SPREAD THE LOAD

Using a two-factor approach to modelling spread options reveals the paradox of negative vegas. Mark Garman shows how this can help users gain a better understanding of these complex instruments

Spread options are quite common in markets where traders want to isolate basis risk – the pricing differential introduced because of differences in location, grade, timing or some other aspect of an underlying asset class.

For example, a corn trader in Nebraska wishing to hedge a local corn position may need to turn to the more liquid Chicago futures market; he is then rightly interested in the spread, or basis, between Nebraska and Chicago corn, which is due to location, grade and perhaps even temporal differences. A bond trader may be interested in the pricing spread between bonds in two different European countries.[1] In the energy market, there are crack spreads (price differentials between refined and unrefined products), Brent-WTI spreads (Brent crude oil versus West Texas Intermediate) and gasoline-heating oil spreads.

Conceptually, any two commodities can be used to form the price differential inherent in a spread, but in practice the more important spreads are between commodities that are closely related, via either demand substitution (for example, in the case of European bonds) or the potential for transformation (such as turning New York gold into London gold by shipping it there). This is important from the analytical viewpoint, since the prices of such closely related commodities are likely to be strongly positively correlated.

Along with basis risk has come the need for financial instruments to manage it. Preeminent among these is the spread option. The standard spread option grants the right to call or put the spread value against a fixed strike price. But our pricing models for the spread option must account for the fact that a spread is the difference between two prices. Early attempts to model the spread have considered the spread itself as an asset price. There are pros and cons to such an approach. Modern methods include true two-price (two-factor) modelling, with attendant complications and curiosities.

This article first examines the competing approaches to modelling spread options, and then discusses a paradoxical situation arising in the two-factor approach. It turns out that spread options can occasionally possess negative vegas; understanding the source of this behaviour may contribute substantially to gaining a more intuitive feeling for the complexity of spread options.

A simplistic approach to modelling the spread option is to consider the spread itself as an asset price, lognormally distributed, and apply the usual Black-Scholes type formulas. This implicitly assumes (because of the lognormal distribution) that the probability that the spread will ever become negative is nil. Unfortunately, negative spreads do occasionally appear in most markets. In addition, the lognormal assumption would further suggest that spread fluctuation sizes would increase for large spreads and decrease for small ones -- often an implication supported by neither evidence nor experience.

Further adverse implications accrue from employing the one-factor approach to spread option pricing. For example, the commodity convenience yield (or equivalently forward price) is a parameter demanded by standard one-factor option pricing models. But what is the convenience yield on the spread? This might be imputed from the forward prices of the two commodities involved in the spread, but unfortunately in the general case, this

The standard spread option grants the right to call or put the spread value against a fixed strike price. But our pricing models for the spread option must account for the fact that a spread is the difference between two prices

depends upon the maturity of the spread option, even in the case of constant individual commodity yields. As a matter of practice, we usually (and rather heroically) simply suppose the convenience yield on the spread is zero within the one-factor model.

The one-factor approach also implies that a single delta should govern spread option hedging activity. For example, if one is hedging a Brent-WTI spread option, it may not seem unreasonable that the appropriate hedge would match each long barrel of Brent with a short barrel of WTI. However, it is clear that in some situations the one-delta implication will be quite outrageous. If the spread comprised a high-volatility asset with a price of five and a low-volatility asset with price one (and therefore a spread of four), it is apparent that the appropriate hedge must seriously differentiate between the two assets, and no single delta would suffice.

Attempts have been made to overcome some of the limitations of the one-factor approach, for example by assuming normal distributions, or by adding a constant amount to the spread and assuming the consequent sum is lognormally distributed. But no one-factor approach can overcome the one-delta limitation. The real solution here is to move up to the true two-factor approach. In this method, the arithmetic difference between two lognormally distributed asset prices is defined as the spread. With the two-factor approach

continued on page 84

[1] Bond traders may also be interested in yield spreads between bonds issued in different currencies. While yield spreads and price spreads are sometimes (but not always) treated via quite different models, the same caveats regarding one-factor versus two-factor approaches and resultant negative vegas generally apply to yield spread environments as well

Mark Garman is professor of finance at the Haas Schools of Business University of California, Berkeley, where he conducts research in financial engineering. He also serves as chief scientist to Financial Engineering Associates, where he has designed risk management software that includes @Global and DerivaTool. He has provided financial institutions with consultancy and training in the areas of second-generation options and advanced risk management systems

1. Probability mass function for equal volatility

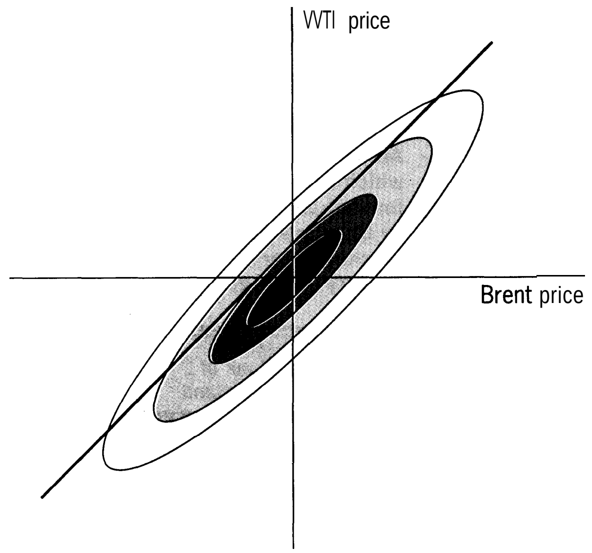

2. Probability mass function for reduced Brent volatility

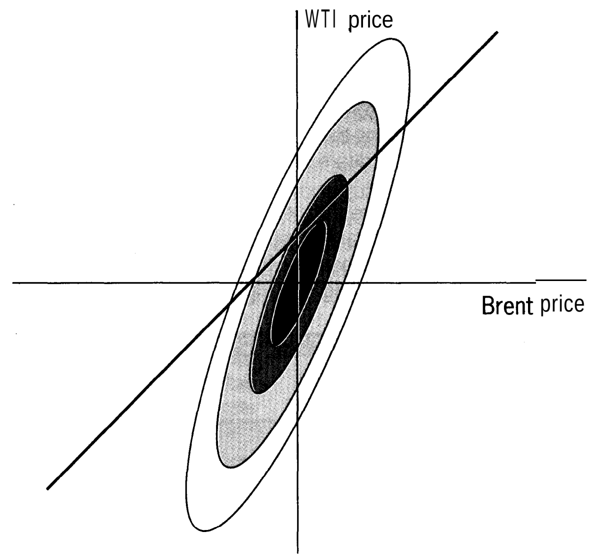

continued from page 68

comes new insights, coupled with fresh complications.

The input parameters of the two-factor spread option pricing model[2] include at least the following: the volatilities of both assets involved in the spread, the correlation coefficient between them, and the convenience (or other) yields of each asset taken individually. All of these are familiar option pricing parameters, with the possible exception of the correlation coefficient. The outputs of the two-factor spread option pricing model also show additional complexity: two deltas, multiple gammas and at least two vegas Such complexity of structure is an unavoidable consequence of the two-factor model; in particular, understanding the relationship between volatilities, correlation and vegas leads to insights in pricing spread options. We begin the discussion with a rather strange phenomenon.

Surprisingly, in many of the common cases encountered, spread options may have negative vegas.[3] This means that lower volatilities can in some circumstances produce higher option prices! Such a phenomenon is virtually unknown in any other form of option pricing, where positive vegas are the rule.

Consider the case of an out-of-the-money Brent-WTI spread option. The individual volatilities of Brent and WTI are likely to be about the same. There is also likely to be a high correlation coefficient in this case, since common market factors drive world oil prices. The payoffs to an out-of-the-money spread option can in part be understood by examining figure 1.

The two axes represent the future prices of WTI and Brent. The axes intersect at the current prices of the two commodities. Only those future price pairs which lie above and to the left of the in-the-money line, shown as a heavy line drawn at a 45° angle, will result in payoff to the spread option. The ovals represent the probability density function of the future price pairs. The darker the oval, the more probability mass is present. The in-the-money line is drawn at a 45° angle as a consequence of the quantity ratio of the spread: one barrel of WTI versus one barrel of Brent. The fact that the axis of the probability density function is also at roughly a 45° angle is a consequence of the approximate equality of the individual volatilities of WTI versus Brent. Finally, the concentration of probability mass along the latter axis is a consequence of the assumed high correlation between the two commodities. Now the value of the spread option is generally related to the amount of probability mass to be found above and to the left of the in-the-money line.[4] But because the in-the-money line and the probability density function axis are roughly colinear, coupled with attendant high correlation, we observe that very little probability mass actually appears in this region.

Next, let us suppose that, ceferis paribus, the volatility of Brent declines. In other words, we hold constant the WTI volatility and the Brent-WTI correlation, while reducing the Brent volatility. The corresponding diagram will be as in figure 2. Notice that considerably more probability mass has been moved into the in-the-money region. This represents the cases where an increase in the spread eventuates when the WTI price increases, since the highly correlated Brent cannot "keep up the pace" due to its diminished volatility. Thus we observe an increase in the spread option's value, although we have reduced the Brent volatility – that is, encountered a negative Brent vega. The same effect can occur when WTI volatility is reduced with other variables held constant – a negative WTI vega – and indeed both negative vegas may be simultaneously present.

Spread options are non-trivial instruments. Traders of these instruments are called upon to manage multiple price risks, an array of gamma risks, correlation risk and a host of similar challenges. Employing the one-factor spread option modelling approach does not remove any of the inherent complexities, it only obscures them. Viewed in this light, the one-factor approach to spread option modelling is, to be frank, the ostrich or "head-in-the-sand," solution: convenient, but not illuminating. It is surely better to face the complications and accompanying paradoxes of the two-factor approach directly. As we have seen, one of these paradoxes is that of negative vegas. Moreover, the negative vegas occur in precisely the same circumstances that realistic spread markets themselves present: high correlation between two asset prices where the volatility ratios are roughly similar to the spread quantity ratios – normally one-to-one. In specifics, traders may wish to evaluate spread options for a range of correlations rather than for a single such value. In general, awareness of such paradoxes can only enhance trading prowess.

2 See, for example, Boyle, Evnine, and Gibbs, 1989, Numerical evaluation of multivariate contingent claims, Review of Financial Studies 2, 2

3 I am grateful to Mark Vonderheide of Vitol for first raising this issue

4 Strictly speaking, the spread option value will equal the integral of the risk-neutral equivalent probability mass multiplied by the distance within the in-the-money region, measured from the in-the-money line