TechHackers' QuantNews

November 25, 1998

Welcome to Quant*News*, Tech Hackers' new newsletter.

We hope to provide you with reports of the latest quantitative techniques being used company-wide. We hope to keep you informed about developments in the consulting practice as well as in software products.

EMU

Changes resulting from European Monetary Union (EMU[1]) will be the highlight of the autumn QuantTools / ***@nalyst*** release. The primary impact will be in the treatment of government bonds for the countries participating in EMU.

In addition to merging their various currencies, the participants in EMU are "harmonizing" their conventions for government bonds. This means that instead of having various calendar conventions (act / act, act / 365, 30E / 360 and variations) all bonds in EMU countries will calculate accrued interest using an act / act calendar. This means that if an interest payment of I is due at the end of a period of D days (D typically being 181, 182, 183, or 184 for semiannual bonds and 365 or 366 for annual bonds), then if we are d days into the period the interest accrued is calculated as dI/D. For act / 365, the denominator D is always fixed at 365 (annual bonds) or 182.5 (semiannual bonds); for 30E / 360 D is 360 or 180 and d is calculated under the assumption that all months have 30 days.

Unfortunately the harmonization is taking place in a very inharmonious way. For example, all Belgian bonds will switch to the new conventions as of January 1, 1999. German bonds, on the other hand, will switch to the new conventions as of the first coupon payment in 1999. Some Finnish bonds will switch, and others won't. Officially Spanish bonds have already switched over, but most quotes still refer to the old methodology.

How does one compare the yield on two bonds which use different conventions? First of all, one needs to use the proper accrued interest rule for each bond. Accrued interest is just a convention for splitting the actual market price of a bond into two pieces, the quoted price and the accrued interest, and it is the actual market price (or the bond's net present value) that needs to be used (at some point) to compute the yield. Then one needs to use only one convention (i.e. one method for measuring elapsed time) for discounting in computing the yield.

For example: If we want to compare the

[1] contrary to popular belief, EMU is not an acronym for "Extra Money for Us."

yield of two German bonds, one still trading according to 30E / 360 conventions and one trading act / act, we could compute the yield of the second bond according to "pure" act / act conventions, and calculate the yield of the first bond by using 30E / 360 for accrued interest and act / act for discounting.

This can result in a substantial effect. If both bonds mature on 1–March–2000, pay a 6% coupon and have a yield of 6%, on 28–February–1999 the act / act bond would be computed to have a price of 99.999518, the 30 / 360 bond would be computed to have a price of 99.998542, and the mixed computation gives the answer 100.033079.

Cliquet Options

We will be adding to our exotic options coverage by supporting *cliquet options*. A cliquet option is an option with a lock-in feature: on a prespecified series of dates, any intrinsic value that the option has reached is guaranteed, regardless of what happens to the asset price after that. In particular, if the lock-in dates correspond to the dates on which you must report earnings, once you've made money you can't then lose it the next quarter.

Suppose a cliquet option has lock-in dates t_1, t_2, ..., t_n, and expiration date t_{exp}. If the asset price on each of these dates is S_{t_i}, the S_{t_i} are each random variables where we know the correlations: the correlation of S_{t_i} and S_{t_j}, $j \geq i$ is $\sqrt{t_i/t_j}$. The we just need to integrate the payoff function,

$$\max\left(\max_i[(S_{t_i} - \text{strike})], 0\right)$$

(for a call) against the bivariate lognormal density, and then discount back to the present time—this is the same procedure as for rainbow ("best of") options.

QuantTools / ***@nalyst*** will support cliquet options with a single lock-in date.

Quasi-Random Numbers

guest column by Maurizio Mondello

My most recent work has involved testing the performance of several quasi-random number generators to be used in place of (pseudo-)random numbers for integration by Monte Carlo sampling. The resulting procedure is known as Quasi-Monte Carlo.

Monte Carlo provides a very general and robust simulation methodology to value financial instruments whose price can be expressed as a (multidimensional) integral. The goal of the Quasi-Monte Carlo methods is to improve the relatively slow rate of convergence of Monte Carlo integration. Implementing these methods in an existing (Monte-Carlo-based) system is fairly straightforward. However, care must be taken of the fact that to perform an n-dimensional integral the quasi-random numbers need to be generated and used as n-dimensional vectors (points of an n-dimensional space), while there is no intrinsic notion of dimensionality in a sequence of (pseudo-)random numbers.

More difficult is to provide reliable sampling errors, because the standard statistical error estimate used in Monte Carlo is based on the fact that sampling points are independent, which, by construction, is not true of quasi-random points. The error in Monte Carlo integration is directly proportional to the total variance of the integral being considered and inversely proportional to the square root of the number of sampling points. It is often advantageous, therefore, to apply variance reduction techniques, even at the cost of some computational overhead. This is true also for Quasi-

Monte Carlo methods, but here it can also make sense to use variance "reshuffling" techniques, if they lead to a concentration of the variance (error) on a smaller number of dimensions. This is because Quasi-Monte Carlo is comparatively more effective at reducing lower-dimensional fluctuation. An example of such variance "reshuffling" is the generation of a sample path (in time) using the Brownian Bridge technique, rather than by incremental steps. The points (steps) in the path could correspond, e.g., to the reset dates of a swap and their total number gives the dimensionality of the associated pricing integral.

In our tests, we have found that this "reshuffling" technique, often referred to as dimensional reduction, can be effective, particularly for paths involving 20 intermediate points or more.

editor's note: For a more comprehensive overview of Quasi-Monte Carlo methods, see S. Galanti and A. Jung, "Low Discrepancy Sequences: Monte Carlo Simulation of Option Prices," *Journal of Derivatives* **5**, no. 1 (Fall 1997), pp. 63–83.

Speed Corner

Your code is computing a 47-dimensional integral which needs to repeatedly compute `exp` and `log` — so that must be where the speed bottleneck lies—
not necessarily...

Some recent profiling determined that some applications we were developing spent the overwhelming majority of their time doing integer division in order to convert dates in the form of month, day and year to serial numbers and back. This is a particularly slow operation on older SPARCstations, which don't support integer division in hardware.

We were able to speed up these apps by almost a factor of four by a combination of

- caching results, and
- changing compilation flags to support hardware integer division on newer machines.

We think we can increase speed further by storing the year as two digits instead of four[2].

Stable Value Products

guest column by Peter Greenberg

Since July, Ken Chan, Vivek Nathan, and I have been working with Rabobank International' Stable Value Product (SVP) business. Rabo sells specialized financial products and services to stable value funds, which are increasingly popular investment options within many companies' pension and 401(k) plans (TechHackers' own 401(k) plan includes an SVP fund which is one of Rabo's big clients). These funds combine the capital stability associated with money market funds with rates of return closer to those found in intermediate bond funds. While SVPs vary quite a bit in structure, the basic concepts are fairly straightforward and the mathmatics quite simple.

Stable Value Funds hold almost all of their assets in SVPs, but they also hold a small amount of cash. The prototypical SVP form is the Guaranteed Investment Contract (GIC) sold by an insurance company. GICs strongly resemble the Certificates of Deposits sold by banks. At contract origination, the investor pays some initial amount (*Initial Contract*

[2] We are porting QuantTools to a *new* system; we are required to pass a major and minor version number. Each is limited to a single digit. We call this the V2D problem.

Value) to the issuer. Each day interest is credited and compounded daily at some rate of interest (*Crediting Rate*) stated in the contract, and added to the *Contract Value*. Maturity is usually in 3 to 5 years, at which time the issuer is obligated to pay the Contract Value to the investor. Under certain (usually quite limited) circumstances, the investor can demand early termination of the contract, at which time the issuer pays the Contract Value to the investor. Some contracts call for partial withdrawals or additional deposits. The basic equations are:

Today's Beginning Contract Value = Yesterday's Closing Contract Value + Deposits − Withdrawals

Today's Closing Contract Value = Today's Beginning Contract Value × (1 + Crediting Rate)$^{1/365}$.

In this context, the Initial Contract Value is a deposit on the first day of the contract, and termination or maturity is a withdrawal on the relevant day.

The insurance company generally invests the proceeds from GIC sales in a portfolio of fixed income investments of similar maturity (of course they, too, hold some small amount of cash). They tend to buy and hold these investments, which tend towards the more illiquid corporate bonds and ABS and MBS tranches of high credit quality. An investor that holds to maturity is always guaranteed to come out ahead, at least if you discount the possibilty of default and the effects of inflation. You may recall that a bond's yield is a rate of return assuming the bond is purchased at a certain price and held to maturity. So long as the average yield-at- purchase (converted to a daily-compounded rate) of bonds in the portfolio exceeds the average crediting rate on the GICs that it sells, the insurance company will make a profit on those GICs.

The risk to the insurance company is that it may be forced to sell bonds prior to maturity. This might happen if the insurance company's incoming cashflows (including insurance premiums, GIC sales, and payments of interest and principal on their investments) exceed their outgoing cashflows (benefit payments and GIC maturities and redemptions). In this case, the cash balance held by the insurance company will begin to decline, and finally the insurance company's ability to pay its creditors (policyholders and GIC investors) will be impaired. It will have to sell its "buy and hold" assets to meet its short-term obligations. If it sells bonds for less than their purchase price, the difference is absorbed by the company's shareholders. Thus, a company's ability to profit from the sale of GICs is related to its ability to predict its cashflows.

Stable Value Funds only buy GICs that contain provisions which ensure their liquidity, and allow it to post gains for its investors every day. Like other types of funds, stable value funds determine their Net Asset Value (NAV) per share each day by adding up the value of all their assets, subtracting any liabilities, and dividing by the number of outstanding shares. They value each of their SVP assets at contract value, and then add the small cash balance that they hold. This method ensures that NAV increases each day, since the value of each SVP asset also increases each day by its crediting rate. Investors can purchase or redeem their shares at NAV on any business day, and this activity, along with purchases and maturities of SVP instruments and other factors, will effect the amount in the fund's cash account. An unexpectedly high number of share redemptions not matched by purchases might draw down the cash balance to the point that the fund will need to liquidate its SVP investments to meet its redemption obligations. Fortunately, the GICs purchased by SVP funds allow early termination at contract value in that case, which is why accountants allow them to value their GICs at contract value.

Recent years have seen the invention of the *synthetic GIC*, which combines an ordinary bond, much like those found in insurance company portfolios, with a *liquidity wrap*. Like a real GIC, this investment also has a contract value that grows each day according to a crediting rate. The liquidity wrap, provided by a bank (like Rabo), obligates the issuer to reimburse the fund for any difference between the contract value and the sale price of the bond, provided that the sale was made necessary by a large number of redemptions unmatched by purchases (a "negative participant cashflow"). In return, the wrap provider is entitled to a regular fee that is based on the contract value and expressed in basis points. The wrap provider bears the risk that the fund's *participant cashflow*, defined as the dollar amount of share purchases minus share redemptions, will turn sharply negative. At the time the wrap is exercised, the wrap provider will only lose money if the bond's market value is below its contract value. In some cases, in fact, the wrap provider is allowed to benefit by the amount that the market value is above the contract value! As a special case, the wrap provider is obligated to pay any remaining contract value on the day that the bond pays off entirely, this being treated as a sale for the purposes of the wrap.

The equations governing synthetic GICs are similar to those for GICs, except that provision is made for the wrap fee, and bond cashflows are essentially treated as withdrawals:

Invoice Amount = total purchase price of underlying bond, including purchase principal and accrued interest, in dollars.

Gross Crediting Rate = The single daily-compounded rate that discounts bond cashflows to the Invoice Amount.

Wrap Fee = some agreed-upon constant expreseed in basis points (daily compounded). Averages about 15bp.

Net Crediting Rate = Gross Crediting Rate − Wrap Fee

Of course, some bonds, like those backed by mortgages, have uncertain future cashflows. In that case, some assumption must be made about future rates of prepayment. The resultant assumed cashflows are discounted to obtain the purchase price.

On the first day of the contract:

Initial Contract Value = invoice amount for bond purchase, including principal and accrued interest, in dollars.

Today's Beginning Gross CV = Today's Beginning Net Contract Value = Initial Contract Value

Subsequently,

Today's Beginning Gross CV = Yesterday's Closing Gross CV − Today's Bond Cashflow

Today's Beginning Net CV = Yesterday's Closing Net CV − (Today's Bond Cashflow − Payment of Yesterday's Closing Accrued Wrap Fee)

On every day, from the first to termination,

Today's Closing Net CV = Today's Opening Net CV $\times$ ((1 + Net Crediting Rate)$^{1/365}$)

Today's Closing Gross CV = Today's Opening Gross CV $\times$ ((1 + Gross Crediting Rate)$^{1/365}$) − Payment of Today's Beginning Accrued Wrap Fee

Today's Closing Accrued Wrap Fee = Today's Closing Gross CV − Today's Closing Net CV

The idea is that the synthetic GIC asset grows essentially at the gross crediting rate, and this is reflected in the Gross CV. However, a liability called *Accrued Wrap Fee* accrues in the background. At any given time, that liability is defined to be the difference between the Gross CV and the Net CV, which accrues at the lower Net Crediting Rate. Stated another way, the Net CV is just the Gross CV asset net of the wrap fee liability. Periodically (often on bond cashflow dates) the Gross CV asset is basically paired off with the accrued wrap fee liability, meaning that the Gross CV is reduced by the payable, and the payable is reduced to zero. At this point, the Net CV and Gross CV are the same. Then they diverge until the next time the accrued wrap fee is paid by the fund. The Net CV is the value of the investment to the fund, and the minimum liquidation value of the bond guaranteed by the wrap provider; the Gross CV exists primarily to compute the wrap fee amount.

There are three major risks associated with issuing a wrap contract:

1. **Participant Cashflow Risk**: if the fund to which the contract is issued has strong negative cashflows, it is more likely that the contract will be "hit", or exercised.

2. **Price Risk**: this is the risk that the market value of the wrapped security will decline, increasing the possibility of loss if the wrap is indeed hit when participant cashflows turn negative.

3. **Prepayment Risk**: not as easy to discern as the prior two risk factors, this is the risk that prepayments will differ significantly from the rates that were used to set the crediting rates. If this happens, Net Contract Value will stray significantly from market price. If this is change is in the "wrong" direction, that is, if a discount security lengthens or a premium security shortens, then Net Contract Value will tend to rise above market price. As with price risk, this increases the chance of a big hit if participant cashflows turn negative.

In conclusion...

We hope you've found our first issue of Quant*News* to be both interesting and useful. Comments, submissions (*hint hint*), and other requests should be sent to your editor at sjanowsky@thi.com.