Options and Option Values
Under Incomplete Information

A Presentation to
The Mathematical Finance Group of the U. of Chicago

by

Joel Gibbons, Logistic Research & Trading Co.

&

Steven Moffitt, Dynamic Trading Systems, Ltd.

1-June-2001

4052 Niles Rd..
St. Joseph, MI, 49085
Tel: 269-408-1511
E-mail: jgibbons@logisticresearch.com

Premise:

We will assume that at any time there is a well defined predictive distribution of probabilities attaching to realizations, but that this distribution is not known.

While we must and will assume that we know something in advance, we will not assume that it is enough to specify the predictive distribution conditional on the current state of the world. What we want to assume is, in some sense, that we have somewhat fuzzy and ambiguous information that bears on, and sheds light on, the predictive distribution.

Promise:

This is not a mathematical problem, as we would usually understand that term. It is a management problem, and what we will talk about here are tools for using what is known, or believed to be known, in a way – in a somewhat impressionistic but nonetheless quite functional way – that minimizes the risk of placing too much faith in the model, while maximizing the usefulness of all prior information.

Business background:

Our business is to develop unconventional option models adapted to the needs of market professionals. The ideas we will talk about here come out of our research on financial markets and our product development work. These ideas have concrete versions that are option models we sell to the public. I mention this not as an advertisement, but by way of truth in public speaking. We expect and hope that some of you in the audience will be inspired to adopt this very flexible strategy for option modeling, and that it will be as productive in your hands as it has been for us. We would however be remiss if we did not extend an offer to work with you and with your employers and employers-to-be on a conventional fee basis.

The ideas we will be presenting are not new or novel to us. In various forms they have been used for years. It is our hope that we can nonetheless advance the state of the art.

Lastly, we want to provide a broad survey of this very extensive subject, and so we will not dwell too long on any single part. We leave it to the audience to stop us with questions that explore in more depth aspects of particular interest. We have also provided some readings that provide further insight into some parts of this broad subject.

Approaches.

We have in mind two quite different approaches

1. **Exact calculations based on models of state-dependent distributions.**

2. **Empirical methods.**

The first topic alone will however keep us fully occupied this evening. We leave the subject of empirical models to a latter talk.

State Dependent Distributions.

The idea is very simple: that the space in which outcomes are realized can be viewed as a sort of manifold of local distributions. Assuming for the moment that we are satisfied that we know what sort of distributions the local distributions are, there remains uncertainty about which locale we will be in. It is necessary in practice to assume that the state space is discrete.

If we accept this model of the world, we see that there is a discrete state space with some sort of matrix of transition probabilities among them. It might be possible to tackle this as an empirical problem, and we will present some findings of that sort. It is also possible to import non-quantitative prior information into the model. While this sounds unscientific, it may in fact be inevitable, because the data demands to estimate the transition matrix tend to be rather massive.

There is a second facet of models of this sort, and that is to use a more flexible family of locale distribution functions, and not just assume a linear diffusion. We have developed models of bounded diffusion which have the property that the "local" distribution actually stays local. We will also present some thoughts on how specific versions of state dependent models correspond to some interesting, intuitively appealing global models.

Organization of This Talk

1. **Examples of State Dependent Distributions**

    a. **Bimodal distribution of common stocks.**
    b. **Leptokurtosis of gold.**

2. Estimation of Transition Probabilities Between States.

3. Option Valuation When the Underlying Price is a Bounded Diffusion.

4. Option Valuation When the Underlying Price Exhibits Partially Reflecting Support and Resistance Points.

1.a. Bimodal Terminal Distributions at Exercise.

We have investigated many terminal distributions of common stocks for evidence of bimodality. A bimodal distribution is highly significant in this context because it arises from a branching process in the means. One common explanation for bimodality, in the case of individual stocks, is that it comes from an attempt by the options market to account for a discrete event risk. The evidence we will discuss here comes from our reconstruction of the implied risk neutral distribution of two examples, Common stock in Dell Corporation, and the implied risk neutral distribution derived from options on the Nasdaq 100 Index.

The following notes summarize our method of estimating the implied risk neutral distribution.

A. Implied Distributions

Suppose that the process of a non-dividend paying stock is a martingale. It then follows that the present value of an option is the discounted expected value of the option at expiration. Given a series of call and put options having strikes $K_1, K_2, \ldots, K_n$, market call prices $C_1, C_2, \ldots, C_n$, and market put prices $P_1, P_2, \ldots, P_n$ with the same expiration date T, one can minimize the expected sum of squared deviations

$$f(\wp) = \sum (C_i - B\, E_{\wp}[\overline{C}_i \mid T])^2 + \sum (P_i - B\, E_{\wp}[\overline{P}_i \mid T])^2$$

where

B	is the value pure discount bond paying 1 at expiration,
$\wp$	is an empirical distribution function constant on intervals $[K_i, K_{i+1}), i = 0,1,...n$ $[K_i, K_{i+1})$,
$\overline{C}_i(\overline{P}_i)$	is the random variable for the call (put) value, and
$E_{\wp}[\overline{C}_i \mid T]$	is the expected call (put) value under $\wp$ at expiration T.

Provided that appropriate conditions are set for this minimization problem, including the addition of initial and terminal pseudo-strikes K_0 and K_{n+1}, a probability distribution $\wp^*$ that minimizes $f(\wp)$ can be found. $\wp^*$ in the distribution we are looking for, at time T.

Examples: the Nasdaq 100 Index.

A proxy for the Nasdaq index is the Nasdaq Index Trust (QQQ) that trades on the AMEX, along with options on the QQQ. This synthetic security is actually a mutual fund that simply holds the one hundred stocks in proportion. Unlike the Nasdaq futures and options on futures, the QQQ pays out all dividends – meager though they are – and thus is a better replica of the index itself for our purposes. A couple of recent empirical risk neutral distributions appear below.

In this plot with the present value of QQQ, the June implied distribution is not unimodal, and in fact, is consistent with a trimodal (tetramodal?) leptokurtic distribution skewed to the left. On the other hand, the September distribution shows a similar shape with dips at 28 and 58, but a 0.14 probability of the price being 60 or greater by expiration! Implied distributions were produced by using options prices at $1 intervals from 34 to 60 in June and 25 to 60 in September. It makes little difference where the upper pseudo-strike is located; the conclusion remains.

One can interpret the 14% probability in September, however, as evidence that the upper tail of this distribution is too heavy due to overpricing of out-of-the-money calls.

Figure 1. Implied QQQ Distributions for June and September expirations based on options closing prices on May 16, 2001. – Closing QQQ Price: 48

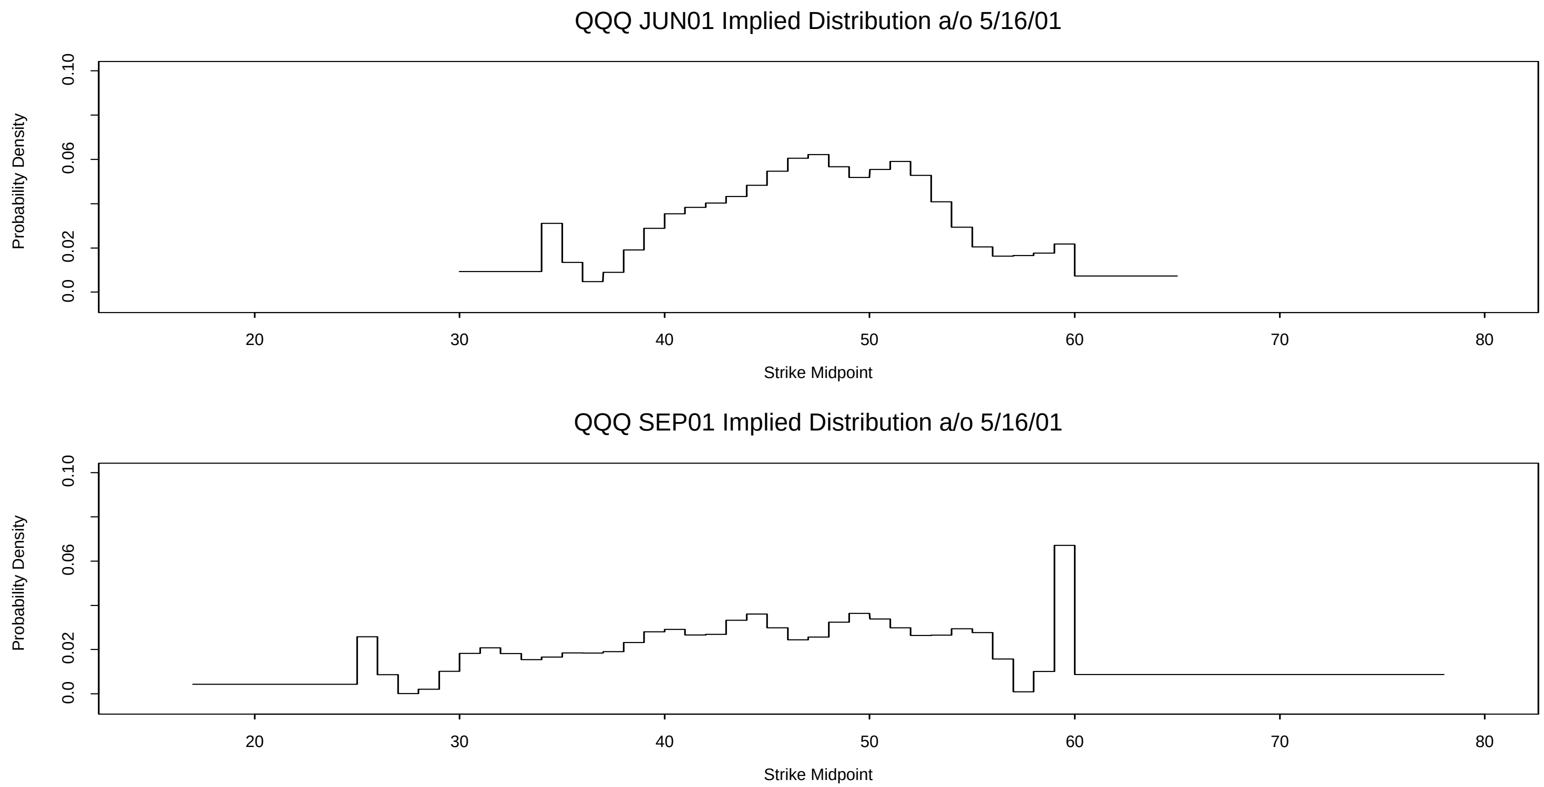

Dell Corporation

Figure 2. Implied DELL Distributions for August 2001, November 2001, and January 2002 expirations based on options closing prices on May 16, 2001. – Closing DELL Price: 25.38.

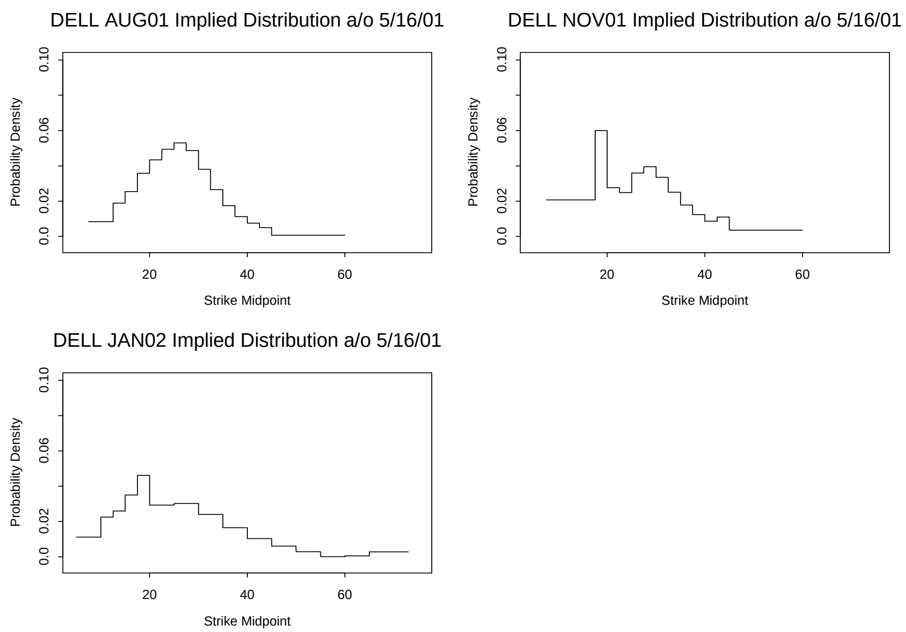

This graph showing a stock with significant "downside risk", DELL implied distributions in August 2001, November 2001, and January 2002. There is currently a price war among PC manufacturers, with DELL undercutting most of its competition by cutting expenses (with layoffs) and selling at low margins. These implied distributions have mean prices of 25.62, 25.93, and 25.03, respectively. August has a mode at 26.25, November has local maxima at 18.50 and 28.75, and January has a mode at 27.50, with probabilities of being below the (higher) mode/maxima of 0.54, 0.63, and 0.60, respectively.

1.b. Leptokurtic Distribution: Mixture of Normals with Unequal Variances.

Gold

As a second, and rather different example we resort to research we have done on the price of gold. The price series is the daily cash settlement price, quoted at about 9:20 A.M. at the Comex exchange in New York. The options data are from options on futures, but that basis difference has no bearing on this research because the broad characterization of the distribution of price changes – as a mixture of normals with the same mean but different variances – must necessarily apply also to the futures on gold.

The mixture of normals distribution is rather familiar. It is a five parameter distribution, with two mean parameters, two variance parameters, and a mixing fraction parameter which corresponds to the frequency of the two embedded normals. In the case of time series data there is a sixth parameter which is the transition probability from one state to the other. More on this later. If these parameters are, respectively, μ_1, σ_1, μ_2, σ_2, and π, the distribution of a single observation is simply

$$Z \approx f(z \mid \mu_1, \sigma_1, \mu_2, \sigma_2, \pi) = z \;\; \pi\, n(z \mid \mu_1, \sigma_1) + (1 - \pi)\, n(z \mid \mu_2, \sigma_2).$$

The parameters of this distribution can be easily and efficiently estimated by maximum likelihood methods. We have applied this distribution to the actual price history of gold. The relevant calculations appear in the Excel spreadsheet entitled “Golddist.xls.” This is not a very elegant or sophisticated implementation, because we developed it for a students in a finance course, but the ideas are all there. Let’s look at the results.

For those of you who would like to know more about our research on gold, and especially on the variance process, we invite you to peruse the research paper entitled “Gold: Volatility and Options, A Reconciliation.”

2. Transitions Between States.

Transitions Between High and Low Vol states of Gold.

Time series data drawn from a mixture distribution presents something of a logical puzzle. The periodic transitions between the states introduce serial dependence into the time series. This is not generally a serious estimation problem if you have long time series. Transitions raise another, more interesting question, which is how to estimate the transition probability. This can always be done, and done efficiently. That is the virtue of mixing Normals, that all maximum likelihood estimators are consistent and highly efficient. There is a rather nice estimator which does not involve maximum likelihood methods (the resulting estimator is still consistent and quite powerful) for the transition probability. It is illustrated in the Golddist spreadsheet, and a derivation and explanation of terms can be found in the note entitled "Transitions Between Volatility States." The proverbial bottom line is the estimator

The key is to select a magnitude cutoff parameter, B, which divides the sample into successive prices that differ by more than B in absolute value and successive price observations that differ by less than B in absolute value. We assume that we have already estimated the parameters of the distribution, and in particular the variances, σ_{HIGH} and σ_{LOW}. In this model we assume that the states have the same mean, which I will further assume to be Zero. Then define the tail probabilities

P_1 = Prob[abs(p(t) – p(t-1)) > B | σ_{HIGH}] and P_2 = Prob[abs(p(t) – p(t-1)) > B | σ_{LOW}]

Lastly, let A denote the frequency of repeated Large price changes (changes larger than B).

P = estimated transition probability from high to high vol = $\{A - [2 P_2\pi + P_1 - 2 P_1\pi] P_1\} / \pi (P_2 - P_1)^2$.

2. Bounded versus Linear Diffusion

Thus far we have talked about mixtures of fairly tame states – conditional distributions – all of which could for instance be Normal. By the same token, theoretical option values can be derived in a fairly straightforward manner if we ignore transitions between states, i.e. if the choice of state is once and for all. Option valuation with transition among states is a very important topic, but not one we will delve into here. In this section we will look elsewhere, broadening the family of states to include distributions with compact support. Because this is a rather tricky case, to get started we will focus on the problem of modeling options subject to a bounded diffusion.

We emphasize at the start that there is no theoretically arbitrage free pricing of options in this case, because option replication imposes the condition that the underlying price evolve as though it drifts are the risk-free rate. Obviously, when the price approaches the upper boundary it must violate this condition. Consequently this model is not exact, but the violation of the usual axioms of option valuation is not necessarily serious, as we shall see. The actual constraints that follow from bounding the distribution are actually stronger than this. It is necessary to assume that the stock has zero drift: that the expected return is equal to zero. While stringent, this is not an unnatural assumption in light of the fact that price is constrained to lie within a bounded interval.

Binomial Tree Computation.

This problem lends itself to a binomial tree expansion. Suppose that the price is bounded within a lower bound of A and an upper bound of B, so $A < p < B$. The calculations are actually carried out on a normalized variable x,

$$x = (p - A) / (B - A).$$

The up-step function, u(x), and the down-step function, d(x) are functions that map the unit interval to itself, and satisfy the further conditions

1. **$f \circ g = g \circ f = \exp\{2\mu t\} = 1$. It is here that we require the assumption that $\mu = 0$. This condition guarantees that the tree recombines.**

2. **$df/dx > 0 > d^2f/dx^2$.**

3. **$f(0) = 0$ and $f(1) = 1$.**

There are infinitely many functions of this type, but for present purposes the choice $f(x) = x^\alpha$ is very convenient, where α is a proper fraction. It follows that $g(x) = x^{1/\alpha}$. The reason why we must exclude a drift is that if the composition of f with g is not the identity, the functions f and g are not inverse to each other. It is not easy to find a commuting pair of functions which compose to $\exp(2\mu t)$ because f and g are highly nonlinear functions.

The implied risk neutral probability of an up move, starting from x_0, is

$$P = (x_0 \exp(rt) - g(x_0)) / (f(x_0) - g(x_0)).$$

A practical problem arises if x_0 is sufficiently close to unity, because then we can violate the condition that P be less than 1. There is no choice but to simply prune these branches by pushing the next step to unity, with probability one. There is, as we said above, no completely satisfying option model possible for a bounded diffusion because it must at some points violate the arbitrage conditions. Note that this difficulty cannot be resolved by any choice of a drift parameter.

One promising application of this model, in a context in which this theoretical blemish is not usually of practical significance, is to Estimate the boundary parameters A and B from actual option prices. The idea is to use the model to infer whether market prices imply either an upper or a lower boundary and to identify the most likely location of these boundaries if they exist.

4. Options of a Price Process that Exhibits Support and Resistance.

Lastly, we want to use the preceding model in a broader context, in which the price diffusion is either in a bounded state, along the lines just discussed, or a linear diffusion. In concept, this adds two further parameters beyond those of the bounded diffusion model. They are the volatility of the linear (i.e. loglinear) diffusion and the probability, π, of being in the bounded state. We are not prepared at this time to contemplate transitions between states. We will assume that whichever state the price process is in, it will stay in that state until the options expire.

As long as the π parameter is constant over time, it is not involved in the dynamic hedging of options, and it is accordingly a risk neutral probability. The value of an option is therefore simply the weighted average of the values implied by the two state-dependent models.

How is this model used? Support and Resistance Levels.

There is an extensive practitioner literature on support and resistance points in trading. Whether or not they are "real" we will leave it to another place and time to ponder at leisure. In any case, there are methodologies which propose to locate them. In any case however they are not impenetrable. They have to be understood in a conditional sense, as levels that attempt to restrain price moves. This is precisely the world with a bounded state and an unbounded state. There are key prices out there which will with some probability constrain the price evolution. In case the price does succeed in breaking out, it will be seen to evolve for a time with a volatility larger – perhaps far larger – than was typical of its process previous to breaking out. The reason is that when the state was unknown, its "volatility" was an average of the volatility of the bounded process – which is almost sure to be rather small – and the volatility of the linear diffusion, but once it has broken out, its volatility is the larger vol with certainty.

The most promising way to utilize this model is to use actual option prices to infer the parameters of this two-state process, and in particular to infer the support and resistance levels, A and B, and the vol of the linear diffusion. This last estimate is of particular interest, because one of the great limitations of the vol estimates from a Black-Scholes Model is that it vastly underestimates the volatility that will apply in the "tails." The conventional way to address this is to assume that price changes are leptokurtic, which they may well be, but this approach is more theoretically satisfying and is also entirely computable.

We mentioned that we can also back out estimates of the support level, A, and the resistance level, B, that are implied by the actual market prices of options. It is of interest to compare the parameters derived in this way against other estimates of these key price levels. The mixing probability provides an interesting further insight on this issue also. Any set of market prices of options will imply some A and B, but not all supports and resistances are equal. The mixing parameter quantifies whether the options are priced as though they are legitimate realities – whether they might actually restrain the price evolution – or are they nearly irrelevant – leaving the price evolution in a Black-Scholes world.