Future Risk

Christopher C. Finger
chris.finger@riskmetrics.com

Luis O'Shea
luis.oshea@riskmetrics.com

February 2006

In November, we discussed risk modeling of credit spreads. We raised two broad questions. First, we asked which market should we look to for information. And when credit is traded in more than one market, should we choose the one with the greatest liquidity, or the one that most closely matches our position? Second, we asked what made a time series useful as a risk factor, and whether we could choose among a variety of definitions of spread to obtain the best properties for forecasting purposes. A third question we could have asked, but did not, was how we should model the volatility once we had obtained a useful time series. We picked up that question in our December note.

In this note, we ask the first two questions again, but for futures contracts rather than credit spreads. As we will discuss, there are modeling choices we have applied for a long time which, while serving us well broadly, are in fact questionable in specific cases. Moreover, it is never a bad thing to return to models that have been around a while, and revisit the thinking that led us to those choices in the past.

Financial futures

There are a myriad ways to partition the many futures contracts. For our purposes, we first make the distinction between financial futures—those futures contracts that entail the delivery of what is essentially an investment product (such as a bond or equity index)— and non-financial futures—those that entail delivery of something (such as energy or food) whose primary purpose is to be consumed.

With financial futures, the fact that the underlying deliverable is an investment product means that we have a good sense of the benefits that accrue to the holder of the deliverable. A bond pays a coupon, for instance, while an equity index pays a dividend. Moreover, we may observe prices of bonds of various maturities, from which we can build a discount curve, and have an idea of the market's expectation of how a particular bond will evolve through time.

It is thus possible to build a native model for financial futures—a model that uses arbitrage arguments to connect prices in the deliverable securities to prices of the futures contracts. This connection between the cash and futures markets gives us a first guess at an independent futures price, but more importantly allows us to utilize risk forecasts on the deliverables to also forecast risk for the futures contracts. Moreover, the connection allows us to articulate sensitivity measures and stress tests consistently across cash and futures positions; it is natural to shift the interest rate curve by some amount and reprice both types of positions.

While the native model is attractive for its consistency and its reduced set of risk factors (one interest rate curve is used to model both cash and futures positions), it is incomplete, in that there is residual risk in the futures that is not described by the native model; there is a well known basis between the forward prices inferred from the cash market and the actual futures markets, and this basis in itself is risky. Some of the basis can be explained by the fact that the cash and futures markets are not perfect substitutes for each other.[1] On the other hand, the two markets are distinct, with different participants, different liquidity, and different supply and demand effects; thus some portion of the basis is purely technical.

The discussion of the adequacy of the native model is similar to our discussion of spread risk modeling across bond and credit default positions. An investor who above all takes long positions in Treasury bonds, and sees the cash and futures markets as equivalent ways to express interest rate views, should opt for model parsimony and risk forecasts based on the best available source of information.

For these investors, the most common issue we confront is whether futures positions imply increased leverage, and therefore increased risk, at least relative to the amount of capital. Again, this depends on the investor: sometimes investors take on futures as an efficient way of gaining a leveraged exposure; other investors treat futures as a substitute for a direct investment in the underlying, and match their futures position with a cash position in order to avoid increasing leverage.

Investors actually trading the basis, through long bond and short futures positions, cannot afford the luxury of model parsimony, as their primary risk is to the factor that the pure native model ignores. For this class of investors, the natural extension is to treat the basis itself as an explicit risk factor, while maintaining the consistency that the native model affords.

Non-financial futures

As we move away from futures on financial securities, and into the realm of energy, metals, and agricultural products, we are forced to make a significant change in our modeling approach. For these futures contracts, there are, for the most part, limited cash markets from which to infer prices; the futures contracts themselves are our primary (if not only) source of market information about prices of the commodity at any time other than the present.

With bonds, we can observe with certainty the cost of financing a cash position, the benefit (in terms of earned interest) of owning the bond outright, and the price of bonds of slightly longer or shorter maturity; all of this information goes into the arbitrage argument and leads to at least a good starting point for the price of a bond futures contract. For non-financial futures, one approach is to apply a similar logic, examining the differences between owning a physical asset and owning a futures position, and using these differences to price the future. These differences include the cost of storing and transporting the commodity, both of which the owner of the physical asset must bear, as well as the benefit (referred to sometimes as the convenience yield) of owning the asset,

[1] For instance, the Treasury bond futures contracts include a cheapest-to-deliver provision, in which the holder of the future may choose to deliver the cheapest of a defined set of bonds; clearly, the owner of an actual bond position does not have that option.

such as having the flexibility to consume or sell the asset at any time.

Unfortunately, these quantities (at least other than the storage cost) are all much more difficult to quantify than their analogues for bonds. In the end, models that apply the convenience yield logic are typically run in reverse: we observe the futures and spot[2] prices, and infer the convenience yield such that our model recovers these prices. We may try to make inferences about the convenience yield, but this is a quantity that is inferred from the futures prices, rather than a fundamental value in itself.

Risk modeling

Where this leads us, from a risk perspective, is to modeling based on the time series of the futures contracts themselves, rather than (wholly or partly) on the deliverable instruments. Simple enough, we might think: we hold positions on individual futures contracts, so we simply model risk based on the historical price series of those contracts. In other words, if we hold the June 2006 crude oil contract, we look at the historical price fluctuations on this specific contract, and make forecasts of how it is likely to move over the next days or weeks.

This approach is in fact quite sensible, and often applied in practice. For risk forecasting, though, it suffers from two drawbacks. The first is practical: for new contracts, by definition, there is no price history. This implies the need to use proxies, or else simply wait until enough history is available.

The second drawback is more subtle. It is not clear that the historical price moves in the June 2006 oil contract are the best source of information to forecast fluctuations in our position in the contract as of today. This issue is not unique to futures. For bonds, we recognize that the volatility in general decreases as the bond ages (and its duration decreases); thus, rather than forecasting based on the bond price history, we create a time series of constant maturity interest rates, and forecast based on these. We must at least question, then, for futures, whether we can do better than simply forecasting based on the price history of specific contracts.

We should remark that even under the relatively simple historical simulation framework, we must confront this issue, and specify how we interpret and apply historical price changes. Consider a risk forecast, performed on February 14, 2006, for a position in the June 2006 crude oil contract. Under historical simulation, we apply the price changes that occurred over the prior year. One data point consists of the price changes on March 23, 2005: the June 2006 contract fell by 1.19, while the August 2005 contract (which had a similar time to expiry on March 23 as the June 2006 contract has as of our analysis) fell by 1.74. Which of these two price changes (which differ by 50%) is most appropriate to apply? Essentially, this is the same question we have stated above.

Crude oil and heating oil

Complicating this picture further is the fact that not all commodities behave similarly. To illustrate this, consider heating oil and crude oil. Both are subject to the same broad supply dynamics; the demand for crude oil, though, is roughly constant throughout a given year, while the demand for heating oil is clearly

[2] That is, the price for immediate (or an approximation thereof) delivery

Figure 1: Price versus time to expiry for crude oil and heating oil futures contracts. Data observed in June (dotted lines) and December (solid lines)

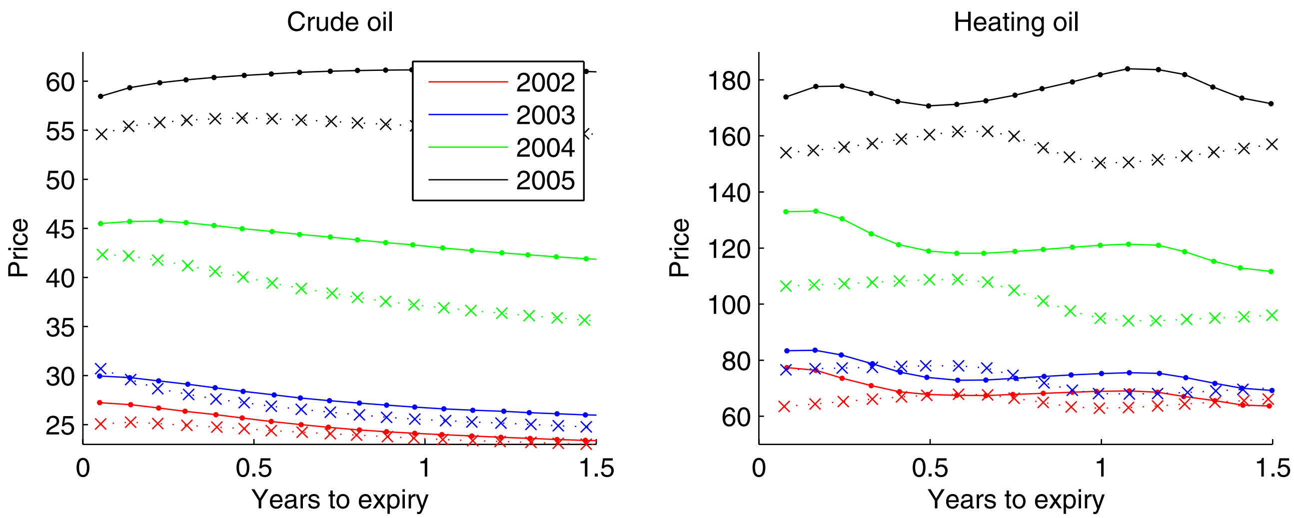

higher in the winter than the summer.

The impact of this difference on the relative price dynamics of the two commodities is evident in Figure 1. On the left, we plot the prices of crude oil futures against time to expiry; on the right, we do the same for heating oil. For crude oil, the shape of the crude oil curves evolve slowly in time, moving from the classic backwardation (that is, downward sloping) profile to contango (upward sloping), but there appears to be no regular pattern to the evolution.

For heating oil, though the term structures follow the same general upward or downward sloping profiles as for crude oil, the most noticeable characteristic is that the profiles are best characterized by the month in which they are observed. Observed in June, the prices of futures expiring in six months time are higher than those of futures expiring in one or twelve months; observed in December, prices are higher for twelve- than for six-month futures. This behavior is not random, but a regular pattern, and entirely consistent with our intuition about demand. We say that heating oil prices exhibit seasonality, while crude oil prices, at least based on these figures, do not.

Sources of volatility

Forecasting volatility for non-financial futures is complicated by the fact that volatility comes from three broad, but distinct, sources: overall market fluctuations, the maturity effect, and seasonality. The market fluctuations are common to all financial time series, and modeling of this phenomenon was the subject of our December note. The other two effects are particular to commodities.

The maturity effect is a decreasing pattern in volatility with respect to time to expiry. The effect was a hypothesis of Samuelson (1965), which went against the standard models of the time, but which has since been validated empirically and incorporated into commodities models.

Seasonality in volatility is a pattern in which contracts with expiry at particular times of the year are inherently more volatile than those with expiry at other times. Importantly, seasonality in volatility is not necessarily the same effect as the seasonality in prices discussed previously. It is reasonable to expect that the two might go hand in hand, but this is not imperative; moreover, the two effects have different implications on modeling.

The appropriate risk model, and in particular, the appropriate time series from which to base forecasts, for a particular commodity depends on the relative importance of the three sources of volatility, as well as on the characteristics of the prices themselves. In order to investigate the volatility sources, we examine the volatility computed over short (50 day) periods for individual commodity futures contracts.[3] We examine the term structures of volatility observed at different points in time.

As we observe the volatility term structures, their levels will certainly fluctuate; this is the market effect. Beyond this, we are interested in the shapes of the term structures. Under the presence of the maturity effect and no seasonality, the volatility term structures should display a downward slope. The shape should evolve, but there should be no specific signature of the term structure in any specific month.

On the other hand, under the presence of seasonality but no maturity effect, we should see patterns in which the volatility term structures observed in June of different years should be similar to each other and distinct from the term structures observed, in December. Moreover, if it is the case that December contracts are inherently more volatile than June contracts, then observing volatilities in December likely produces a downward sloping term structure (from which we might erroneously identify the maturity effect). Observing volatilities in June, though, should produce the opposite.

We plot the volatility term structures of crude oil contracts in Figure 2 and heating oil contracts in Figure 3. Crude oil exhibits the classic maturity effect. The effect was most pronounced in the months after September 11, 2001, and is less pronounced recently. On average, we see that there is roughly a ten percent difference in volatility between nearby and one-year contracts. At typical levels (30%), this difference could account for one third of a contract's total volatility.[4] And as with crude oil prices, there is no systematic pattern to the term structures, but rather a slow evolution through time.

For heating oil, whereas for prices, we saw strong seasonality coupled with a weak tendency toward backwardation or contango, we see for volatility a strong maturity effect coupled with weak seasonality. The maturity effect has comparable magnitude as for crude oil, with a one year difference in time to expiry accounting for roughly a ten percent difference in annualized volatility. We can also see a weak seasonality: the term structures observed in June and September decrease monotonically, while those observed in March and December decrease until about one year to expiry, and then turn up. This signature of the term structures appears persistent across the six years presented in the figure.

[3] These figures should not be interpreted as forecasts, but simply as characterizations of how volatile prices are at a point in time.

[4] In fact, the ten percent difference understates the problem, since we are comparing volatilities realized over periods (two months) that are far from short relative to the one year over which we are making the comparison. This actually smoothes the effect somewhat.

Figure 2: Crude oil futures contracts. Volatility (past 50 days) versus time to expiry

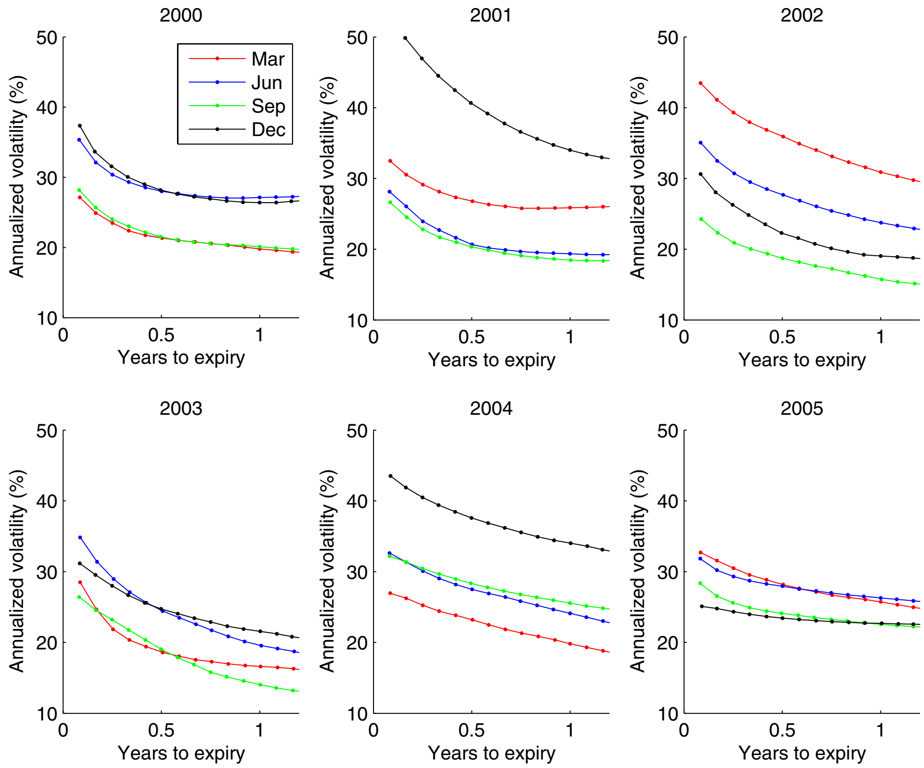

Figure 3: Heating oil futures contracts. Volatility (past 50 days) versus time to expiry

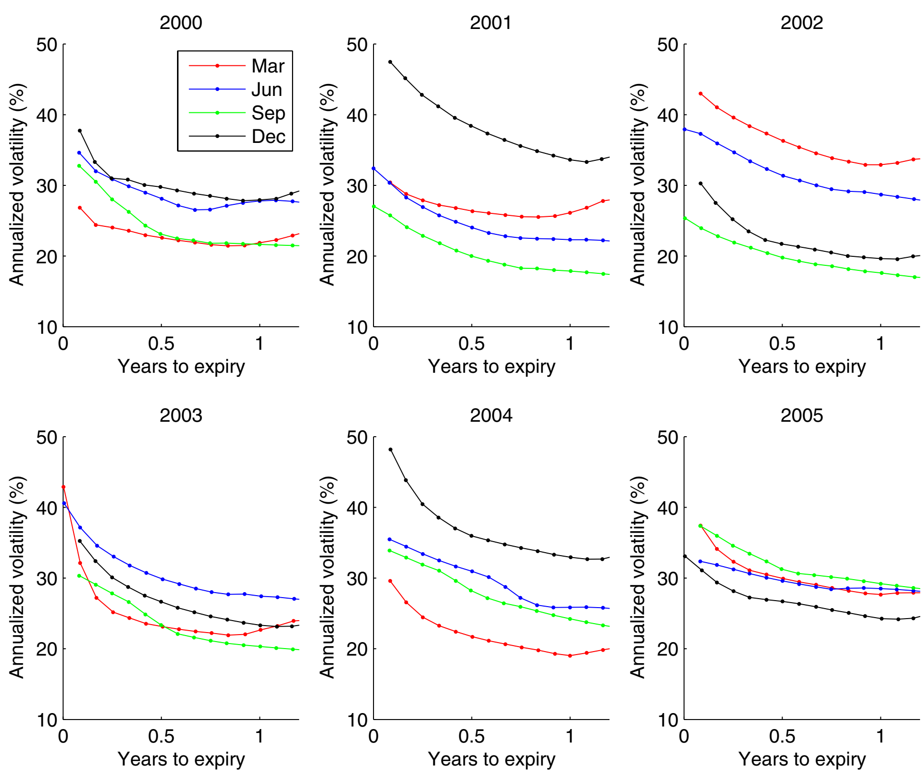

Forecasting implications

Recall that our main goal here is to forecast the range of likely price moves for a specific futures contract, whether by volatility forecasting or by historical simulation. In either case, the prominent maturity effect is of primary concern. Under historical simulation, we conclude that the most relevant price changes to apply are those that come from historical changes on a contract with similar maturity to what we hold. In our earlier example, we were forecasting, as of February 2006, the price changes in the June 2006 contract. From the March 23, 2005 price data, the most relevant price change to apply is not that from the June 2006 contract itself (which at the time of the data had fifteen months to expiry) but rather the larger price change from the August 2005 contract (which at the time had roughly four months to expiry).

For more formal volatility forecasting, the problem is a bit more subtle, but the intuition is the same: we should base our forecasts on the most relevant set of historical returns. Standard volatility forecasting methods are intended to capture market fluctuations, but typically assume no systematic evolution of the volatility. In our case, in addition to the market fluctuations, we have two systematic effects: the maturity effect, which implies that contracts become more volatile as they age, and seasonality in volatility, which implies that contracts expiring in some parts of the year are systematically more volatile than others.

Forecasting volatility based on a weighted average of past squared returns, we encounter the problem that some of the return data derives from a time period in which the contract had inherently lower volatility. We could attempt to address this through the weighting scheme; its purpose, however, is to utilize the historical data optimally to capture market fluctuations, not to handle systematic effects. To rectify this problem, we have two options: explicitly incorporate the maturity effect and seasonality into our volatility forecasts,[5] or transform our price data into something that presents only the market fluctuations in volatility, and neither of the systematic effects.

Price transformations

At this point, we introduce another modeling constraint: we would like, to the extent possible, to apply a consistent modeling approach across all asset classes. This steers us to the price transformation approach. Further, this approach resolves our issues with historical simulation, in that it provides a specific set of shocks to apply. The key, of course, is how we transform the prices. Our goal is to obtain a transformation that accounts for the maturity effect, is robust to commodities with strong price seasonality, and is amenable to standard volatility forecasting schemes; we will not address volatility seasonality for now, as it appears to be of secondary importance.

An obvious option is to use time series of "constant nearby-ness". For example, we create a time series of returns on the "first nearby" contract by taking each day the return on whichever is the nearest contract to expiry at that time; similarly, we create return time series for the second, third, and so forth nearby contracts. With this transformation, our volatility forecasts are based on historical returns on contracts

[5] For example, we could fit a long-run average level of volatility as a function of time to expiry and delivery month of the contract, and then use more standard volatility forecasting technology to model deviations from this average level.

of similar, though not precisely the same, time to expiry. To the extent that the maturity effect is very steep, then the reduction of time to expiry within a specific month could pose problems; and if one contract behaves differently from the next, then the discrete jump across contracts may also give us difficulties.

To move beyond this first approach, we must recognize a separation of the notions of risk and valuation; the value of a position in a specific contract clearly comes from the price of that contract, while the risk on the position may well be characterized by a historical data other than that contract's prices.

A second transformation approach (one that we have applied historically) is to calculate prices on hypothetical constant maturity futures contracts: on each day, for example, we calculate a one month futures price by interpolating between the two contracts that straddle one month to expiry. This approach addresses the maturity effect more precisely, in that the returns we observe always come from (albeit hypothetical) contracts with the same time to expiry, and ensures that the move from one contract to the next within the time series is smooth. As with the prior approach, though, this time series does not represent an investable strategy.

Of greater concern, this approach leads to biased forecasts in the presence of price seasonality. Consider the hypothetical constant six-month maturity heating oil futures price. From Figure 1, we know that this price will increase from December (when the six-month maturity contract is the June contract) to June (when the six-month maturity contract is the December contract). Standard volatility forecasts utilize a weighted sum of past squared returns, with the assumption that the mean return (that is, the drift in prices) is negligible. Here, the price drift is certainly not negligible, and thus leads to overstated squared returns, and consequently a biased volatility forecast.

What is most desirable is to maintain the constant maturity aspect of the previous approach (and thus address the maturity effect), but to avoid the bias from calculating returns across a series of prices that is systematically increasing. We can achieve this by reversing the order of the averaging (that is, the interpolation across two futures contracts) and the return calculation. Rather than averaging first, and then computing returns on hypothetical (and biased) price series, we may compute returns first (on the actual contracts) and then average to maintain a constant effective maturity.

This logic leads us to the notion of an investable constant maturity strategy. To do this, we create a one-month index as follows: at the beginning of the time series, we invest a total of \$100 in the two nearest contracts, in such proportion that the average time to expiry of the total position is one month; each subsequent day, we sell some of the shorter contract and buy some of the longer one in order to maintain the one-month average maturity. The returns on this index address the maturity effect, but should not have the same structural drift as the prior approach for commodities with strong price seasonality.

Ultimately, the best transformation is the one that leads to the best risk forecasts. The investable indices have no obvious deficiencies, but empirical tests are needed. Furthermore, it is not clear that any of these approaches would perform well in the presence of strong seasonality of volatility. Whether there is a transformation that addresses this feature is a topic for further research.

Further reading

- Finger, C. (2005). Spread values, Research Monthly, November.
- Finger, C. and Zumbach, G. (2005). Forecasating for solvency risk, Research Monthly, December.
- Finnegan, J. (2004). Back to basics: Commodity futures pricing—contango and backwardation. *Financial Engineering News*, January/February.
- Gibson, R. and Schwartz, E. (1990). Stochastic convenience yield and the pricing of oil contingent claims. *Journal of Finance*, **45**(3): 959–976.
- Liu, N. (2005). What's new in futures, Presentation, RiskMetrics Group European Client Conference.