

Quantitative Strategies Research Notes

July 1999

Strike-Adjusted Spread:

*A New Metric For Estimating
The Value Of Equity Options*

Joseph Zou
Emanuel Derman



QUANTITATIVE STRATEGIES RESEARCH NOTES

Copyright 1999 Goldman, Sachs & Co. All rights reserved.

This material is for your private information, and we are not soliciting any action based upon it. This report is not to be construed as an offer to sell or the solicitation of an offer to buy any security in any jurisdiction where such an offer or solicitation would be illegal. Certain transactions, including those involving futures, options and high yield securities, give rise to substantial risk and are not suitable for all investors. Opinions expressed are our present opinions only. The material is based upon information that we consider reliable, but we do not represent that it is accurate or complete, and it should not be relied upon as such. We, our affiliates, or persons involved in the preparation or issuance of this material, may from time to time have long or short positions and buy or sell securities, futures or options identical with or related to those mentioned herein.

This material has been issued by Goldman, Sachs & Co. and/or one of its affiliates and has been approved by Goldman Sachs International, regulated by The Securities and Futures Authority, in connection with its distribution in the United Kingdom and by Goldman Sachs Canada in connection with its distribution in Canada. This material is distributed in Hong Kong by Goldman Sachs (Asia) L.L.C., and in Japan by Goldman Sachs (Japan) Ltd. This material is not for distribution to private customers, as defined by the rules of The Securities and Futures Authority in the United Kingdom, and any investments including any convertible bonds or derivatives mentioned in this material will not be made available by us to any such private customer. Neither Goldman, Sachs & Co. nor its representative in Seoul, Korea is licensed to engage in securities business in the Republic of Korea. Goldman Sachs International or its affiliates may have acted upon or used this research prior to or immediately following its publication. Foreign currency denominated securities are subject to fluctuations in exchange rates that could have an adverse effect on the value or price of or income derived from the investment. Further information on any of the securities mentioned in this material may be obtained upon request and for this purpose persons in Italy should contact Goldman Sachs S.I.M. S.p.A. in Milan, or at its London branch office at 133 Fleet Street, and persons in Hong Kong should contact Goldman Sachs Asia L.L.C. at 3 Garden Road. Unless governing law permits otherwise, you must contact a Goldman Sachs entity in your home jurisdiction if you want to use our services in effecting a transaction in the securities mentioned in this material.

Note: Options are not suitable for all investors. Please ensure that you have read and understood the current options disclosure document before entering into any options transactions.

SUMMARY

Investors in equity options experience two problems that compound each other. In contrast to fixed-income and currency markets, there are thousands of underlyers and tens of thousands of options, and each underlyer can have a potentially large volatility skew. How can an options investor gauge which option provides the best relative value?

In this paper, we make use of a method for estimating the fair volatility smile of any equity underlyer from information embedded in the time series of that underlyer's historical returns. We can then compute the relative richness or cheapness of any particular strike and expiration by examining the option's *Strike-Adjusted Spread*, or *SAS*, the difference between its market implied volatility and its estimated historically-fair volatility.

We obtain fair volatility smiles by estimating the appropriate risk-neutral distribution for valuing options on any equity underlyer from that underlyer's historical returns. The distribution includes the effect of both past price jumps and past shifts in realized volatility. Using this distribution, we can estimate the fair volatility skews for illiquid or thinly-traded single-stock and basket options. We can also forecast changes in the skew from changes in a single options price.

Joseph Zou (212) 902-9794
Emanuel Derman (212) 902-0129

Editorial: Barbara Dunn

Table of Contents

THE RICHNESS AND CHEAPNESS OF OPTIONS.....	1
Current vs. Past Implied Volatilities	1
Implied vs. Historical Volatilities	1
Strike-Adjusted Spread (SAS)	1
ATM Strike-Adjusted Spread	3
OPTIONS PRICES AND IMPLIED DISTRIBUTIONS.....	5
STOCK RETURNS AND HISTORICAL DISTRIBUTIONS.....	6
MAXIMAL UNCERTAINTY AND MARKET EQUILIBRIUM.....	8
Entropy as a Measure of Uncertainty	8
The Risk-Neutralized Historical Distribution	9
APPLICATIONS OF THE RISK-NEUTRALIZED DISTRIBUTION.....	11
Is The Index Implied Volatility Skew Fair?	11
Strike-Adjusted Spread As A Measure Of Options Value	13
Valuing Options on Baskets of Stocks	15
Forecasting The Shape of The Skew	
From A Change In A Single Option Price	17
End-of-Day Mark To Market	19
Filling Gaps In Investors' Market Views	19
CONCLUDING REMARKS	20
APPENDIX A: INFORMATION AND ENTROPY	21
APPENDIX B: DETERMINING THE RISK-NEUTRAL	
DISTRIBUTION FROM HISTORICAL RETURNS	24
APPENDIX C: A DERIVATIVE ASSET ALLOCATION MODEL	
AND THE EQUILIBRIUM RISK-NEUTRAL DENSITY	25
REFERENCES.....	29

**THE RICHNESS AND
CHEAPNESS OF
OPTIONS**

The equities world is a mass of data. Surrounded by fluctuating share prices, dividend yields, earnings forecasts, P/E ratios, and hosts of more sophisticated measures, analysts, investors are in need of some gauge or metric with which to compare the relative attractiveness of different stocks. Into the breach, in newsletters, books and websites, step countless economists, technical analysts, fundamental analysts, chartists, wave theorists, alpha-maximizers and other optimists, hoping to impose order and rationality, to tell you what to buy and sell.

Investors in equity options face an equally difficult task, with less resources. For each underlying stock, basket or index, many standard strikes and expirations are available. For a given underlying, each strike and expiration trades at its own implied volatility, all of which, together, comprise an implied volatility surface [Derman, Kani and Zou, (1996)] that moves continually. Each underlying has its own idiosyncratic surface. In addition, underlyers can be grouped to create baskets, new underlyers with their own (never before observed) volatility surface.

For a given stock or index, how is an investor to know which strike and expiration provides the best value? What metric can options investors use to gauge their estimated excess return? What is the appropriate volatility surface for an illiquid basket? Help is sparse.

**Current vs. Past
Implied Volatilities**

The most common gauge of options value has been the spread between current and past implied volatilities. This is the metric of options speculators, who hope to get in at historically low volatilities, hedge for a while, and get out high. When all options of a given expiration trade at the same implied volatility, it is not too hard to compare changes in implied volatility over time. Since the advent of the volatility smile, however, it has become harder to have a clear opinion of the relative richness of two complex volatility surfaces.

**Implied vs. Historical
Volatilities**

A second gauge is the spread between current implied and past realized volatilities. This is the metric of options replicators, who hope to lock in the difference between future realized and current implied volatilities by delta-hedging their options to expiration. This comparison, becomes imprecise in the presence of a volatility skew, when there are a range of implied volatilities, varying by strike, that must be compared with only a single historical realized volatility.

Strike-Adjusted
Spread
(SAS)

The historical time series of a stock's returns contains much useful information. In this paper we try to come to the practical aid of options investors by estimating the fair value of options from the historical returns of their underlyers. This method for options pricing has been extensively developed by Stutzer (1996), and also employed by Derman, Kamal, Kani & Zou (1997), and Stutzer and Chowdhury (1999). Here we apply it in the practical situations that occur on an equity derivatives trading desk, where options on many different underlyers must be valued daily.

This method leads us to the notion of *Strike-Adjusted Spread*, or *SAS*, a natural one-dimensional metric with which to rank the relative value of all standard equity options, irrespective of their particular strike or expiration. We propose to use *SAS* in roughly the same way that stock investors use "alpha" and mortgage investors use *OAS* (option-adjusted spread). To be specific, the *SAS* of an option is the spread between the current market implied volatility of that option and our model's estimate of its historically appropriate volatility. Our estimate includes both the effect of past price jumps and the influences of changes in volatility and correlations for basket options.

Theoretically, the historically appropriate implied volatility for a given option is determined by the cost of replicating that option throughout its lifetime. Not only is this replication cost difficult and time-consuming to simulate, but, in our experience, the hedging errors due to inaccurate volatility forecasting and infrequent hedging make the resulting statistics inconclusive. Instead, our method for obtaining the appropriate implied volatility of a stock option involves the estimation of an appropriate *risk-neutral* distribution from the past *realized* return distribution of the stock. We will explain the method in more detail below, and describe its application to *SAS*. The same technique can be used to mark and hedge illiquid equity options whose market prices are unknown.

The strike-adjusted spread of an option depends on both its strike K and time to expiration T , and can be written more precisely as $SAS(K, T)$. *SAS* can be thought of as an extension of the commonly quoted implied-to-historical volatility spread, which is unique only in the absence of skew. In non-skewed worlds, both spreads become identical.

In brief, the *SAS* of a stock option is calculated as follows.

1. First, choosing some historically relevant period, we obtain the distribution of stock returns over time T . This empirical return distribution characterizes the past behavior of the stock.

2. Option theory dictates that options are valued as the discounted expected value of the option payoff over the risk-neutral distribution. We do not know the appropriate risk-neutral distribution. However, we use the empirical return distribution as a statistical prior to provide us with an estimate of the risk-neutral distribution by minimizing the entropy¹ associated with the difference between the distributions, subject to ensuring that the risk-neutral distribution is consistent with the current forward price of the stock. We call this risk-neutral distribution² obtained in this way the **risk-neutralized historical distribution**, or **RNHD**.
3. We then use the RNHD to calculate the expected values of standard options of all strikes for expiration T , and convert these values to Black-Scholes implied volatilities. We denote the Black-Scholes implied volatility of an option whose price is computed from this distribution as Σ_H . This is our estimated fair option volatility.
4. For an option with strike K and expiration T , whose market implied volatility is $\Sigma(K, T)$, the strike-adjusted spread in volatility is defined as

$$SAS(K, T) = \Sigma(K, T) - \Sigma_H(K, T)$$

This spread is a measure of the current richness³ of the option based on historical returns.

ATM Strike-Adjusted Spread

The volatility skew, the relative gap between at-the-money and out-of-the-money implied volatilities for a given expiration, is more stable than the absolute level of at-the-money implied volatilities. Often, therefore, irrespective of historical return distributions, the current level of at-the-money implied volatility is the most believable estimate of future volatility. It is likely that historical distributions tell us more about the higher moments of future distributions than it does about their standard deviation.

Therefore, we will often use a modified version of *SAS* for which the risk-neutralized historical distribution is further constrained to repro-

-
1. As we explain later, markets in equilibrium are characterized by maximum uncertainty or minimal information, and minimal entropy change is an expression of minimal information.
 2. Stutzer (1996) refers to this as the “canonical distribution.” and this method of options valuation as “canonical valuation.”
 3. A positive *SAS* connotes richness only for standard options whose value is a monotonically increasing function of volatility. Exotic options may have values that decrease as volatility increases.

duce the current market value of at-the-money options. We call this (additionally constrained) distribution the **at-the-money adjusted, risk-neutralized historical distribution**, or RNHD_{ATM} . The strike-adjusted spread computed using this distribution, denoted $SAS_{\text{ATM}}(K, T)$, is a measure of the relative value of different strikes, assuming that, by definition, at-the-money-forward implied volatility is fair.

We propose using $SAS_{\text{ATM}}(K, T)$ to rank options on the same underlying, in order to determine which strikes provide the best value by historical standards. More radically, we can also use the same measure to compare options of different underlyers.

In the remainder of this paper, we flesh out these concepts. The next section explains the relation between options prices and implied distributions. Thereafter, we compare implied distributions to historical return distributions. We then explain that markets in equilibrium are characterized by maximal investor uncertainty, and, introducing the notion of entropy, show that we can obtain an estimate of the risk-neutral distribution from the historical distribution by minimizing the entropy difference between the distributions. The main body of the paper then develops several applications of the risk-neutralized historical distribution, including SAS . After some concluding remarks, we provide several mathematical appendices.

OPTIONS PRICES AND IMPLIED DISTRIBUTIONS


According to the theory of options valuation, stock options prices contain information about the market's collective expectation of the stock's future volatility and its return distribution. If no riskless arbitrage can occur, there exists a risk-neutral return probability distribution Q such that the value V of an option on a stock with price S at time t is given by the discounted expected value of the option's payoff, written as

$$V(S, t) = e^{-r(T-t)} E_Q[\text{Option Payoff at } T \mid S, t] \tag{EQ 1}$$

where r is the risk-free interest rate and $E_Q[\mid \mid]$ denotes the expected value of the future payoff at time T , given that the stock price at time t is S .

In the Black-Scholes theory, the risk-neutral implied probability distribution Q is the lognormal density function with a specified volatility. In implied tree models⁴, Q is skewed relative to the Black-Scholes density, and can be estimated at any time from a set of traded European option prices. Figure 1a illustrates an implied volatility skew for S&P 500 index options, typically about five volatility points for a 10% change in strike level; Figure 1b shows the correspondingly skewed risk-neutral implied distribution Q . Whenever the shape of the skew changes, there is a corresponding change in the distribution. Knowing Q , you can calculate the fair value of any standard European option.

FIGURE 1. (a) The three-month implied volatility skew for S&P 500 index options on 3/10/99. (b) Q , the corresponding risk-neutral implied probability distribution of returns. We assume a riskless interest rate of 5%.



4. See for example, Derman, Kani, and Zou (1996).

STOCK RETURNS AND HISTORICAL DISTRIBUTIONS

Stock options' prices determine the implied distribution of stock returns. Independently, we can also observe the actual distribution of stock. Consider the historical series S_i of daily closing prices of a stock or stock index. We can construct the rolling series of continuously compounded stock returns R_i from day i for a subsequent period of N trading days by calculating


$$R_i = \log \left[\frac{S_{i+N}}{S_i} \right] \tag{EQ 2}$$

Figure 2 shows the distribution of actual three-month S&P 500 returns for periods both before and since the 1987 stock market crash, *where the latter period includes the crash itself*.

The pre-crash return distribution is approximately symmetric and normally distributed. In contrast, the post-crash distribution (1987-crash data included) has a higher mean return and a lower standard deviation, as well as an asymmetric secondary peak at its lower end.

There is a rough similarity in shape between the implied distribution of Figure 1b, whose mean reflects the risk-free rate at which its options were priced, and the historical distribution of Figure 2b, whose (different) mean is the average historical return over the post-crash period.

FIGURE 2. Three-month, S&P 500 index, observed return distributions. (a) Pre-1987 crash (Jan. 1970 to Jan. 1987); (b) Post-1987 crash (June 1987 to June 1999)



Options theory does not enforce an unambiguous link between historical and implied distributions. Nevertheless, historical distributions, suitably interpreted, can provide plausible information about fair options prices. Our aim in this paper is to develop a heuristic but logical link between the two distributions, utilizing the notions of market equilibrium and uncertainty.

MAXIMAL
UNCERTAINTY AND
MARKET EQUILIBRIUM

Markets are supposed to settle into equilibrium when supply equals demand, when there are equal numbers of buyers and sellers at some price. In an efficient market, the potential buyers of a stock must think the stock is cheap, and potential sellers must think it rich. This difference of opinion means that, in equilibrium, the distribution of expected returns displays great uncertainty.

How do we quantify this simple intuition that equilibrium involves uncertainty in the expected return distribution?

Entropy as a Measure
of Uncertainty

The probability of a *single* event is a measure of the uncertainty of its occurrence. *Entropy* is a mathematical function that measures *the uncertainty of a probability distribution*⁵. The entropy of a random variable R , whose i^{th} occurrence in the distribution has probability p_i , is defined to be

$$H(R) = - \sum_{i=1}^n p_i \log p_i \tag{EQ 3}$$

Since any probability p_i is less than or equal to 1, the entropy is always positive. If the distribution R collapses to one certain single event j , whose probability $p_j = 1$ with all other $p_i = 0$, then $H(R) = 0$. Therefore, certainty corresponds to the lowest possible entropy. You can also show that the entropy takes its maximum value, $\log(n)$, when $p_i = 1/n$ for all i , that is, when all outcomes have an equal chance and uncertainty reigns. This is consistent with the notion that maximum entropy corresponds to maximum uncertainty and minimum information.

$H(R)$ is the entropy of a single distribution R . We can also define the *relative entropy* $S(P, Q)$ between an initial distribution P and a subsequent distribution Q . S measures the decrease in entropy (or the increase in information) between the initial distribution P and the final distribution Q , and is given by

$$S(P, Q) = E_Q[\log Q - \log P] = \sum_x Q(x) \log \left(\frac{Q(x)}{P(x)} \right) \tag{EQ 4}$$

In Appendix A we show that the relative entropy is always non-negative, and is zero if and only if the two distributions P and Q are identi-

5. In Appendix A we explain the link between entropy and information.

cal. This agrees with our intuition that any change in a probability distribution conveys some new information. The relative entropy between two distributions measures the information gain (or reduction in uncertainty) after a distribution change. Thus, minimum relative entropy corresponds to the least increase in information.

The Risk-Neutralized Historical Distribution

Consider a stock option with time to expiration T on a stock whose spot price is S_0 . To value the option, we need to average the option payoff over the risk-neutral probability density $Q(S_0, 0; S_T, T)$. In theory, $Q(\cdot)$ is found by solving the differential equation that constrains the instantaneously hedged option to earn the instantaneously riskless return. In the Black-Scholes world, a stock's future probability distribution is assumed to be lognormal, and consequently, though not obviously, $Q(\cdot)$ itself is a lognormally distributed probability density, and its options prices have no volatility skew.

This theoretical lack of skew conflicts with the data from markets, where stocks and indexes that have sufficiently liquid out-of-the-money strikes display clear, and often large, skews. How can we estimate a suitable risk-neutral probability density that is more consistent with market skews than the Black-Scholes lognormal distribution?

It is natural to turn for insight to the distribution of actual returns, $P(S_0, 0; S_T, T)$. The two distributions $Q(\cdot)$ and $P(\cdot)$ cannot be strictly identical, because the expected value of the stock price under the risk-neutral distribution $Q(\cdot)$ at any time must be the stock's current forward price, as determined by the current risk-free rate, whereas the expected value of the stock price under $P(\cdot)$ is the average historical forward price, which bears no relation to current risk-free rates.

The rigorous way to obtain $Q(\cdot)$ from the past evolution of stock prices is to obtain fair historical options prices for a variety of strikes by simulating the instantaneously riskless hedging strategy over the life of these options, and to then infer the risk-neutral density that matches these prices. This requires a detailed knowledge of every past instant of the stock price evolution, at all times and market levels, and is time-consuming, difficult, error-prone and ultimately impractical.

Instead, we will estimate the current risk-neutral return distribution $Q(\cdot)$ for a stock from its historical distribution $P(\cdot)$ by assuming that the latter is a plausible estimate for the former, and then requiring that the relative entropy $S(P, Q)$ between the distributions is minimized. We impose this criterion in order to avoid any spurious increase in apparent information in creating the risk-neutral distribution from the historical distribution. We perform the minimization subject to the risk-

neutrality constraint, that is, the condition that the expected value of the stock price under the risk-neutral distribution $Q(\cdot)$ is consistent with the stock's current forward price⁶. We call $Q(\cdot)$ found in this way the **risk-neutralized historical distribution**⁷, or the **RNHD**. It is our plausible guess for the distribution to use in options valuation, given our knowledge of the past. Our knowledge of a stock's historical volatility, the second moment of its distribution, is often used to estimate options values using the Black-Scholes formula. Here we go one step further by using the entire historical return distribution. A description of the general approach outlined here can also be found in Stutzer (1996).

It is possible to impose further constraints on $Q(\cdot)$. If you believe that the current at-the-money volatility for some particular stock is fair, you can constrain the distribution $Q(\cdot)$ to match not only the stock forward price, but also to match the current at-the-money implied volatility. We denote this additionally constrained distribution by $Q_{atm}(\cdot)$, and refer to it as the **at-the-money-consistent, risk-neutralized historical distribution**, or **RNHD_{ATM}**. It can be used to compare the relative values of options with different strikes on one underlying, assuming that at-the-money volatility is fair.

Appendix B states the minimization condition on $Q(\cdot)$ in mathematical terms. In Appendix C we present a model of an Arrow-Debreu economy and show that it is possible to obtain the risk-neutralized historical distribution by optimally allocating investors' wealth under an equilibrium condition with an exponential utility function.

Having obtained our estimate of the risk-neutral distribution, we can estimate the fair price for any standard option as the discounted expected value of its payoff at expiration. We then extract the fair implied volatility as the volatility which equates the Black-Scholes option price to the estimated fair price. This procedure can be repeated for all strikes and maturities to yield an entire fair implied volatility surface.

6. Several authors have studied the relevance of entropy in financial economics and derivatives pricing. See Stutzer (1996), Derman et al. (1997), Buchen and Kelly (1996), and Gulko (1996).

7. For a normal historical distribution of simply compounded returns, one can show that the risk-neutralized historical distribution obtained by entropy minimization is equivalent to a translation of the historical distribution to re-center it at the appropriate risk-neutral rate, without altering its shape. This translation invariance of the shape in moving from the historical to the risk-neutral distribution does not hold in general.

APPLICATIONS OF THE RISK-NEUTRALIZED HISTORICAL DISTRIBUTION

Is The Index Implied Volatility Skew Fair?

The RNHD contains information which can be used to estimate the value of illiquid options whose prices are unobtainable, as well as to compare the relative value of options with known market prices. We present several representative examples below.

Since the 1987 crash, equity index markets have displayed a pronounced, persistent implied volatility skew. Is this skew fair? Are the options prices determined by the skew justified by historical returns? Figure 3a shows the risk-neutralized three-month S&P 500 return distribution for the pre-crash period corresponding to Figure 2a, constructed using our method of relative entropy minimization. Figure 3b shows the same distribution corresponding to the post-crash era of Figure 2b. The post-crash distribution has a substantially longer tail at low returns than the pre-crash distribution.

Skew slopes seem more stable than volatility levels. Therefore, we will focus here on the relation between the implied volatilities of different strikes that follows from these distributions, and pay little attention to the prevailing absolute level of implied volatility. We estimate the fair volatility skew by using the distributions of Figure 3 to calculate options prices, and by then converting these options prices to Black-Scholes implied volatilities.

The results are shown in Figure 4. The pre-crash skew is approximately flat, but the post-crash volatilities increase for low strikes, with a slope similar to actual index skews in stable markets. The observed degree of skew, about five to six volatility points per 10% change in

FIGURE 3. The three-month risk-neutral distribution of S&P 500 returns constructed from the empirical distributions of Figure 2 using a 6% riskless rate. (a) Pre-crash (b) Post-crash.





FIGURE 4. The fair implied volatility skew for three-month S&P 500 options as calculated from the risk-neutralized historical distributions of Figure 3.



strike level, seems approximately fair in the light of post-crash market behavior. Our fair post-crash skew is bilinear and more convex than the recent skew of Figure 1a, but index markets do sometimes display skews like that of Figure 4.

We find that estimated one-month skews tend to resemble a smile more than a skew: our fair implied volatilities of both out-of-the-money calls and puts for one-month expirations exceed at-the-money volatilities. Short-dated index options often display this type of behavior. We have applied our method to several other major stock indexes and found that their fair volatility skews are roughly consistent with observed market skews during normal market periods, as shown in Table 1.

TABLE 1. Comparison of actual skews with estimated fair volatility skews for three major indexes. The spread shown is the difference in volatility points between a 25-delta put and a 25-delta call.

Index	Normal ^a Spread	Extreme ^b Spread	Fair ^c Spread
SPX	4-7%	14%	6.0%
DAX	3-6%	10%	3.5%
FTSE	2-6%	10%	4.0%

a. Average during normal market conditions (excluding the periods of extreme volatility in late October 1997 and August-September 1998).
 b. Average during periods of extreme market volatility.
 c. Based on historical returns over the period June 1987 to June 1999.

Strike-Adjusted
Spread As A Measure
Of Options Value

The Strike-Adjusted Spread for an option with strike K and expiration T is defined as

$$SAS_{ATM}(K, T) = \Sigma(K, T) - \Sigma_H(K, T)$$

where $\Sigma(K, T)$ is the Black-Scholes market implied volatility of the option, and $\Sigma_H(K, T)$ is the implied volatility computed from the RNHD over some chosen relevant period. SAS_{ATM} is constrained to be consistent with the market's at-the-money-forward implied volatility for that particular underlying and expiration, so that $SAS_{ATM}(S_F[T], T) = 0$, where $S_F[T]$ is the forward value of the underlying at time T . This spread is a measure of the current richness, relative to history, of an option, assuming that at-the-money-forward options, usually the most liquid, are fairly valued.

Figure 5a shows a plot of fair and market skews for Sept. 1999 S&P 500 options, on May 18, 1999, using the twelve years of historical returns from May 1987 to May 1999 to calculate $\Sigma_H(K, T)$. Figure 5b shows the SAS_{ATM} for the same options. For out-of-the-money puts, the entropy-adjusted volatilities slightly exceed the market volatilities, which suggests that out-of-the-money puts are slightly cheap. Conversely, out-of-the-money calls seem rich.

Figure 6 shows the same plots based on a historical return distribution taken from May 1988 through May 1999, thereby excluding the 1987 global stock market crash. In this case, out-of-the-money puts seem much too rich, while out-of-the-money calls are slightly cheap.

In our view, SAS is a quantitative tool for ranking the relative value of options, but this does not absolve the user from choosing the historical period relevant to the computation of the risk-neutralized distribution. There is no escaping the judgement necessary to decide which past period is most relevant to the current market from both a fundamental and psychological point of view.

In Figure 7, we plot the skews and SAS_{ATM} for the same set of options used in Figure 5 and 6, but evaluated one month later. Although at-the-money volatility has now fallen from 25.5% to 21%, the size of skews has remained relatively stable. Roughly irrespective of which historical distribution was used, the strike-adjusted spreads have changed so that out-of-the-money puts have become about two SAS points cheaper, whereas the SAS of out-of-the-money calls has changed less. If you had thought the relevant historical distribution was the crash-inclusive one of Figure 5, and had bought cheap puts, you would

FIGURE 5. (a) Fair and market skews for S&P 500 index options on May 18, 1999. (b) SAS_{ATM} for the same options.

The options considered expire on September 17, 1999. Both fair and market implied volatilities are constrained to match at the money, forward. The RNHD is constructed using returns from May 1987 to May 1999, including the 1987 crash.






FIGURE 6. (a) Fair and market skews for S&P 500 index options on May 18, 1999. (b) SAS_{ATM} for the same options.

The options considered expire on September 17, 1999. Both fair and market implied volatilities are constrained to match at the money, forward. The RNHD is constructed using returns from May 1988 to May 1999, thereby excluding the 1987 crash.



have lost *SAS*. If, on the other hand, you had thought that the relevant distribution was the crash-exclusive one of Figure 6, and sold rich puts, you would have gained several points of *SAS*

FIGURE 7. Re-evaluated SAS_{ATM} on June 21, 1999 for September 17, 1999 S&P 500 options. The top two figures correspond to the crash-inclusive distributions of Figure 5; the bottom two correspond to the crash-exclusive distributions of Figure 6.



Valuing Options on Baskets of Stocks

The value of an OTC option on a custom basket of stocks is difficult to estimate, since there is no liquid options market from which to extract pricing information. Consider an investor interested in buying a collar on a basket of bank stocks he owns. Suppose he wants to buy a 10% out-of-the-money put and finance it by selling a 10% out-of-the-money call on the basket. What volatility spread or skew should one use to price the collar?


One of the problems in valuing basket options is that correlations among component stocks vary widely with the stock levels. When there is a sharp downward market move, the correlations tend to increase. Consequently, a basket may exhibit large volatility skews even if each component stock shows little skewness in its distribution. Thus OTM puts on a basket should, in general, trade at a premium because the

increasing correlations in a down market makes the basket return distribution skewed to the lower end. How do we model this effect? We utilize the information embedded in the historical time series of the basket.

To be specific, we consider an example in which the basket consists of an equal number of shares of five bank stocks: J.P. Morgan, Wells Fargo, Bank One, Bank America, and Chase. We first retrieve the historical data for all five stocks and aggregate them to form the time series of basket returns and their historical distribution. We use historical data from June 1987 to June 1999 in this example. By minimizing the relative entropy, we convert the historical distribution into an estimate of the risk neutral distribution. Figure 8 displays the estimated three-month implied volatility skew for the bank basket calculated from the risk-neutral distribution. The volatility spread between the 10% OTM call and the 10% OTM put is approximately seven volatility points. In the absence of any market information on the price of options on this basket with a variety of strikes, this seems a useful method of obtaining some sense of the appropriate skew. We note that in using this approach, we managed to bypass the problem of predicting future correlations between the component stocks in the basket, a major hurdle in valuing basket options. In this particular example, the three-month correlations between the stocks in the basket almost doubled during the Fall of 1998 following the Russian currency devaluation. Our approach takes into account the changes in correlations embedded in the basket time series.

We have also applied our model to options on the BKX index (a basket of 24 large U.S. banks with options listed on the Philadelphia exchange). We constructed a basket with the same weighting as the BKX index and calculated both its empirical return distribution and its estimated risk-neutral distribution. The resulting three-month volatility skew is close to the skew observed in the listed options market, even when the at-the-money volatility levels differ. This further demonstrates the reasonableness of our approach.

FIGURE 8. The estimated fair three-month implied volatility skew for the basket of five bank stocks listed in the text, estimated from the risk-neutralized historical distribution using returns from June 1987 to June 1999.



Forecasting The Shape of The Skew From A Change In A Single Option Price

Market makers in index options keep a steady eye on the skew. Suppose that for a given expiration there are n options, with strikes K_i and known implied volatilities $\Sigma_i(K_i)$, that characterize the skew. Suppose that implied volatilities and the skew have been relatively stable; then, one of the option's implied volatilities suddenly changes in response to new market sentiments or pressures. How should the market maker adjust the quotes for all other options given the sudden change in the price of one? This question is particularly relevant for automated electronic market-making systems. The maximum entropy method provides a possible answer.

We start with the implied distribution⁸ computed from the known implied volatilities $\Sigma_i(K_i)$. Now suppose the implied volatility of one option with strike level K_j has changed to a new value $\tilde{\Sigma}(K_j, T)$. We would like to regard this one move in implied volatility as the visible tip of the iceberg, the observable segment of a new skew that will soon manifest. To identify this new skew, we seek to find the new risk-neutral distribution $\tilde{Q}(S_T, T|S_0, 0)$ that is consistent with the single new and known implied volatility, while minimizing the entropy change

8. See Derman, Kani, and Zou (1996).

between the old and new distributions. Once the new risk-neutral distribution is obtained, we can update the quotes for the rest of the options by valuing them off the new distribution.

Here are some examples. Consider a hypothetical index whose current value is 100. Suppose the three-month, at-the-money volatility is 24%, and the three-month skew is linear in strike with a slope corresponding to a two-volatility-point increase per ten-strike-point decline, as displayed in Figure 10a. The heavy X in the figure shows the one newly observed implied volatility, assumed to rise of four implied volatility points, from 26% to 30%, for the 90-strike put. Figure 10b shows the change in the risk-neutral implied distribution obtained by minimizing the change in distributional entropy consistent with one new implied volatility. The increase in 90-strike implied volatility has led to a significant hump in the risk-neutral distribution below the 90 level. Finally, Figure 10c shows both the old and new skews, the latter computed from the new risk-neutral distribution. The new estimated skew differs from the old in a non-obvious way: it has not shifted parallel to accommodate the one new item of information, but instead suggests that the skew slope will increase as a response to this shock.

Figure 9 shows several more examples where the volatility at one particular strike is shocked. Note that shifts in at-the-money volatility seem to lead to parallel shifts in the skew, whereas shifts in out-of-the-money volatility lead to changes in slope as well as level.

End-of-Day Mark To Market

At the end of a trading day, volatility traders need to re-mark all of their options positions. Often, only the liquid strikes have traded close to the end of the day, and, if the last traded option has undergone a significant change in implied volatility, one needs to estimate the appropriate skew for the remaining, less liquid options, based on the new information. Our model of minimizing the relative entropy provides a ready solution to the problem. In the examples shown in Figure 9, one could plausibly mark to market the rest of options using the solid lines as our best guesses for closing volatilities.

Filling Gaps In Investors' Market Views

Risk arbitrageurs, speculators, and other situation-driven investors often have very specific, but not necessarily complete, views on the market. For instance, risk arbitrageurs taking a position in the stocks of two companies involved in a merger or acquisition may estimate a 90% probability that the deal will be completed and the stock will move

FIGURE 9. The shifts in skew that result from minimizing the entropy change of the new risk-neutral distribution to accommodate a shift in one implied volatility, as shown by the arrows below.





FIGURE 10. (a) The initial skew (dashed line) and the one newly observed implied volatility (X). (b) The initial (dashed line) risk-neutral distribution and the new (solid line) risk-neutral distribution. (c) The new skew (solid line), with greater negative slope.



above a certain target level within three months. Suppose that they would like to take positions in stock and options to implement this belief.

Although the arbitrageurs have a firm opinion about only one segment of the probability distribution, a more complete distribution can be helpful in constructing strategies and determining reasonable prices. How can we fill the gap and extract the arbitrageurs' market distribu-

tion based on their specific prediction that the market has a 90% chance of moving above the target level? One good starting point is to use the option implied distribution, P , as a prior distribution, and then find a new distribution, Q , that satisfies the arbitrageur's 90% probability estimate

$$\int_{S_T}^{\infty} Q(S, T; S_0, 0) dS = 0.9$$

where S_T is the arbitrageur's target stock price. Again the model of minimizing the relative entropy provides a natural solution to this problem.

**CONCLUDING
REMARKS**

Investors in equity options have two problems that compound each other: the many thousands of equity underlyers and the presence of a unique volatility skew for each of them. For many thinly-traded single stock and basket options, it is difficult or impossible to get adequate information on the market skew.

In this paper, we have employed a systematic, semi-empirical method for estimating the risk-neutral distribution of any underlyer, stock or basket, whose historical returns are available. This method, originally used by Stutzer (1996), involves the determination of a new, risk-neutralized historical distribution (RNHD) for an underlyer by minimizing the relative entropy between the historical distribution and the risk-neutral distribution.

Using the RNHD, we can compute the estimated fair implied volatilities of options of any strike and expiration. We can apply this method to illiquid or thinly-traded derivatives where market prices are unavailable.

We have defined a new metric, the strike-adjusted spread, or SAS, for gauging the value of options whose prices are known. SAS is the difference between an option's implied volatility and its fair volatility as estimated using the RNHD. This spread represents the richness in volatility points of an option, compared to the history of its underlyer. Most often, in liquid markets, we calibrate the SAS to be consistent with current at-the-money volatility, so that it becomes a measure of skew richness as compared with history. The SAS ranking cannot be used blindly; it depends on the user's selection of the historical period most relevant to the current market.

There are many other applications of the method of minimal relative entropy we have illustrated in this paper. One may choose as a prior an existing options' implied distribution, or any other distribution reflecting subjective market views. We hope that this practical method and its extensions will help investors make more rational decisions about value in volatility markets.

APPENDIX A:
INFORMATION AND
ENTROPY

Probability measures the uncertainty about the occurrence of a single random event. For a given random variable X , what can we deduce from a single observation that $X = x$?

Information

Information⁹ changes our view of the world. It seems obvious that the amount of information conveyed by the observation that $X = x$ should depend on how likely this event was previously assumed to be. If a stock was expected to go up the next day by everyone, and it actually went down, the surprising outcome is certainly more informative than the expected outcome. We want to quantify the notion of events providing information.

We seek a function $I(p)$ that can represent the information provided by the occurrence of the event $X = x$ whose probability was assumed to be p . We require that $I(p)$ be a non-negative and decreasing function of p ; our intuition says that $I()$ is *non-negative* because the occurrence of the event must provide some information; similarly, $I()$ must *decrease with increasing p* because the more likely we thought the event was, the less information its occurrence provides.

Consider X and Y to be two independent random variables, and assume that

$$P(X = x) = p \quad \text{and} \quad P(Y = y) = q \tag{A 1}$$

Since X and Y are independent, we have the joint probability

$$P(X = x; Y = y) = pq \tag{A 2}$$

When both independent events, $X = x$ and $Y = y$, occur, the associated information $I()$ of each must add to the total, so that

$$I(pq) = I(p) + I(q) \tag{A 3}$$

Differentiate Equation 3 first with respect to p and similarly with respect to q to obtain

$$\begin{aligned} q \frac{\partial}{\partial(pq)} I(pq) &= \frac{\partial}{\partial p} I(p) \\ p \frac{\partial}{\partial(pq)} I(pq) &= \frac{\partial}{\partial q} I(q) \end{aligned} \tag{A 4}$$

Dividing the first equation by the second, we obtain

9. See Kullback, S. (1967), and Cover et al. (1991).

$$\frac{q}{p} = \frac{\frac{\partial}{\partial p} I(p)}{\frac{\partial}{\partial q} I(q)}$$

or

$$q \frac{\partial}{\partial q} I(q) = p \frac{\partial}{\partial p} I(p) \tag{A 5}$$

Since q and p are independent variables, each term in Equation 5 must be a constant, denoted $-c$, so that

$$I(q) = -c \ln(q) + A$$

Since an event with probability $q = 1$ provides no information, $A = 0$. Since $0 \leq q \leq 1$, and since we required that $I(\cdot)$ be a non-negative and decreasing function of q , the constant c must be positive. From now on we set $c = 1$, which defines the conventional size of a unit of information. The information provided by the occurrence of an event whose probability was p is therefore given by

$$I(p) = -\ln(p) \tag{A 6}$$

Entropy

The probability assigned to a *single* event is a measure of the uncertainty of its occurrence. The *entropy* of a random variable R , whose i^{th} occurrence in the distribution has probability p_i , is defined to be the expected value of the information from the occurrence of an event in the distribution, namely

$$H(R) = - \sum_{i=1}^n p_i \log p_i \tag{A 7}$$

Since any probability p_i is less than or equal to 1, the entropy is always positive.

A large expected value of information means that the distribution was broad, with a wide spread of probabilities. A small expected information means the distribution was relatively narrow, so that not much information can be gained from the occurrence of an expected event. Qualitatively, therefore, you can see that H represents the uncertainty in the distribution: large (small) average information H corresponds to

high (low) uncertainty. *Entropy* is the mathematical function that measures *the uncertainty of a distribution*. This is consistent with our intuition that maximum entropy corresponds to maximum uncertainty.

If the distribution collapses to one certain single event j whose $p_j = 1$, all other $p_i = 0$, then $H = 0$, a minimum. You can also show that the entropy takes its maximum value, $\log(n)$, when $p_i = 1/n$ for all i , i.e. when all outcomes have an equal chance and there is maximal uncertainty.

Having established the notion that the entropy of a probability distribution reflects the expected amount of information, we can now quantify the information gained upon changing the distribution as a result of new information. Let us assume we have a prior distribution of a random variable X which we denote by P . Upon arrival of new information, a posterior distribution Q is established. What is the consequent reduction in uncertainty (decrease in entropy) in this process? One obvious choice is the relative entropy:

$$S(P, Q) = E_Q[\log Q - \log P] = \sum_i q_i \log \frac{q_i}{p_i} = -\sum_i q_i \log \frac{p_i}{q_i} \quad (\text{A } 8)$$

The $-\log(\)$ function is convex, so that, by Jensen's inequality, the average of the $-\log(p_i/q_i)$ is greater than the $-\log(\)$ of the average of (p_i/q_i) . Therefore,

$$S(P, Q) > -\log \sum_i \left(q_i \frac{p_i}{q_i} \right) = -\log \sum_i (p_i) = -\log 1 = 0$$

Therefore, $S(P, Q)$ is strictly non-negative, and is zero if and only if $P \equiv Q$ identically. $S(P, Q)$ can be thought of as a "distance" between two probability distributions.

To maintain maximum uncertainty given some new information, our goal is to minimize this relative entropy or "distance" between the prior and posterior distributions.

APPENDIX B:
DETERMINING THE
RISK-NEUTRAL
DISTRIBUTION FROM
HISTORICAL RETURNS

Consider a stock whose spot price is S_0 . Assume that investors believe the future return distribution of the asset is given by a prior $P(S_0, 0; S_T, T)$. Our goal is to infer the risk-neutral distribution $Q(S_0, 0; S_T, T)$ from $P(S_0, 0; S_T, T)$ subject to the constraint that the mean of the risk-neutral stock distribution must equal the stock forward price. We determine the risk-neutral distribution by minimizing the relative entropy between $P()$ and $Q()$ subject to $Q()$ satisfying the forward condition.

We seek $Q(S_0, 0; S_T, T)$ so that

$$\text{Min} \left\{ S(P, Q) = E_Q \left[\log \left(\frac{Q(S)}{P(S)} \right) \right] \right\} \quad (\text{B } 1)$$

such that

$$\int Q(S_T) S_T dS_T = S_0 e^{r_f T} \quad (\text{B } 2)$$

and

$$\int Q(S_T) dS_T = 1 \quad (\text{B } 3)$$

where r_f is the current riskless interest rate.

In some cases we will also have a good idea of the current value of the stock's at-the-money implied volatility. In that case we will add to Equation B2 and Equation B3 the further constraint that the at-the-money implied volatility produced by the distribution $Q()$ is equal to the at-the-money implied volatility of the stock.

Solving the above equations, we obtain the risk-neutral distribution

$$Q(S_0, 0; S_T, T) = \frac{P(S_0, 0; S_T, T)}{\int P(S) \exp(-\lambda S) dS} \exp(-\lambda S_T) \quad (\text{B } 4)$$

where the constant λ can be found numerically by ensuring that it satisfies the forward condition

$$\int Q_\lambda(S_0, 0; S_T, T) S_T dS_T = S_0 e^{r_f T} \quad (\text{B } 5)$$

Because $P(S_T)$ in Equation B4 is always non-negative, so is $Q(S_T)$.

APPENDIX C:
A DERIVATIVE ASSET
ALLOCATION MODEL
AND THE EQUILIBRIUM
RISK-NEUTRAL DENSITY

In this appendix, we present a simple derivative asset allocation model and give a possible financial economic interpretation of the results obtained in this paper based on the maximal entropy principle. Consider an economy in equilibrium with total investor wealth W_0 . There is a market for an equity index, a riskless bond, and a complete range of derivative instruments on the index. We assume that there is a representative investor¹⁰ with a subjective market view expressed through a conditional probability density $P(S_0, 0; S_T, T)$, where S_0 is the spot index level. We also assume the return on the riskless bond over the period T is r_f . The representative investor allocates a portion of the total initial wealth to the riskless bonds and the remaining wealth to the risky assets. Let α be the portion allocated to riskless bonds, and $1 - \alpha$ be the portion allocated to risky assets. Since the risky assets can be replicated with a portfolio of Arrow-Debreu securities, the investor can simply find an optimal portfolio of Arrow-Debreu securities. An Arrow-Debreu security with parameter E has a price given by

$$\pi(S_p, t; E, T) = DQ(S, t; E, T) \quad (C 1)$$

where $D = 1/(1 + r_f)$ is the discount factor and $Q(S_p, t; E, T)$ is the risk-neutral probability density. The payoff of the Arrow-Debreu security at the expiration date is, by definition,

$$\pi(S_T, T; E, T) = \delta(S_T - E) \quad (C 2)$$

Let $\omega(E)dE$ be the portion of the $1 - \alpha$ fraction allocated to risky assets that is invested in Arrow-Debreu security with parameter E . At the end of period T , the total investor wealth will be

$$W_T(S_T) = W_0[1 + \alpha r_f + (1 - \alpha) \int \omega(E) r_E(S_T) dE] \quad (C 3)$$

where

$$r_E(S_T) \equiv \frac{\pi(S_T, T; E, T) - \pi(S_p, t; E, T)}{\pi(S_p, t; E, T)} = \frac{\delta(S_T - E)}{DQ(S, t; E, T)} - 1 \quad (C 4)$$

Note that

$$\int dE r(E) Q(S_p, t; E, T) = r_f \quad (C 5)$$

10. See Constantinides (1982).

If the representative investor changes the allocation $\omega(E)dE$, the supply and demand for the Arrow-Debreu security will also change and thus so will the shape of the risk-neutral density function $Q(S_t, t; E, T)$. Therefore, to achieve a market equilibrium, the representative investor must solve the asset allocation problem by maximizing the expected utility $U(W_T)$:

$$\begin{aligned} \text{Max } \{ E_p[U(W_T)] \} \\ \{ \alpha, \omega(E), Q(E) \} \end{aligned} \tag{C 6}$$

subject to the constraints:

$$\begin{aligned} \text{budget constraint: } \int \omega(E) dE &= 1 \\ \text{normalization: } \int Q(E) dE &= 1 \\ \text{forward constraint: } \int Q(E) E dE &= S_0(1 + r_f) \end{aligned} \tag{C 7}$$

where the expectation operator E_p is under measure P. The first order conditions for the constrained optimization problem are

$$E_p[U(W_T) W_0(r_f - \int \omega(E) r_E(S_T) dE)] = 0 \tag{C 8}$$

$$E_p[U(W_T) W_0(1 - \alpha) r_E(S_T)] = \lambda_1 \text{ for } \forall E \tag{C 9}$$

$$\int \omega(E) dE = 1 \tag{C 10}$$

$$E_p \left[U(W_T)' W_0(1 - \alpha) \frac{\omega(E) \delta(S_T - E)}{DQ^2(E)} \right] = \lambda_2 + \lambda_3 E \text{ for } \forall E \tag{C 11}$$

$$\int Q(E) dE = 1 \tag{C 12}$$

and

$$\int Q(E) E dE = S_0(1 + r_f) \tag{C 13}$$

where λ_1, λ_2 , and λ_3 are three Lagrange multipliers corresponding to the three constraints in Equation C7, respectively. From Equation C8 and C9 we obtain

$$E_P[U(W_T)r_E(S_T)] = \frac{\lambda_1}{W_0(1-\alpha)} \quad (\text{C 14})$$

and

$$E_P[U(W_T)] = \frac{\lambda_1}{W_0(1-\alpha)r_f} \quad (\text{C 15})$$

From Equations C8, C14 and C15 we have

$$Q(S_0, t; S_T, T) = \frac{U(W_T(S_T))}{E_P[U(W_T)]} P(S_0, t; S_T, T) \quad (\text{C 16})$$

where

$$W_T(E) = W_0 \left[\alpha(1+r_f) + (1-\alpha) \frac{\omega(E)}{DQ(E)} \right] \quad (\text{C 17})$$

Using Equations C11, C15 and C16, we get

$$\frac{\omega(E)}{Q(E)} = \frac{r_f}{1+r_f} \left[\frac{\lambda_2}{\lambda_1} + \frac{\lambda_3}{\lambda_1} E \right] \quad (\text{C 18})$$

Using the constraints (C12) and (C13), Equation (C16) leads to

$$\frac{\lambda_2}{\lambda_1} + \frac{\lambda_3}{\lambda_1} S_0(1+r_f) = \frac{1+r_f}{r_f} \quad (\text{C 19})$$

These results so far are independent of the particular functional form of the utility function as long as the utility function is monotonically increasing and concave. We now specialize in an exponential utility given by

$$U(W_T) = -\exp(-bW_T) \quad \text{with } b > 0 \quad (\text{C 20})$$

Using Equations C16, C17, C18 and C20, one can show that

$$Q(S_0, t; E, T) = P(S_0, t; E, T) \exp(c_0 - c_1 E) \quad (\text{C 21})$$

where constant c_0 and c_1 satisfy

$$c_0 = -bW_0\alpha(1+r_f) - \log(E_P[U/b]) - \frac{\lambda_2}{E_P[U/b]} \quad (\text{C 22})$$

$$c_1 = \frac{\lambda_3}{E_P[U/b]} \quad (\text{C } 23)$$

and that we have, using Equation C17,

$$E_P[U/b] = \exp \left\{ -bW_0(1+r_f) - \int Q(E) \log \left[\frac{Q(E)}{P(E)} \right] dE \right\} \quad (\text{C } 24)$$

We can now carry out the optimization program as follows.

Step 1. From Equations C21, C12 and C13, we can numerically solve for constant c_0 and c_1 , and thus obtain the risk-neutral distribution Q . We point out that the constant c_0 and c_1 (and thus the risk-neutral distribution Q) is independent of the parameter, b , of the utility function. It only depends on the prior distribution P , and the forward price constraint Equation C13! The parameter b is characteristic of the representative investor's risk aversion. It is essential that the risk-neutral distribution be independent of the investor's risk aversion! In our derivation of distribution Q , we relied only on the fact that the investor's utility function is of exponential form.

Step 2. From Equation C24, we can solve for $E_P[U/b]$, which does depend on parameter b .

Step 3. From Equation C23, we can solve for λ_3 .

Step 4. From Equations C15, C19 and C22 we solve for λ_1 , λ_2 , and α .

Finally, the representative investor's allocation of the Arrow-Debreu security with parameter E , $\omega(E)$, can be found from Equation C18, which of course depends on the risk aversion parameter b .

The most important feature of this solution to the asset allocation problem is that the risk-neutral probability density function $Q(S_t, t; S_T, T)$ is of the form given by Equation C21 subject to the normalization constraint and Equation C13. This is the same solution as that given by the minimal relative entropy approach in Equation B4, provided that the representative investor chooses the realized historical distribution as the subjective prior distribution P .

REFERENCES

- Buchen, P.W. and M. Kelly, 1996, *Journal of Financial and Quantitative Analysis*, **31**, pp. 143 - 159.
- Constantinides, G.M., 1982, *The Journal Of Business*, **55**, No. 2.
- Cover, T.M. and J.A. Thomas, 1991, *Elements of Information Theory*, New York, John Wiley & Sons.
- Derman, E., I. Kani and J. Zou, 1996. The Local Volatility Surface, *Financial Analysts Journal*, July/August, pp. 25-36.
- Derman, E., M. Kamal, I. Kani, and J. Zou, 1997, Is the Volatility Skew Fair?, *Goldman, Sachs Quantitative Strategies Research Notes*
- Gulko, L., 1996, Yale University Working Paper.
- Kullback, S., 1967, *Information Theory and Statistics*, New York, Dover Publications, Inc.
- Stutzer, M., 1996, A Simple Nonparametric Approach to Derivative Security Valuation, *J. of Finance*, **51**. pp. 1633 - 1652.
- Stutzer, M., and M. Chowdhury, 1999, A Simple Non-Parametric Approach to Bond Futures Options Pricing, *Journal of Fixed Income*, March 1999, pp. 67-76.

SELECTED QUANTITATIVE STRATEGIES PUBLICATIONS

- June 1990 *Understanding Guaranteed Exchange-Rate Contracts In Foreign Stock Investments*
Emanuel Derman, Piotr Karasinski and Jeffrey Wecker
- Jan. 1992 *Valuing and Hedging Outperformance Options*
Emanuel Derman
- Mar. 1992 *Pay-On-Exercise Options*
Emanuel Derman and Iraj Kani
- June 1993 *The Ins and Outs of Barrier Options*
Emanuel Derman and Iraj Kani
- Jan. 1994 *The Volatility Smile and Its Implied Tree*
Emanuel Derman and Iraj Kani
- May 1994 *Static Options Replication*
Emanuel Derman, Deniz Ergener and Iraj Kani
- May 1995 *Enhanced Numerical Methods for Options with Barriers*
Emanuel Derman, Iraj Kani, Deniz Ergener and Indrajit Bardhan
- Dec. 1995 *The Local Volatility Surface: Unlocking the Information in Index Option Prices*
Emanuel Derman, Iraj Kani and Joseph Z. Zou
- Feb. 1996 *Implied Trinomial Trees of the Volatility Smile*
Emanuel Derman, Iraj Kani and Neil Chriss
- Apr. 1996 *Model Risk*
Emanuel Derman
- Aug. 1996 *Trading and Hedging Local Volatility*
Iraj Kani, Emanuel Derman and Michael Kamal
- Apr. 1997 *Is the Volatility Skew Fair?*
Emanuel Derman, Michael Kamal, Iraj Kani and Joseph Zou
- Apr. 1997 *Stochastic Implied Trees: Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility*
Emanuel Derman and Iraj Kani

- Sept. 1997 *The Patterns of Change in Implied Index Volatilities*
Michael Kamal and Emanuel Derman
- Nov. 1997 *Predicting the Response of Implied Volatility to Large
Index Moves: An October 1997 S&P Case Study*
Emanuel Derman and Joe Zou
- Sept. 1998 *How to Value and Hedge Options on Foreign Indexes*
Kresimir Demeterfi
- Jan. 1999 *Regimes of Volatility: Some Observations on the
Variation of S&P 500 Implied Volatilities*
Emanuel Derman
- Mar. 1999 *More Than You Ever Wanted to Know About Volatility
Swaps*
Kresimir Demeterfi, Emanuel Derman, Michael Kamal
and Joseph Zou