfitting by complicated models with many free parameters, the problem is that daily ISD surfaces are highly serially correlated—as, indeed, is predicted by our models with persistent underlying state variables. The best forecast of tomorrow’s ISD surface is roughly today’s ISD surface, as discussed by Berkowitz (2001). Consequently, any model with multiple free parameters and a good in-sample fit will perform roughly as well out-of-sample at short horizons as any other model—including the ad hoc BSM model with money- and maturity-specific ISDs often used as the “straw man” in these tests.¹³ The out-of-sample

¹³ Christoffersen and Jacobs (2001) point out and illustrate the critical importance of matching up in-sample and out-of-sample loss functions when estimating implicit parameters.

option hedging tests also used by Bakshi et al. and Dumas et al. are probably more informative.

It is not obvious that overfitting is currently a problem for empirical research on option pricing models. In-sample fit over a fixed option data set is not the only diagnostic typically used; models are also tested for a hypothesized ability to describe dynamic asset price and/or option price evolution. As those time series properties typically do not directly affect implicit parameter estimation, there is no evidence of a bias towards model acceptance from in-sample methodology. It has in fact proven difficult to devise a model capable of matching in-sample both cross-sectional option prices and the joint distribution of asset and option returns.

1.4.4. Presentation

Presenting the results of empirical option pricing research intelligibly is one of the biggest challenges facing any researcher. The problem is the abundance of data, and of tests that can be run. Options differ by strike price and by maturity; and it is all too easy to generate tables upon tables of average/percentage/absolute option pricing errors sorted by moneyness and maturity. Some early studies attempted to avoid the tedium of such tables by summarizing the results in regressions with moneyness and maturity as regressors.

At present, there is no unambiguously preferable metric for computing and presenting in-sample fit. Both dollar and percentage option pricing errors exhibit cross-sectional heteroskedasticity, with standard errors falling and exploding, respectively, for options deeper out-of-the-money. Furthermore, dollar errors depend upon the magnitude of the underlying asset price, making it difficult to compare results across options on different assets. Dividing dollar errors by the underlying asset price makes results more comparable, and is also appropriate given that option prices are theoretically nonstationary but option price/asset price *ratios* are stationary under most hypothesized processes. Using ISD errors (estimated minus observed ISD) presents results in a commonly understood metric, but an ISD metric also exhibits severe cross-sectional heteroskedasticity. Furthermore, synchronization error with the underlying asset price can generate apparent arbitrage violations for in-the-money options, precluding ISD computation for unfiltered data.

One of the great advantages of affine models is that they create a model-specific factor structure that compactly summarizes option prices and how option prices evolve. For instance, Bates (2000) estimated a two-factor stochastic volatility/jump-diffusion model, and used the implicit factors to summarize the evolution of implicit distributions over 1988–1993. The reduction in order from a moneyness/maturity grid to a pair of summary state variables has great expositional advantages illustrated in Fig. 5 of that paper. While not as model-independent as the factor analysis used in elsewhere in finance,¹⁴ examining the “cumulant factor loadings” and factor evolutions of alternative affine models may be a useful and intuitive way of framing the debate over appropriate model specification.

¹⁴ See Skiadopoulos et al. (1999) for a factor analysis of ISD surface evolution.

2. Current issues and concerns

At present, empirical options research seems substantially focused on stock index options, and the “volatility smirk” evidence of negatively skewed risk-neutral distributions. This empirical focus is certainly reasonable. Equity is the riskiest portion of most investors’ portfolios, and options have gained acceptance as one method of managing that risk—as is evident in the sharply increased demand for put options on stock indexes since the crash of 1987 (see Fig. 3 of Bates, 2000), and the emergence of instruments such as equity-linked annuities that offer downside-protected equity investments.

As we move towards more complicated models, it is useful to keep in mind some key results from the earlier tests of the BSM model. First, implicit volatilities are *biased* forecasts of subsequent stock market volatility. Regressing some volatility measure upon corresponding volatility assessments inferred from option prices typically yields results such as

$$\frac{365}{T} \sum_{\tau=t+1}^{t+T} (\Delta \ln F_{\tau})^2 = \begin{matrix} 0.0027 + 0.681 \text{ } ISD_t^2 + \varepsilon_{t+T}, & R^2 = 0.33, \\ (0.0033) & (0.161) \end{matrix} \quad (2)$$

$$\sqrt{\frac{365}{T} \sum_{\tau=t+1}^{t+T} (\Delta \ln F_{\tau})^2} = \begin{matrix} 0.0160 + 0.756 \text{ } ISD_t + \varepsilon_{t+T}, & R^2 = 0.45, \\ (0.0142) & (0.102) \end{matrix} \quad (3)$$

with slope coefficients significantly less than 1.¹⁵ Since the intercepts are small, implicit volatilities typically exceeded realized volatilities—by about 2% on average over 1988–1998. And although implicit variances inferred using the geometric Brownian motion are biased estimates of expected average variance under alternate specifications, the bias appears to be small and is of the wrong sign to explain the volatility bias puzzle. This bias has generated substantial speculative profit opportunities from writing options in the post-crash era; see Fleming (1998), Jackwerth (2000), and Coval and Shumway (2001).

Conditional variances of asset returns can be estimated more accurately than other aspects of conditional distributions (e.g. means, higher moments), while risk-neutral conditional variances inferred from near-the-money options are not especially sensitive to model specification. The basic volatility bias results are therefore of key importance to more complicated model specifications that pool information from near-the-money options and asset returns. For instance, affine models imply that objective and risk-neutral conditional variances are maturity-dependent linear transforms of the latent state

¹⁵ The above regression used ISDs from 30-day noon at-the-money options on S&P 500 futures over January 1988 through November 1998, and realized volatility computed from noon futures prices over the lifetime of the options. The data were consequently monthly, and essentially non-overlapping. Heteroskedasticity-consistent standard errors are in parentheses.

variable(s) hypothesized common to asset returns and option prices:

$$\begin{aligned} \text{Var}_t[\ln(F_{t+T}/F_t)] &= a(T) + b(T)Y_t, \\ \text{Var}_t^*[\ln(F_{t+T}/F_t)] &= a^*(T) + b^*(T)Y_t. \end{aligned} \quad (4)$$

The conditional variance bias estimated in Eq. (2) requires a slope coefficient $\hat{\beta} \approx b(T)/b^*(T)$ less than one, or $\hat{b}^* > \hat{b}$. This can be achieved in two ways: by a volatility risk premium that implies slower volatility mean reversion under the risk-neutral than under the objective measure (and therefore greater sensitivity b^* to the initial spot variance Y_t), or by jump risk premia on a time-varying jump component increasing risk-neutral jump variance and total variance relative to objective jump variance.¹⁶ Thus, traditional volatility regressions can substantially determine the risk premium estimates in more complicated stochastic volatility/jump diffusion models, such as Pan (2002).¹⁷

The second major empirical puzzle is the pronounced volatility smirk implicit in options on stock indexes or stock index futures, which implies strongly negatively skewed risk-neutral distributions. As discussed by Bates (2000), the parameter values necessary to match the smirk in S&P 500 futures options appear inconsistent with the time series properties of the stock index futures and of option prices. Moderately infrequent large crashes occurring 0.5–5 times a year (under the risk-neutral distribution) can match the smirk, but it takes substantial risk adjustments to reconcile those parameters with the absence of large crashes in the 14 years since the crash of 1987. A similar point emerges from Jackwerth's (2000) discussion of the substantial crash frequencies needed to eliminate profit opportunities from writing out-of-the-money put options.

Current models attribute to risk premia any discrepancies between objective properties of stock index returns, and risk-neutral probability measures inferred from the volatility smirk. The fundamental theorem of asset pricing implies that provided there are no outright arbitrage opportunities in cross-sectional option prices, there exists *some* pricing kernel that can reconcile the two probability measures. This result has prompted examinations by Ait-Sahalia and Lo (2000), Jackwerth (2000), and Rosenberg and Engle (2002) inter alia regarding general characteristics of the pricing kernel, while Chernov (2003) and Pan (2002) examine pricing kernel and risk premia issues within the class of affine models.

Too much weight has been given to the no-arbitrage foundations of pricing kernels. It is the function of the asset pricing branch of finance to ascertain appropriate risk premia on financial risks—including the jump and stochastic volatility risks spanned by traded options. To blithely attribute divergences between objective and risk-neutral probability

¹⁶ Assuming a time-varying jump intensity, as in Bates (2000) and Pan (2002), is critical. Jump risk premia on a time-invariant jump process such as in Bates (1996a) would affect the intercepts but not the slope coefficients.

¹⁷ Both Pan (2002) and Jones (2003) find that estimated volatility risk premia modify implicit volatility drift sufficiently that the risk-neutral volatility process actually appears locally explosive rather than mean-reverting. As noted by both, such estimates generate a sharply upward-sloped term structure of implicit volatilities, which is incompatible with observed term structures. Pan finds a jump risk premium fits the cross-sectional moneyness and maturity properties of option prices substantially better than a volatility risk premium, but does not otherwise examine the plausibility of the estimated risk premium magnitudes.

measures to the free “risk premium” parameters within an affine model is to abdicate one’s responsibilities as a financial economist. Of course, judging appropriate prices of risks is complicated by the empirical rejections of standard asset pricing models such as the CAPM and consumption CAPM. Cochrane and Saá-Requejo (2000) provide one possible starting point, by deriving option pricing bounds from plausible instantaneous conditional Sharpe ratio constraints on the risk premia on untraded risks.

Furthermore, the “representative agent” construction and calibration associated with implicit pricing kernels and risk premia appear inconsistent with the industrial organization of stock index options markets. Such models imply risk dispersal across all agents in the economy, until all agents are indifferent at the margin to the systematic risks they bear given the prices of those risks. The stock index options markets, by contrast, appear to be functioning as an insurance market. A large clientele of assorted institutional investors predominantly buy stock index options, whereas a relatively small number of specialized option market-making firms predominantly write them. These firms can delta-hedge their positions against small market fluctuations, but are perforce exposed in aggregate to jump and volatility risks.¹⁸ Consequently, market makers limit their exposure to these risks by risk management techniques such as value at risk and worst-case scenario analysis. The limited capitalization of the market makers implies a limited supply of options, and the pricing of the extra risks spanned by options at the shadow prices of investors with relatively high aversions to those risks.

3. Suggestions for future research

As the above discussion suggests, I believe a renewed focus on the explicit financial intermediation of the underlying risks by option market makers is needed. Since such market makers are effectively writing crash insurance, there may be lessons to be learned from the insurance literature. Froot (1999, 2001), for instance, has looked at the financial intermediation of catastrophe risk. From a finance perspective, that market for rare events appears rather inefficient (with premia occasionally nine times expected losses), given the risks are in principle diversifiable. The banking literature has also been concerned with the financial intermediation of risk, and it is possible there are useful models to be found there. Both banking and insurance have greater entry restrictions than are theoretically present in options markets, however.

As part of this research, further scrutiny of market makers’ risk management issues appears warranted. There has, of course, been substantial scrutiny of option hedging under frictions such as discrete-time hedging or transactions costs, but most of that literature has been partial-equilibrium in nature, and premised on BSM assumptions of geometric Brownian motion for the underlying asset price. Devising plausible models of market maker behavior under more general risks and incorporating them into equilibrium models of risk pricing is desirable.

¹⁸ Those “gamma” and “vega” risks can be hedged by buying options, but that just transfers the risk to another market maker.

It is possible that transaction costs may be relevant for some of the observed option pricing puzzles. However, the fact that the deviations between objective and risk-neutral probability measures appear substantially more pronounced for stock index options than for currency options suggests that such costs are not the major factor. Furthermore, those models are difficult to work with even in a geometric Brownian motion framework, and have not yielded especially useful insights. One conclusion is that it appears unproductive to combine transaction costs and the no-arbitrage foundations of BSM. Such approaches implicitly assume option market makers are infinitely risk-averse, and generate excessively broad bid-ask spreads—e.g., the asking price for a call option is the value of the underlying stock. Transaction costs models need to operate in conjunction with a more plausible risk-return calculation on the part of option market makers.

There is, as always, some scope for pricing exotic options under alternative distributional hypotheses. We certainly could use more efficient methods for pricing American options under jumps and/or stochastic volatility; most of the existing literature (see, e.g., [Broadie and Detemple, 1996](#)) has focused on geometric Brownian motion. However, I am generally reluctant to encourage students to explore exotic option pricing issues. First, since those markets are predominantly over-the-counter, data are not readily available for empirical work. Second, there do exist methods for pricing derivatives under jumps ([Amin, 1993](#)) or stochastic volatility ([Wiggins, 1987](#); [Ritchken and Trevor, 1999](#)). They may be inefficient, but they work. More efficient methods or approximations may be of interest to practitioners, but otherwise have a somewhat limited audience.

In terms of alternate model specifications, the Markovian assumptions of standard affine models appear excessively strong. It is assumed that market participants know the instantaneous values Y_t of all relevant underlying state variables as well as the parameter values θ of stochastic processes with certainty, so those values are theoretically inferable from option prices. But whereas the instantaneous conditional volatility may be inferable from high-frequency data in a standard stochastic volatility model, do we really know the time-varying “central tendency” towards which it is headed in a multifactor generalization? Or the instantaneous jump intensity if time-varying—or even if constant? These assumptions place strong testable restrictions on the relationship between time series evidence and cross-sectional option prices; perhaps too strong.

It would be interesting to see equilibrium option pricing models that incorporate uncertainty about the correct model, the parameters, and/or the values of relevant latent variables. My suspicion is that there would still be latent variables inferable from option prices that summarize the market’s best guess of the true distributional characteristics, as suggested by the Bayesian approach of [Dothan and Feldman \(1986\)](#). However, uncertainty might break the perfect-information identity between parameters inferred from option prices and those estimated from time series analysis. The robust decision-making approach of [Anderson et al. \(1998\)](#) and [Cagetti et al. \(2002\)](#) is an alternate approach to specification risk that could generate some interesting equilibrium option pricing implications.

Finally, it is safe to predict that numerical methods will be used increasingly in future empirical options research. Computers continue to get faster at a remarkable pace, and

the need for analytically tractable models diminishes as alternative brute-force methods become cheaper. An expositional barrier will still remain, however: explaining the methodology and the results given the many methods discussed above whereby option pricing models can be tested. Implementing any of the above suggestions and presenting them intelligibly is left as an exercise to the readers.

References

- Ahn, D., Dittmar, R., Gao, B., Gallant, A.R., 2003. Purebred or hybrid? Reproducing the volatility in term structure dynamics. *Journal of Econometrics*, this issue.
- Aït-Sahalia, Y., Lo, A.W., 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9–51.
- Alizadeh, S., Brandt, M.W., Diebold, F.X., 2002. Range-based estimation of stochastic volatility models. *Journal of Finance* 57, 1047–1091.
- Amin, K.I., 1993. Jump diffusion option valuation in discrete time. *Journal of Finance* 48, 1833–1863.
- Anderson, E.W., Hansen, L.P., Sargent, T.J., 1998. Risk and robustness in general equilibrium. University of Chicago working paper.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2001a. The distribution of stock return volatility. *Journal of Financial Economics* 61, 43–76.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2001b. Modeling and forecasting realized volatility. National Bureau of Economic Research working paper 8160.
- Bakshi, G., Madan, D.B., 2000. Spanning and derivative-security valuation. *Journal of Financial Economics* 55, 205–238.
- Bakshi, G., Cao, C., Chen, Z., 1997. Empirical performance of alternative option pricing models. *Journal of Finance* 52, 2003–2049.
- Bakshi, G., Cao, C., Chen, Z., 2000. Do call prices and the underlying stock always move in the same direction? *Review of Financial Studies* 13, 549–584.
- Bates, D.S., 1996a. Jumps and stochastic volatility: exchange rate processes implicit in PHLX deutsche mark options. *Review of Financial Studies* 9, 69–107.
- Bates, D.S., 1996b. Testing option pricing models. In: Maddala, G.S., Rao, R. (Eds.), *Handbook of Statistics*, Vol. 14. Elsevier, Amsterdam, pp. 567–611.
- Bates, D.S., 2000. Post-'87 Crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94, 181–238.
- Berkowitz, J., 2001. Getting the right option price with the wrong model. University of California (Irvine) working paper.
- Black, F., 1976. Studies of stock price volatility changes. *Proceedings of the 1976 Meetings of the American Statistical Association*, pp. 177–181.
- Black, F., Scholes, M., 1972. The valuation of option contracts in a test of market efficiency. *Journal of Finance* 27, 399–417.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–659.
- Bollerslev, T., 1987. A conditional heteroskedastic time series model for speculative prices and rates of return. *Review of Economics and Statistics* 69, 542–547.
- Broadie, M.N., Detemple, J., 1996. American option valuation: new bounds, approximations, and a comparison of existing bounds. *Review of Financial Studies* 9, 1211–1250.
- Cagetti, M., Hansen, L.P., Sargent, T.J., Williams, N., 2002. Robustness and pricing with uncertain growth. *Review of Financial Studies* 15, 363–404.
- Campa, J.M., Chang, P.H.K., 1995. Testing the expectations hypothesis on the term structure of implied volatilities in foreign exchange options. *Journal of Finance* 50, 529–547.
- Campbell, J.Y., 1995. Some lessons from the yield curve. *Journal of Economic Perspectives* 9, 129–152.
- Chernov, M., 2003. Empirical reverse engineering of the pricing kernel. *Journal of Econometrics*, this issue.

- Chernov, M., Gallant, A.R., Ghysels, E., Tauchen, G., 2003. Alternative models for stock price dynamics. *Journal of Econometrics*, this issue.
- Christoffersen, P., Jacobs, K., 2001. The importance of the loss function in option pricing. McGill working paper.
- Cochrane, J.H., Saá-Quejeo, J., 2000. Beyond arbitrage: “good-deal” asset price bounds in incomplete markets. *Journal of Political Economy* 108, 79–119.
- Coval, J.D., Shumway, T., 2001. Expected option returns. *Journal of Finance* 56, 983–1009.
- Cox, J.C., Ross, S.A., 1976. The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145–166.
- Cox, J.C., Rubinstein, M., 1985. *Options Markets*. Prentice-Hall, Englewood Cliffs, NJ.
- Cox, J.C., Ingersoll Jr., J.E., Ross, S.A., 1985. A theory of the term structure of interest rates. *Econometrica* 53, 385–407.
- Dai, Q., Singleton, K.J., 2000. Specification analysis of affine term structure models. *Journal of Finance* 55, 1943–1978.
- Diz, F., Finucane, T.J., 1993. Do the options markets really overreact? *Journal of Futures Markets* 13, 298–312.
- Dothan, M.U., Feldman, D., 1986. A theory of asset prices and the term structure of interest rates in a partially observable economy. *Journal of Finance* 41, 369–382.
- Duffie, D., Pan, J., Singleton, K.J., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Dumas, B., Fleming, J., Whaley, R.E., 1998. Implied volatility functions: empirical tests. *Journal of Finance* 53, 2059–2106.
- Fleming, J., 1998. The quality of market volatility forecasts implied by S&P 100 index option prices. *Journal of Empirical Finance* 5, 317–345.
- Froot, K.A., 1999. The evolving market for catastrophic event risk. *Risk Management and Insurance Review* 2, 1–28.
- Froot, K.A., 2001. The market for catastrophe risk: a clinical examination. *Journal of Financial Economics* 60, 529–571.
- Gallant, A.R., Tauchen, G., 2001. Simulated score methods and indirect inference for continuous-time models. University of North Carolina working paper.
- Garcia, R., Ghysels, E., Renault, E., 2001. The econometrics of option pricing. University of Montréal working paper.
- Geske, R., 1979. The valuation of compound options. *Journal of Financial Economics* 7, 63–81.
- Heath, D., Jarrow, R.A., Morton, A.J., 1992. Bond pricing and the term structure of interest rates: a new methodology for contingent claims valuation. *Econometrica* 60, 77–105.
- Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6, 327–344.
- Hull, J., White, A., 1987. The pricing of options on assets with stochastic volatility. *Journal of Finance* 42, 281–300.
- Jackwerth, J.C., 2000. Recovering risk aversion from option prices and realized returns. *Review of Financial Studies* 13, 433–451.
- Jacquier, E., Polson, N.G., Rossi, P.E., 1994. Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* 12, 1–19.
- Jagannathan, R., Kaplin, A., Sun, S.G., 2003. An evaluation of multi-factor CIR models using LIBOR, swap rates, and cap and swaption prices. *Journal of Econometrics*, this issue.
- Johannes, M., Polson, N.G., 2001. MCMC Methods for financial econometrics. Columbia University working paper.
- Johnson, H., Shanno, D., 1987. Option pricing when the variance is changing. *Journal of Financial and Quantitative Analysis* 22, 143–151.
- Jones, C.S., 2003. The dynamics of stochastic volatility. *Journal of Econometrics*, this issue.
- Merton, R.C., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.
- Naik, V., 1993. Option valuation and hedging strategies with jumps in the volatility of asset returns. *Journal of Finance* 48, 1969–1984.

- Pan, J., 2002. The jump-risk premia implicit in options: evidence from an integrated time-series study. *Journal of Financial Economics* 63, 3–50.
- Pearson, N.D., Sun, T., 1994. Exploiting the conditional density in estimating the term structure: an application to the Cox, Ingersoll, and Ross model. *Journal of Finance* 49, 1279–1304.
- Ritchken, P.H., Trevor, R., 1999. Pricing options under generalized GARCH and stochastic volatility processes. *Journal of Finance* 54, 377–402.
- Rosenberg, J., Engle, R.F., 2002. Empirical pricing kernels. *Journal of Financial Economics* 64, 341–372.
- Rubinstein, M., 1983. Displaced diffusion option pricing. *Journal of Finance* 38, 213–217.
- Scott, L.O., 1987. Option pricing when the variance changes randomly: theory, estimation, and an application. *Journal of Financial and Quantitative Analysis* 22, 419–438.
- Scott, L.O., 1997. Pricing stock options in a jump-diffusion model with stochastic volatility and interest rates: applications of Fourier inversion methods. *Mathematical Finance* 7, 413–426.
- Skiadopoulos, G., 2000. Volatility smile consistent option models: a survey. *International Journal of Theoretical and Applied Finance* 4, 403–437.
- Skiadopoulos, G., Hodges, S., Clewlow, L., 1999. The dynamics of the S&P 500 implied volatility surface. *Review of Derivatives Research* 3, 263–282.
- Stein, J.C., 1989. Overreactions in the options market. *Journal of Finance* 44, 1011–1023.
- Stein, E.M., Stein, J.C., 1991. Stock price distributions with stochastic volatility: an analytic approach. *Review of Financial Studies* 4, 727–752.
- Tompkins, R.G., 2001. Implied volatility surfaces: uncovering regularities for options on financial futures. *The European Journal of Finance* 7, 198–230.
- Wiggins, J.B., 1987. Option values under stochastic volatility: theory and empirical estimates. *Journal of Financial Economics* 19, 351–377.