TechHackers' Quant*News*

December 17, 1998

Welcome to the second issue of Quant*News*.

PIK bonds

Payment In Kind (PIK) bonds are bonds which pay a portion of their interest via the issuance of new bonds, rather than the payment of a cash coupon. They are a cross between ordinary coupon bonds and a zero coupon bonds[1], paying out some interest but having a significant fraction of their "value" in the increased principal amount.

Consider a bond which pays annually, with a 2% cash coupon and a 10% PIK coupon. Based on $100 par value, on the first coupon date the holder will receive $2.00 cash and an additional $10 par value worth of bonds. Then on the second coupon date the holder will receive 2%×$110 = $2.20 cash and 10%×$110 = $11 additional par value worth of bonds. This continues until maturity (or redemption).

Usually PIK bonds will have a changing coupon schedule. Early in the life of the bond most of the interest is payed in kind, with a larger proportion being paid in cash later on.

[1] PIK bonds may have a tax advantage over zero coupon bonds in that their "income" will be taxed at current market value, rather than via a pre-defined accrual (OID) schedule

Analytics for PIK bonds are quite similar to that for ordinary bonds—in fact, once one has generated the cash flows, the formulae are the same. One just needs to adjust the size of the cash flows to account for the PIK feature.

QuantTools / ***@nalyst*** will support PIK bonds beginning with the version 99 release.

Non-Deliverable Forwards

A forward contract is a contract to buy (or sell) an asset at a particular price at a particular future time. For example, I could contract now to have 400 gallons of heating oil delivered to my house on December 21 at a price of $0.97 / gallon.

Of course I'd be in trouble if I did this, as there is no oil tank at my house, and I'd rather not fill the swimming pool with number 2 oil. But suppose I'd still like to lock in the $0.97 / gallon, as I'm planning on buying a lot of gasoline, plastic, and other commodities whose prices are closely linked to that of oil. What can I do?

I can enter into a *non-deliverable forward* contract. This is a contract that is *cash settled*, i.e. instead of receiving the oil, and paying $0.97 / gallon, I simply receive (or pay) the current spot price of oil minus 0.97. Note that this payment is the same as the terminal value of an ordinary forward contract, and thus both ordinary and non-deliverable forwards are priced the same: the value of such contracts is just the difference between the current market forward price and the locked-in forward price, discounted back to the current time.

In order for non-deliverable forwards to "work," there needs to be a recognized reference price available, so that the contract can determine an accurate settlement price[2].

Non-deliverable forwards have recently become popular in foreign currency markets as a way of escaping currency controls. Rather than taking delivery of a non-freely-convertible currency, one takes or makes a dollar payment representing the appreciation or depreciation of that currency with respect to the dollar.

Improving Monte Carlo

In the last issue Maurizio Mondello wrote about Quasi-Monte Carlo. These new approaches do not, however, replace "ordinary" Monte Carlo, particularly if one remembers some simple methods for improving the convergence properties of Monte Carlo methods.

Antithetic Variables. The use of antithetic variables means that we choose our Monte Carlo paths in pairs: if we have a sample path S_t, we should also price the option on the sample path $\hat{S}_t$ so that $(S_t\hat{S}_t)^{1/2} = S_{\text{forward}}e^{-v^2t/2}$. For example, if the forward price of an asset is 20 and one path puts us at the asset price 25, we should also consider the path that puts as at the price 16.

What do we gain by this? First of all, if generating the random values is "expensive", we produce twice as many samples with little effort. We've also guaranteed that the (geometric) mean of our sample distribution matches the forward price exactly, and thus we have eliminated all errors from the leading term in Taylor-like expansion.

This method is also useful in a multivariable setting. For example, if we have two assets, we can consider paths 4 at a time: (S_1, S_2), $(\hat{S}_1, S_2)$, $(S_1, \hat{S}_2)$, $(\hat{S}_1, \hat{S}_2)$.

Control Variates. Suppose we want to compute $E[A(x)]$, and we have an exact formula for $E[A_0(x)]$. Then we can write $A(x) = [A(x) - A_0(x)] + A_0(x)$ and thus

$$E[A(x)] = E[A(x) - A_0(x)] + E[A_0(x)].$$

If A_0 is a good approximation for A, the variance of $A(x) - A_0(x)$ will be much smaller than that of $A(x)$ and the simulation will converge much more rapidly.

One common example where this is used is in the pricing of Asian options. Here $A(x)$ is the option value using arithmetic averaging (where there is no closed form solution), and $A_0(x)$ is the option value using geometric averaging (where a closed form solution exists).

Change of Measure. When using Monte Carlo to evaluate an option that is far out-of-the money, most of the terms will be zero.

[2] The reference source should not be subject to excessive manipulation. Unfortunately many of the markets where a NDF is most useful are subject to such manipulation: e.g. the ruble rose to 7.5 to the dollar on Sept. 15, 1998, dropping back to 13.5 to the dollar the next day; Sept. 15 was the expiration date of about $2 billion dollars worth of NDF contracts. NDFs are also more susceptible to political risk than ordinary forwards, since they may be pegged to an "official" artificial rate rather than a true (black-market) rate.

Thus, most of the computation is a waste. Changing the probability measure allows most have the samples to have a nontrivial contribution; a multiplicative factor then adjusts for the change of measure.

For example, suppose the risk neutral measure has drift μ. Instead of evaluating $E_\mu[F(x)]$, one can evaluate

$$E_{\mu+m}[e^{-m^2t/2}F(x)e^{-mx}]$$

which is equivalent according to Girsanov's theorem.

By choosing $\mu + m$ so that the process drifts to the strike price K, i.e. so that $\mu + m = [\log(K/S_0)]/t$, where S_0 is the starting asset price and t is the time to expiration, we will be sure that our sample paths visit the "interesting" part of the distribution.

The Problem With Credit Derivatives

The problem with credit derivatives is that there are two methods of pricing them, and they're both wrong.

Using bond prices. The simplest method of pricing credit derivatives is to use the difference between prices of bonds of different levels of credit-worthiness as a judge of the value of a credit derivative. For example: consider two zero coupon bonds, both maturing in one year. One is a U.S. Treasury issue, considered immune from default, and the other is a corporate issue with a given credit rating. One might say that the Treasury bond is just a corporate bond with the addition of a credit put which pays off the value of the bond in case of default. Thus, the value of the put = value of Treasury bond − value of corporate bond.

Assume the Treasury is yielding 4.5% and the corporate is yielding 5.3%. A little use of our favorite software package (***@nalyst***) turns those yields into prices of $95.647 and $94.903; the difference is $0.744. Thus one would say the credit put insuring against default is worth $0.744 — **WRONG!**

The difference in price between the two bonds is a result of many different factors, including liquidity risk, tax preferences, regulatory requirements, along with credit risk. Given that Treasuries fare favorably against corporates on all these fronts, the $0.744 represents only an upper bound for the price of the credit put, and probably not a particularly tight bound.

Using historical rating information. Another approach to pricing credit derivatives makes use of the historically known probabilities for credit rating upgrades and downgrades. In other words, we have a state space consisting of the different possible ratings a bond can have, and assume that we have a stochastic process acting on that space, where the transition probabilities are (say) the observed historical one-year transition rates. For example, historically, A rated bonds, after one year, have ended up with ratings according to the following table:

AAA	.0010
AA	.0234
A	.9154
BBB	.05082
BB	.00614
B	.00264
CCC	.0001
default	.0005

Thus there is a 0.05% chance of the A bond defaulting after one year. Default rates for longer periods can be found by raising the transition matrix (the above table represents one row of that table) to the appropriate power.

There are two problems with this approach. First of all, the true stochastic process needn't be Markovian, so that the probabilities might be path-dependent. Related to this is the fact that the probabilities might be inconsistent with shorter time periods.

Secondly, and more importantly, knowing the default probabilities doesn't allow you to compute the value of credit derivatives. One needs the *risk-neutral* default probabilities, not the "real-world" historical probabilities (just as when computing Black-Scholes option prices the implied volatility is often significantly different from the historical volatility).

If the market is risk-averse, then the risk-neutral default probabilities will be higher than those observed historically, resulting in higher prices for credit puts. This is the current situation, with bonds trading at prices that reflect unrealisticly high chances of default. On the other hand, as of several months ago, the market was risk-seeking, with, for example, one-year bonds that historically had a 20% chance of defaulting trading with a price only $3 below one year Treasury zeroes.

Hedging: How To Benefit Stockholders Rather Than Bondholders

When a company hedges its exposure to risk, who benefits? Typically, it is the bondholders who benefit, at the expense of the stockholders. Certain types of innovative financing structures can turn this around and allow the stockholders to benefit. A detailed example is given in "Credit enhancement through targeted risk management: Freeport-McMoRan's gold-denominated depositary shares," N.K. Chidambaran, C.S. Fernando and P.A. Spindt.[3]

First of all, what does hedging typically accomplish? One obtains protection against a "bad event", but in return, must pay a premium or otherwise give up some of the upside potential of a possible "good event". Protection against the bad event helps both stockholders and bondholders, but the cost is incurred only by the stockholders. Thus bondholders benefit and stockholders lose (on a risk-neutral basis).

Since the bondholders benefit from the company's hedging activities, the trick is to get them to pay for it. One can't coerce the existing bondholders, but one can structure a new debt offering with a hedge built in. In the Freeport-McMoRan case, the structure was a gold-denominated bond, which is effectively a combination of a bond and a long forward position in gold. The combination of the bond and forward position sold at a higher price as a combined asset than it would have sold as two separate positions since Freeport's short forward position enhanced the credit-worthiness of the bonds. This additional value went to the stockholders of the company. If they had first sold bonds, and then hedged, the value would have gone to the (new and old) bondholders. If they had first hedged, and then sold bonds, the value would have gone to the senior bondholders.

Amazingly, Freeport must have stumbled in to this type of structured financing without understanding it—in addition to the gold-denominated notes, they also issued silver-denominated notes. The gold-denominated notes were an effective risk management tool as Freeport had gold backing up its hedge. Since they did not have significant silver reserves, the silver-denominated notes actually

[3] in the Proceedings of the 1998 Conference of the International Association of Financial Engineers

increased their risk exposure and resulted in a financing cost higher than that which would have been obtained from issuing “ordinary” bonds.

A Brief Introduction To The BGM Model

What follows is a synopsis of a talk given by Sven Sandow at the TechHackers seminar on October 15, 1998.

The Brace, Gatarek and Musiela (BGM) model is a model for the time evolution of forward interest rates. The basic variables of the model are the forward rates for a discrete set of time periods: for example, one might model the rates $f_t(k\tau, [k+1]\tau)$ for positive integer k, which are the forward rates measured at time t applicable to the time period beginning at $k\tau$ and ending at $[k+1]\tau$, where τ is, for example, 3 months. Thus one does not attempt to model the entire yield curve structure, but only a subset corresponding to observables at discrete times.

Identifying $f_t^{(k)} = f_t(k\tau, [k+1]\tau)$, the time evolution of the rates is given by

$$\frac{df_t^{(k)}}{f_t^{(k)}} = \mu_t^{(k)} dt + \sum_{j=1}^{M} \sigma_t^{(k,j)} dz_t^{(j)}$$

where the M factors $dz_t^{(j)}$ are independent Brownian increments. If one integrates the above stochastic differential equation, one sees that the forward rates have correlated lognormal distributions.

The coefficients $\mu_t^{(k)}$ and $\sigma_t^{(k,j)}$ are not determined independently but are related by a no-arbitrage relationship:

$$\mu_t^{(m)} = \tau \sum_{k:k\tau>t}^{m} \sum_{j=1}^{M} \frac{\sigma_t^{(m,j)} \sigma_t^{(k,j)} f_t^{(k)}}{1 + \tau f_t^{(k)}}.$$

In order to use the model for pricing, one must now choose to model the time dependence of the volatilities (e.g. piecewise constant in time). Then one calibrates the model against the market, and proceeds to price assets using, for example, Monte Carlo simulations.

OAS

What is OAS?

OAS stands for Option Adjusted Spread.

OAS a a mechanism for comparing the yields of of financial assets that have embedded options. For example, if a Treasury bond yields 5% and a mortgage-backed security yields 6%, the spread is 1%—but the OAS will generally be (significantly) less, as much of the additional (apparent) yield of the MBS represents the value of the prepayment option built-in to the security.

How is OAS computed? Here we will neglect issues involving the yield curve, and only consider the simplest case.

Suppose the 6% MBS is priced at 102. If an analysis of the prepayment option indicates that it is valued at 1.50, then the MBS, if stripped of its option, would be worth 103.50. At this price one computes the yield to be 5.25%, and thus the OAS is 0.25% or 25 basis points.

TechHacker's Street**Math**.com computes OAS for CMOs and other mortgage-backed securities (using a full treatment of the yield curve and a proprietary model for prepayments).

In conclusion…

We hope you've enjoyed issue 2 of Quant*News*. Comments, submissions, and other requests should be sent to your editor at sjanowsky@thi.com.