Empirical Option Pricing: A Retrospection

David S. Bates
University of Iowa and the
National Bureau of Economic Research

January 2001

Henry B. Tippie College of Business
University of Iowa
Iowa City, IA 52242-1000

Tel: (319) 353-2288
email: david-bates@uiowa.edu

On October 20-21, 2000, a conference was held at Duke University on the subject of risk-neutral and objective probability measures. Some of the fifteen papers examined recent econometric techniques for estimating conditional distributions, others discussed "risk-neutral" models consistent with observed bond and/or options prices, while most drew inferences from the interplay of time series analysis and derivatives prices.[1] The papers demonstrated the considerable advances of the last decade with regard to developing tractable models for pricing bonds and options, and with regard to developing numerically intensive econometric techniques compatible with the continuous-time framework commonly employed in derivatives research.

It is my intent in this article to reflect on what we have achieved, what are the issues in empirical option pricing research, and where we may be heading. I begin in Section I 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 II, I discuss some major remaining puzzles in option pricing. I conclude in Section III with some suggestions for possible research topics for those interested in working in the area.

---

[1]Seven papers examined option pricing issues: Aït-Sahalia and Duarte (2000), Bakshi, Kapadia and Madan (2000), Chernov (2000), Garcia, Luger and Renault (2000), Rosenberg and Engle (2000), Jones (2000) and Pan (2000). Three papers examined bond pricing issues: Ahn, Dittmar, Gao, and Gallant (2000), Dai and Singleton (2000b), and Duffee and Stanton (2000). Jangannathan, Kaplin and Sun (2000) examined the interplay between bond and option pricing models. Chernov, Gallant, Ghysels, and Tauchen (2000) and Zhou (1999) examined time series estimation, while Cagetti, Hansen, Sargent and Williams (2000) and Stutzer (2000) advocated robust and entropy-based estimation techniques, respectively.

I. Empirical derivatives research

I.A. 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 derivative 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

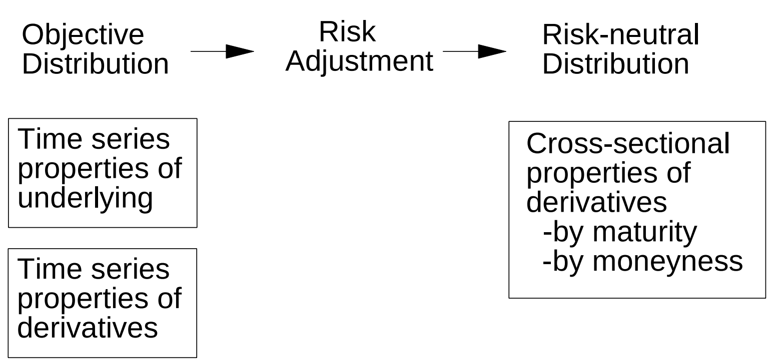

Figure 1. Objective and risk-neutral distributions.

derivatives returns *conditional* on information inferred from the cross-section of derivatives prices. The latter approach has been somewhat 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.

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 in Campbell (1995) examine whether yield spreads between *n*-period 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 interest rate models of Pearson and Sun (1994), Ahn *et al* (2000), and Jagannathan *et al* (2000), and the Heath, Jarrow and Morton (1992) method of modelling stochastic forward rate curve evolution consistently with an initial term structure of forward rates.

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

1) *Time series properties of underlying asset*. 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 (ISD's) 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 ISD's should also be constant across time -- a hypothesis clearly rejected by time series graphs of ISD's.

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 ISD's. Typical results summarized in Bates (1996b, pp. 593-5) are that ISD's 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 ISD's predicts subsequent evolution in short-term ISD's.

6) *Joint distribution of asset returns and option returns.* Bakshi, Cao, and Chen (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,

$$\frac{\Delta \tilde{C}}{C} = \left.\frac{S\, C_S^{BS}}{C^{BS}}\right|_{\hat{\sigma}} \frac{\Delta \tilde{S}}{S} + O(\Delta t) \tag{1}$$

with the leverage estimate based typically on an ISD inferred from recent option prices.[2] Profitability tests examine the average returns from self-financing portfolios of options and accompanying asset hedge positions, while 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 Black-Scholes-Merton 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 ISD's initially generated a debate on how to construct weighted ISD estimates for use in examining the conditional distribution of asset and/or options returns,[3] rather than being taken as evidence against the model. Similarly, the use of contemporaneous or lagged ISD's when forecasting volatility and computing hedge ratios was inconsistent with the BSM prediction that ISD's 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.

I.B. Alternative models

The empirical deficiencies of the Black-Scholes-Merton model have prompted research along three dimensions. The *univariate diffusion* models maintained the diffusion and no-arbitrage 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, Pan, and Singleton, 1998). *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 in

[2] The original Black-Scholes (1972) test of their model used time-series based volatility estimates, which falls within the sixth test classification above.

[3] See Bates (1996b, pp. 587-9).

Skiadopoulos (1999), 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-Scholes (1973) paper, and continued in the 1980's with the stochastic volatility models, actually testing those alternatives using options data was handicapped by several factors. First, options data weren't readily available until after the advent of options trading on centralized exchanges. Second, the new models were perforce more complicated. Pricing options was a substantial computational burden for the (mainframe) 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 1990's 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 (1981) and Heston (1983). 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,[4] 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, Ingersoll, and Ross (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.

[4]See, e.g., Bates (1996a, 2000), Scott (1997), Duffie, Pan and Singleton (1998), Dai and Singleton (2000a), and Bakshi and Madan (2000).

Affine models have two major advantages for econometric work. First, European call (and put) option prices $c(S_t, \boldsymbol{Y}_t,\ T;\ \theta,\ \Phi)$ with maturity T can be computed rapidly conditional upon the observed underlying asset price S_t, the values of relevant underlying state variables $\boldsymbol{Y}_t$, and the various model parameters θ and risk premia Φ that determine the risk-neutral probability measure. Second, the joint characteristic function associated with the joint conditional transition density $p(S_{t+T}, \boldsymbol{Y}_{t+T} \mid S_t, \boldsymbol{Y}_t, \theta)$ has an analytic solution, implying objective transition densities can also be evaluated via Fourier inversion or other methods.[5] Consequently, it becomes relatively straightforward to infer $\boldsymbol{Y}_t$ 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 Jaquier, Polson and Rossi (1994), and Gallant and Tauchen's EMM/SNP approach. 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 $\hat{\boldsymbol{Y}}_t$ from times series data. As illustrated in Ahn *et al* (2000), 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.

[5]The Cox, Ingersoll, and Ross (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.

I.C. Testing alternative models

So what have we learned from proposing alternatives to the Black-Scholes-Merton model of geometric Brownian motion with constant volatility and interest rates? In terms of the criteria discussed in connection with Figure 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 Black-Scholes-Merton 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; a goal 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 in Bakshi, Cao and Chen (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 introduce a time dimension into option pricing in a fashion inconvenient for empirical work. This last property is shared by bond pricing models that rely on time-dependent processes to 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 in 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.[6] 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 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 cannot easily match observed volatility smiles and smirks. Bakshi, Cao and Chen (1997) find that stochastic volatility models do as well as jump and stochastic interest rate models with regard to an option hedging criterion.

Since every simple 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 modelling 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 in Skiadopoulos (1999).

Using option pricing models with stochastic latent variables reduces the *ad hoc* structure of many early option pricing tests. The early models were sufficiently inconsistent with key empirical findings that periodic recalibration was common: infer key parameter(s) from option prices, test a particular time series prediction for the next period, and repeat the inference and testing for subsequent periods. Latent variable models by contrast create a state space structure consistent with the hypothesized data generating process. Parameters are constant, current values of specific state

---

[6]See Tompkins (2000) for an examination of observed ISD patterns from stock index, bond, currency, and interest rate options for various countries.

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.[7]

Of course, these generalized models in conjunction with general continuous-time estimation techniques for multiple series can get quite complicated. Since an affine model has cross-sectional *and* time series implications, it is now possible to combine multiple implications in the estimation methodology in an internally consistent fashion. This approach has been used in testing bond pricing models since Pearson and Sun (1994), and in option pricing research since Bates (1996a). Chernov (2000) provides a recent illustration of how affine models can exploit multiple information sources.

However, econometric techniques that pool cross-sectional and time series information should be treated with caution. While valid under the null hypothesis of correct model specification, this approach has two disadvantages. First, the approach is mixing fundamentally different statistical techniques: time series estimation, and parameter inference from derivatives prices. While the statistical properties of the former approach are well-understood, the latter essentially involves inversion rather than estimation. And while "measurement error" gives parameter inference an aura of statistical respectability, hypothesizing measurement error can obscure a fundamental incompatibility between the time series properties and cross-sectional derivatives prices.

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.

---

[7]One frequently voiced concern is that the "out-of-sample" properties of the early tests are lost when parameter invariance is imposed over, e.g., a 10-year horizon for both objective and risk-neutral probability measures. However, one can run subsample stability tests such as the Chow test to check for parameter drift.

II. Current issues and concerns

At present, empirical options research seems substantially concentrated 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 Figure 1 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 Black-Scholes-Merton model. First, implicit volatilities are *biased* forecasts of subsequent stock market volatility. Regressing some volatility measure upon corresponding volatility assessments inferred from options prices typically yields results such as

$$\frac{365}{T}\sum_{\tau=t+1}^{T}(\Delta \ln F_\tau)^2 = \underset{(.0028)}{.0027} + \underset{(.086)}{.681}\, ISD_t^2 + \varepsilon_{t+T}, \quad R^2 = .33 \qquad \textbf{(2)}$$

with slope coefficients significantly less than 1.[8] Since the intercept is small, implicit volatilities typically exceeded realized volatilities on average over 1988-98. 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 at-the-money options in the post-crash era; see Fleming (1998) and Jackwerth (2000).

Pan (2000) illustrates the relevance of keeping the volatility bias results in mind when working with more complicated affine models. She uses stock index returns and the evolution of

---

[8]The above regression used ISD's from 30-day noon ATM options 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. Using standard deviations instead of variances generates a higher slope coefficient of 0.756 (standard error of .073), and an R^2 of .45. The average annualized divergence between implicit and realized standard deviations over 1988-98 was 2%.

a latent variable inferred from ATM option prices to estimate the actual and risk-neutral parameters under the Heston (1993) stochastic volatility model, and under the Bates (2000) stochastic volatility/jump-diffusion model with time-varying jump risk. Affine models imply that objective and risk-neutral conditional variances are maturity-dependent linear transforms of the latent state variable hypothesized common to asset returns and option prices:

$$\begin{aligned} Var_t[\ln(F_{t+T}/F_t)] &= a(T) + b(T)Y_t \\ Var_t^*[\ln(F_{t+T}/F_t)] &= a^*(T) + b^*(T)Y_t . \end{aligned} \tag{3}$$

The volatility bias results of equation (2) require $\hat{b}^* > \hat{b}$. 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 value Y_t), or by jump risk premia increasing risk-neutral jump variance and total variance relative to objective jump variance.[9] Pan finds the latter approach fits the cross-sectional moneyness and maturity properties of option prices substantially better than the former approach, but does not otherwise examine the plausibility of the estimated risk premia.

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 in 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 thirteen 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

[9]Assuming a time-varying jump intensity is critical. The jump risk premia on a time-invariant jump process such as in Bates (1996a) would affect a^* but not b^*.

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 Aït-Sahalia and Lo (1998), Jackwerth (2000), and Rosenberg and Engle (2000) *inter alia* regarding general characteristics of the pricing kernel, while Chernov (2000) and Pan (2000) 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 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 (1996) provide one possible starting point, by deriving option pricing bounds from plausible 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 marketmaking 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.[10] Consequently, marketmakers 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 marketmakers implies a limited supply of options, and the pricing of the

[10]Those "gamma" and "vega" risks can hedged by buying options, but that just transfers the risk to another marketmaker.

extra risks spanned by options at the shadow prices of investors with relatively high aversions to those risks.

III. 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 marketmakers is needed. Since such marketmakers are effectively writing crash insurance, there may be lessons to be learned from the insurance literature. Froot (1999a,b), 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 originally 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 marketmakers' 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 Black-Scholes-Merton assumptions of geometric Brownian motion for the underlying asset price. Devising plausible models of marketmaker behavior under more general risks and incorporating them into equilibrium models of risk pricing is desirable.

It is possible that transactions 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 transactions 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 is unproductive to combine transactions costs and the no-arbitrage foundations of Black-Scholes-Merton. Such approaches implicitly assume option marketmakers 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. Transactions costs models need to operate in conjunction with a more general risk-return calculation on the part of option marketmakers.

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 focussed 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. Coming up with more efficient methods or approximations may be of interest to practitioners, but otherwise has 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 $\boldsymbol{Y}_t$ 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 options 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 could break the perfect-information identity between parameters inferred from option prices and those estimated from time series analysis. The robust decision-making approach of Andersen, Hansen, and Sargent (1998) and Cagetti *et al* (2000) 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, Dong-Hyun, Robert Dittmar, Bin Gao, and A. Ronald Gallant (2000). "Purebred or Hybrid? Reproducing the Volatility in Term Structure Dynamics." University of North Carolina working paper, September.

Aït-Sahalia, Yacine and Jefferson Duarte (2000). "Nonparametric Option Pricing under Shape Restrictions." Princeton working paper, October.

Aït-Sahalia, Yacine and Andrew W. Lo (1998). "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." *Journal of Finance* **53**, 499-547.

Amin, Kaushik I. (1993). "Jump Diffusion Option Valuation in Discrete Time." *Journal of Finance* **48**, 1833-1863.

Anderson, Evan W., Lars P. Hansen, and Thomas J. Sargent (1998). "Risk and Robustness in General Equilibrium." University of Chicago working paper, March.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen (2000). "Do Call Prices and the Underlying Stock Always Move in the Same Direction?" *Review of Financial Studies* **13**.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip B. Madan (2000). "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." University of Maryland working paper, June.

Bakshi, Gurdip and Dilip B. Madan (2000). "Spanning and Derivative-Security Valuation." *Journal of Financial Economics* **55**, 205-238.

Bates, David S. (1996). "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in PHLX Deutschemark Options." *Review of Financial Studies* **9**, 69-107.

Bates, David S. (1996). "Testing Option Pricing Models," in G.S. Maddala, C.R. Rao, eds., *Statistical Methods in Finance,* Amsterdam: Elsevier, 567-611.

Bates, David S. (2000). "Post-'87 Crash Fears in the S&P 500 Futures Option Market." *Journal of Econometrics* **94**, 181-238.

Black, Fischer (1976). "Studies of Stock Price Volatility Changes." *Proceedings of the 1976 Meetings of the American Statistical Association* , 177-181.

Black, Fischer and Myron Scholes (1972). "The Valuation of Option Contracts in a Test of Market Efficiency." *Journal of Finance* **27**, 399-417.

Black, Fischer and Myron Scholes (1973). "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* **81**, 637-659.

Broadie, Mark N. and Jerome Detemple (1996). "American Option Valuation: New Bounds, Approximations, and a Comparison of Existing Bounds." *Review of Financial Studies* **9**, 1211-1250.

Cagetti, Marco, Lars P. Hansen, Thomas J. Sargent, and Noah Williams (2000). "Robustness and Pricing with Uncertain Growth." University of Chicago working paper, October.

Campa, Jose M. and P. H. Kevin Chang (1995). "Testing the Expectations Hypothesis on the Term Structure of Implied Volatilities in Foreign Exchange Options." *Journal of Finance* **50**, 529-547.

Campbell, John Y. (1995). "Some Lessons from the Yield Curve." *Journal of Economic Perspectives* **9**, 129-152.

Chernov, Mikhail (2000). "A Case of Empirical Reverse Engineering: Estimation of the Pricing Kernel." Columbia University working paper, August.

Chernov, Mikhail, A. Ronald Gallant, Eric Ghysels, and George Tauchen (2000). "Alternative Models for Stock Price Dynamics." University of North Carolina working paper, October.

Cochrane, John H. and Jesús Saá-Requejo (1996). "Beyond Arbitrage: "Good-Deal" Asset Price Bounds in Incomplete Markets." National Bureau of Economic Research working paper 5489.

Cox, John C., Jonathan E. Ingersoll,Jr., and Stephen A. Ross (1985). "A Theory of the Term Structure of Interest Rates." *Econometrica* **53**, 385-407.

Cox, John C. and Stephen A. Ross (1976). "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics* **3**, 145-166.

Cox, John C. and Mark Rubinstein (1985). *Options Markets*, Englewood Cliffs, New Jersey: Prentice-Hall.

Dai, Qiang and Kenneth J. Singleton (2000a). "Specification Analysis of Affine Term Structure Models." *Journal of Finance* **55**, 1943-1978.

Dai, Qiang and Kenneth J. Singleton (2000b). "Expectation Puzzles, Time-varying Risk Premia, and Dynamic Models of the Term Structure." New York University working paper, September.

Diz, Fernando and Thomas J. Finucane (1993). "Do the Options Markets Really Overreact?" *Journal of Futures Markets* **13**, 298-312.

Duffee, Gregory R. and Richard H. Stanton (2000). "EMM Estimation of Affine and Nonaffine Term Structure Models." University of California (Berkeley) working paper, September.

Duffie, Darrell, Jun Pan, and Kenneth J. Singleton (1998). "Transform Analysis and Option Pricing for Affine Jump-Diffusions." Stanford University working paper, August.

Fleming, Jeff (1998). "The Quality of Market Volatility Forecasts Implied by S&P 100 Index Option Prices." *Journal of Empirical Finance* **5**, 317-345.

Froot, Kenneth A. (1999). "The Market for Catastrophe Risk: A Clinical Examination." National Bureau of Economic Research working paper 7286.

Froot, Kenneth A. (1999). "The Evolving Market for Catastrophic Event Risk." National Bureau of Economic Research working paper 7287.

Garcia, René, Richard Luger, and Eric Renault (2000). "Empirical Assessment of an Intertemporal Option Pricing Model with Latent Variables." University of Montréal working paper, September.

Geske, Robert (1979). "The Valuation of Compound Options." *Journal of Financial Economics* **7**, 63-81.

Heath, David, Robert A. Jarrow, and Andrew J. Morton (1992). "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation." *Econometrica* **60**, 77-105.

Heston, Steve 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, John and Alan White (1987). "The Pricing of Options on Assets with Stochastic Volatility." *Journal of Finance* **42**, 281-300.

Jackwerth, Jens Carsten (2000). "Recovering Risk Aversion from Option Prices and Realized Returns." *Review of Financial Studies* **13**, 433-451.

Jacquier, Eric, Nicholas G. Polson, and Peter E. Rossi (1994). "Bayesian Analysis of Stochastic Volatility Models." *Journal of Business and Economic Statistics* **12**, 1-19.

Jagannathan, Ravi, Andrew Kaplin, and Steve Guoqiang Sun (2000). "An Evaluation of Multi-Factor CIR Models Using LIBOR, Swap Rates, and Cap and Swaption Prices." Northwestern University working paper, September.

Johnson, Herb and David Shanno (1987). "Option Pricing When the Variance Is Changing." *Journal of Financial and Quantitative Analysis* **22**, 143-151.

Jones, Christopher S. (2000). "The Dynamics of Stochastic Volatility." University of Rochester working paper, May.

Merton, Robert C. (1976). "Option Pricing When Underlying Stock Returns are Discontinuous." *Journal of Financial Economics* **3**, 125-144.

Naik, Vasanttilak (1993). "Option Valuation and Hedging Strategies with Jumps in the Volatility of Asset Returns." *Journal of Finance* **48**, 1969-1984.

Pan, Jun (2000). "Integrated Time-Series Analysis of Spot and Option Prices." Stanford University working paper, August.

Pearson, Neil D. and Tong-Sheng Sun (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, Peter H. and Rob Trevor (1999). "Pricing Options under Generalized GARCH and Stochastic Volatility Processes." *Journal of Finance* **54**, 377-402.

Rosenberg, Joshua and Robert F. Engle (2000). "Empirical Pricing Kernels." New York University working paper, July.

Rubinstein, Mark (1983). "Displaced Diffusion Option Pricing." *Journal of Finance* **38**, 213-217.

Scott, Louis O. (1987). "Option Pricing When the Variance Changes Randomly: Theory, Estimation, and an Application." *Journal of Financial and Quantitative Analysis* **22**, 419-438.

Scott, Louis O. (1997). "Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Interest Rates: Applications of Fourier Inversion Methods." *Mathematical Finance* **7**, 413-426.

Skiadopoulos, George (1999). "Volatility Smile Consistent Option Models: A Survey." University of Warwick working paper, November.

Stein, Elias M. and Jeremy C. Stein (1991). "Stock Price Distributions with Stochastic Volatility: An Analytic Approach." *Review of Financial Studies* **4:** 727-752.

Stein, Jeremy C. (1989). "Overreactions in the Options Market." *Journal of Finance* **44**, 1011-1023.

Stutzer, Michael (2000). "A Large Deviations Approach to Portfolio Analysis: The Growth - Security Efficiency Frontier." University of Iowa working paper.

Tompkins, Robert G. (2000). "Implied Volatility Surfaces: Uncovering Regularities for Options on Financial Futures." Vienna University of Technology working paper, July.

Wiggins, James B. (1987). "Option Values Under Stochastic Volatility: Theory and Empirical Estimates." *Journal of Financial Economics* **19**, 351-377.

Zhou, Hao (1999). "Jump-Diffusion Term Structure and Itô Conditional Moment Generator." Duke University working paper, October.